White Paper

Industry
Manufacturing

Boon Logic's Amber Transforms IIoT Condition Monitoring Training and Inference

Based on Boon Logic's Nano, the unsupervised machine learning anomaly detection solution trains itself without data scientists, runs 1000x faster than current methods, and scales easily—only on Intel® processors with Intel® SGX

Authors

Brian Turnquist
PhD
CTO Boon Logic

Executive Summary

AI-powered anomaly detection is changing how companies and business operations achieve higher yield and operational efficiency while providing better quality in their products. But not all AI methods are effective for detecting anomalies in the Industrial Internet of Things (IIoT).

Much of today's AI methods are designed around supervised machine learning, where anomalies, rather than a normal operating state, become part of the model. Amber from Boon Logic is a high-speed, highly accurate anomaly detection solution that uses unsupervised machine learning to learn a machine's or process' normal states of operation. During inference, anomalies are any condition that does not match normal conditions.

Amber is 1000x faster than current unsupervised machine learning methods,[1] enabling it to train and inference at the edge of the network. In many deployments, it has shown to train in minutes rather than weeks or months—allowing it to continuously train if needed—and has predicted failure events months in advance of existing means.

Amber is powered by Boon Logic's Nano algorithm. The solution runs exclusively on Intel® processors with Intel® Software Guard Extensions (Intel® SGX), without accelerators—outperforming accelerators such as GPUs and FPGAs.

This white paper describes Amber and how it can perform anomaly detection 1000x faster than competing methods, why Intel processors are used, presents speed and accuracy benchmarks, and offers case study examples of real-world deployments and the value achieved by customers.

Contents

Automating Quality Compliance and Reliability Through Anomaly Detection

The Industrial Internet of Things (IIoT) uses instrumentation in machinery and processes. Instrumenting for anomaly detection helps ensure quality and regulatory compliance; increase yield, efficiencies, and reliability; and reduce cost and losses. An anomaly is simply something that isn't "normal" within the context of a set of features, e.g., engine temperature, rpm, fuel mixture, vibration, etc. Any feature that is considered important and monitorable could lead to detecting an anomaly and enable the benefits IIoT brings to business operations.

Anomalies Provide Actionable Data

The more you know about an anomaly, the better you can address problems early. Anomalies can be differentiated in two ways—probability and distance—relative to what is considered normal for the features being monitored. These two characteristics provide a way to measure anomalies.

- Probability anomalies are very rare events.
- Distance anomalies are varyingly yet importantly different events from normal.

Solutions that can measure anomaly distance and probability offer actionable information that can lead to early detection of potential problems before they become critical.

Anomalies occur unexpectedly, can be idiosyncratic, and are often unique to each asset. For example, a set of elevators, while designed by the same manufacturer, may present slightly different motor currents for various reasons, some indicative of an upcoming failure mode. Therefore, sufficient anomaly detection that can offer actionable data must go beyond statistical approaches and subjective threshold settings for all assets. The approach must account for what is considered to be "normal" within a given dataset from each asset. That dataset can become very complex, especially as monitored features scale up. Therefore, anomaly detection solutions are increasingly automated with artificial intelligence (AI) methodologies. But not all AI solutions are alike, and some are not the best method for effective anomaly detection in IIoT.

Supervised Machine Learning Is Not Always the Best AI-Based Anomaly Detection

Today's anomaly detection solutions use telemetry from sensors, imagery, and aural monitoring, feeding a model trained to look for abnormalities. AI Models are typically trained using neural networks and supervised machine learning. Thus, failure mode data is collected and labeled by data scientists and ingested by the network, resulting in a trained model. During real-time inferencing, the model classifies anomalies as they occur.

Supervised machine learning assumes a supervisor has seen these anomalous events in the past, have accurately characterized them with the correct feature sets, and have an abundance of references—labeled data—to train the model. This means events previously unseen or infrequently seen will be difficult—or impossible—to identify in the future.

Importantly, these methods don't measure an anomaly's probability and distance.

Supervised machine learning is expensive. The approach requires a combination of skills difficult to find in industry. Neural networks depend on curated, labeled data, requiring expertise in data science and the particular domain to effectively build models. But each asset's feature data can be unique to the system, which makes scaling one-off models to thousands of assets impractical. Additionally, GPUs are often used—sometimes required—to train in a reasonable amount of time, adding cost. Most importantly, supervised learning and neural networks aren't always the best approach.

Amber AI Anomaly Detection from Boon Logic

Finding anomalies is straightforward if a high-dimensional model can be trained for normal. Amber is an automated system that detects anomalies in complex assets and environments by autonomously training for normal instead of abnormal using unsupervised machine learning.

Amber has been designed, and has proven in the field, to detect equipment noncompliance earlier than other predictive analytics tools, giving operators more time to schedule maintenance, order parts, or change production settings.[2]

Starting from normal is the only way to have insights you can truly trust.

Amber is designed for complex machinery and their unique datasets. It trains in real time from data collected during normal operations instead of relying on failure modes and programmed logic. Amber's auto-learning algorithm trains itself, enabling deployment in minutes, not weeks. During training, Amber autonomously learns hundreds of relationships between features selected during configuration. As training progresses, Amber's learning curve starts to level, indicating that Amber is becoming familiar with the asset's normal operations. Once the learning curve plateaus, training is complete.[3]

During normal operations, Amber inferences telemetry data in real time. The solution measures anomalies and presents their significance within the users' familiar operating environment through compliance score and feature significance interfaces.

Amber's Compliance Score (Figure 1) is based on an individualized, high-dimensional, unsupervised machine learning model of an asset. The score indicates ongoing asset health and deviations from normal operating behavior.

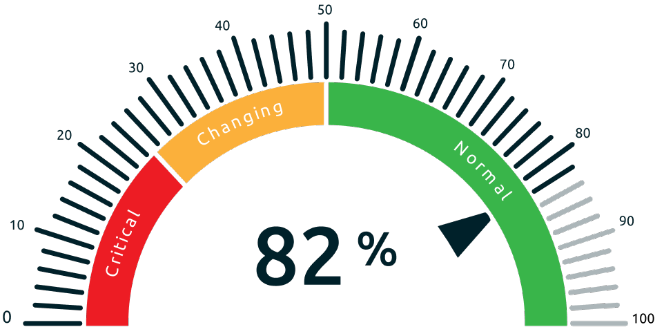

Figure 1. Typical Compliance Score readout.

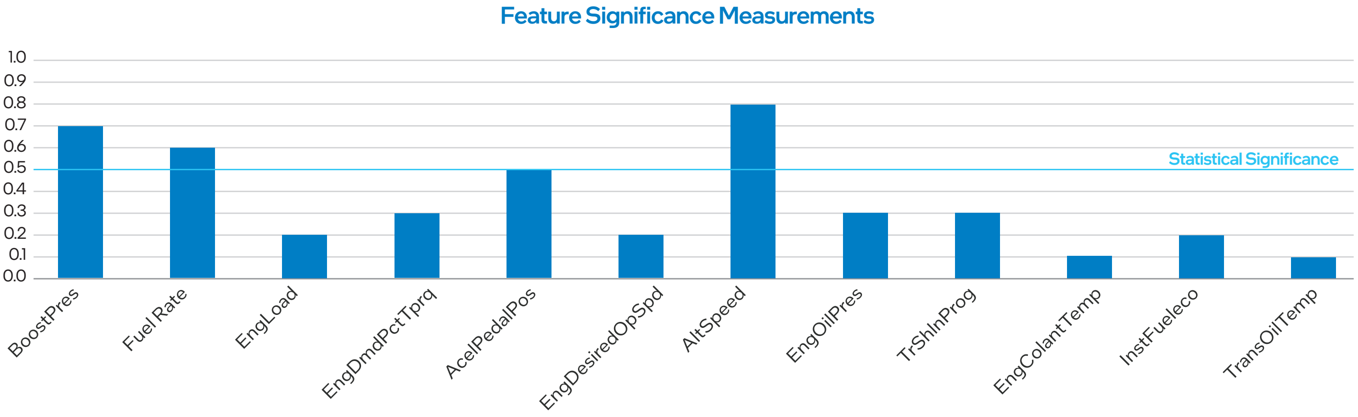

Figure 2. Feature significance indicates each feature's impact on the system.

Feature Significance measurements (Figure 2) shed light on which features are suspect in causing a drop in Compliance Score. Being able to swiftly pinpoint the origin of the problem allows for prompt and cost-effective issue resolution.

Amber Trains With Unsupervised Machine Learning

Unsupervised machine learning, as the name implies, is an algorithm-first approach that assesses and organizes data prior to a human reviewing the results. Amber does not require labeled data to learn what is normal for the asset. It analyzes the data on its own.

Amber's unsupervised machine learning is based on a proprietary algorithm called Nano developed by Boon Logic. Using Nano, Amber quickly and extremely efficiently builds models for each asset or process from each feature's data stream. Each model defines what Amber learns as normal for the asset, making anomaly detection fast, easy, and measurable.

Amber is an n-dimensional segmentation technology that gives similar results to k-means. But it is 1000x faster than k-means,[1] easily scales with the number of features, auto-tunes hyperparameters, and is designed for complex data sets. Thus, it does not need domain or data science expertise. That makes it cheaper to build a model, easier to implement because a reliability or process engineer can create the model from any appropriate data stream, and faster to deploy.

Critically, Amber does not require GPUs. It leverages the performance and technologies of Intel® Xeon® processors with Intel SGX for training and inferencing.

The Boon Nano

Nano from Boon Logic is a next-generation unsupervised machine learning algorithm. The Nano is an n-dimensional segmentation (clustering) algorithm. Clustering is a data science method applied when categories in the data set don't exist, and the scientist does not have very much knowledge of it—perfect for IIoT applications.

Segmentation is Complex in IIoT

Segmentation is often illustrated in a simplistic manner with few clusters. But, from a data science perspective and IIoT application it can quickly become quite complex. In real-world, semi-structured IIoT data, many data sources are streamed rather than batched. Data sources are idiosyncratic rather than universally represented. And streamed data creates temporal relationships between the features being monitored. Table 1 lists several domains where segmentation becomes complex due to these data characteristics. This time relationship is what transforms segmentation into multivariate time-series analysis.

Domain	Source	Sample Count	Dimensionality	Cluster Count	Separation
Medical Biomarkers	Multi-electrode EEG and ECG	Thousands per second per electrode	12 dimensions	Hundreds	Poor (e.g., at seizure onset)
Medical Biomarkers	Volumetric imaging (MRI, CT, Ultrasound)	Tens of millions per scan	Up to 1024 dimensions	Hundreds to thousands	Poor (e.g., at tissue gradients)
Manufacturing	Asset sensor telemetry (vibration, temperature, etc.)	Up to thousands per second per sensor	5 to 100 dimensions	Hundreds	Fair, depending on the asset operational modes
Network Monitoring and Security	Ethernet packets or network flow measurements	Millions per second	10 to 50 dimensions	Hundreds	Fair
Image Processing	N x M images, either black and white or color	Millions per image	Up to thousands depending on segmentation technique	Thousands for video data	Poor due to natural lighting gradients

Table 1. Examples of domains and their characteristics that make segmentation difficult and complex.

The Boon Nano is designed for these complexities and time relationship to quickly and efficiently build accurate models from unlabeled data in real time.

k-Means and The Nano

k-means is a common clustering algorithm that has been implemented in various applications since the 1970s. It is an effective algorithm for simpler datasets. But finding the right k (number of clusters) for a complex dataset is time-consuming. And increasing the number of features or increasing the value of k significantly decreases performance of the algorithm and increases the training time, thus k-means lacks scalability as feature count and model complexity expand.

The Boon Nano automatically tunes for the number of clusters as it analyzes the dataset. It scales easily as the dataset becomes more complex. And, it has shown to run up to 1000x faster than other solutions, such as k-means, supervised learning with neural networks, and autoencoders—while providing high accuracy.[1]

For more details about The Boon Nano, please read the white papers, Nano™ by Boon Logic and Nano for Data Scientists.

Amber Runs Only on Intel® Xeon® Processors

Amber is designed to run only on Intel processors with Intel SGX. And it runs best on these CPUs—without additional accelerators. GPUs are not only unnecessary, they do not offer the same performance benefits as Intel Xeon processors. Segmentation of large datasets needs to have a global view of the data. GPUs, while having many cores, cannot see the overall data with their limited memory. In benchmarks, Amber on one core of an Intel processor outperformed k-means running on 3500 cores of a GPU by 40 percent, the equivalent of 4900x faster per compute core.[4]

FPGAs are often used to accelerate certain operations, such as anomaly detection. But testing Amber on an FPGA resulted in only a 6x speedup compared to k-means versus 1000x speedup on Intel processors.[5]

Performance Advantages of Intel Xeon Processors

4th Gen Intel® Xeon® processors are the latest generation of Intel data center processors. This generation of Intel® CPUs offers important performance enhancements designed to accelerate AI, data streaming, and analytics.

Intel Xeon processors feature the broadest and widest set of built-in accelerator engines for today's most demanding workloads. Whether on-prem, in the cloud, or at the edge, Intel® Accelerator Engines can increase application performance, reducing costs and improving power efficiency.

Amber is embarrassingly parallel. 4th Gen Intel Xeon processors are highly scalable with up to 60 cores per socket in the Intel® Xeon® Platinum 8490H Processor. Thus, as feature sets and data expand, Amber performance scales linearly on Intel CPUs—you just spin up more cores. However, Amber runs excellent on fewer cores as shown in Boon Logic benchmarks using Intel CPUs.[1]

Built-In Security with Intel® SGX

Amber must be run on CPUs with Intel SGX in order to protect sensitive data, trained models, and code.

Intel SGX Overview

Intel SGX provides a set of software extensions to the Intel CPU instruction set. Intel SGX offers hardware-based protection that isolates specific application code and data in memory. It allows user-level code to allocate private regions of memory, called secure enclaves, which are designed to be protected from processes running at higher privilege levels. These enclaves can be designed to protect sensitive client data, models, codes, and intellectual property (IP)—such as the Boon Nano—from tampering, attack, and readout.

Amber Protects Models and Data with Intel SGX

Boon Logic developers designed specific Intel SGX code within Amber to create secure enclaves. The enclaves help guard against attack and corruption of data and models.

- **Protecting customer data** — Amber is built to run in the most data-sensitive environments, such as defense and national security. Using Intel SGX secure enclaves allows deploying the technology at the very edge of the network—from shipboard at sea to enterprise edge environments—with the highest security against data exfiltration. Only the authorized application code can access the data in the secure enclaves.
- **Protection against model poisoning** — Model poisoning is a covert method of contaminating AI. Poisoning occurs when an adversary feeds data into a model to trick the model into thinking features are presented as "normal" when they are not or "abnormal" when they are compliant. Intel SGX secures the model in an enclave and helps protect the model against poisoning.

Deploying Amber

Amber can be used in manufacturing, heavy industrials, oil and gas, and anywhere sensor telemetry is being collected or produced. Amber is fast enough to continuously train and inference at the edge—at the site of data generation and accumulation. However, some industries don't want or allow continuous training. Amber is flexible for both scenarios.

Built as a microservice, Amber easily plugs into any IIoT platform, such as AVEVA PI, Ignition, Cumulocity, Rockwell's Factory Talk, and Azure IoT Hub (Figure 3). Using supported APIs and messaging protocols (Table 2), Amber integrates with an existing customer database and visualization engine to give the user an intuitive look and feel, utilizing an already familiar workflow. Amber integrates powerful, innovative AI without the headache and cost of developing an environment around it.

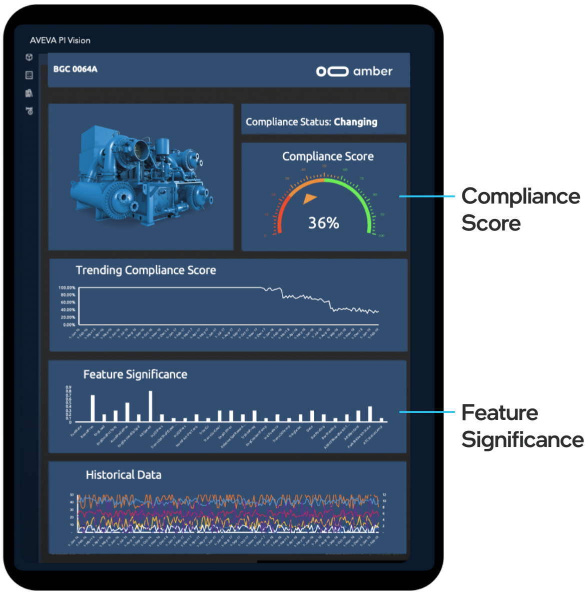

Compliance Score

Feature Significance

Figure 3. Example Amber dashboard in AVEVA PI Vision.

Programming Interfaces	Messaging Protocols
▪ REST API ▪ Python SDK ▪ JavaScript SDK ▪ C++ SDK ▪ Go SDK ▪ R SDK ▪ C# SDK ▪ Java SDK	▪ MQTT ▪ OPC UA ▪ MODBUS ▪ SQL

Table 2. Amber-supported APIs and protocols.

Benchmarks

Boon Logic performed two benchmarks with The Nano to evaluate it against k-means.

KDD Cup 1999—Nano is 1000x Faster than k-Means

KDD Cup 1999 is a well-established network intrusion detection benchmark. Its dataset allows evaluating both accuracy and segmentation speed. With 4.9 million rows, each one has 32 features, representing a network connection. The benchmark tests how well a solution can differentiate benign from malicious connections.

Boon Logic ran this benchmark with results illustrated in Figure 4. It shows the average time (in microseconds) to assign a cluster ID to each row of the dataset as the number of clusters increases. For high accuracy, the number of clusters increases in order to achieve sufficient separation to distinguish between benign and malicious traffic. The computational complexity of the two algorithms is linear versus the number of clusters, but the complexity growth constant of k-means is 1000x larger than the Nano. In deployment, the Nano processes 1000 rows in the time for k-means to process one.[1]

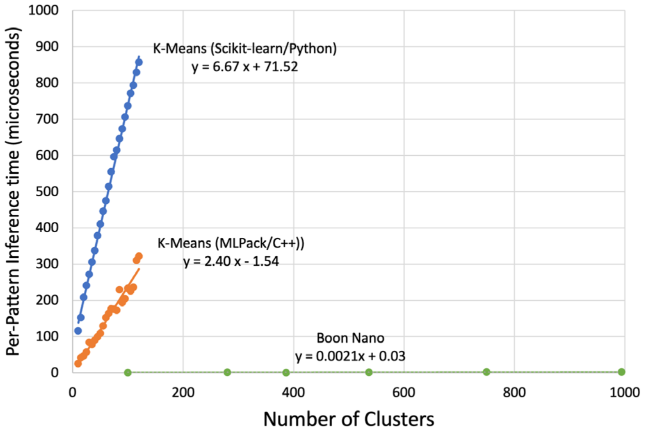

Figure 4. KDD Cup 1999 benchmark of Nano vs. k-means on a single Intel® processor core (courtesy of Boon Logic).[1]

Internal Accuracy Benchmark—99 Percent Accuracy for Nano vs. 85.4 Percent for k-Means

This benchmark was generated internally by Boon Logic to evaluate accuracy of the two algorithms. The data set contains 100,000 25-dimensional vectors that are separated by ground truth into 100 clusters, ranging in size from 1 to 3899. Datasets like this can be generated using Python, MATLAB, and other tools. The dataset is also available from Boon Logic for self-evaluation.

It is important to note that for this benchmark, the target cluster count (k) was set to 100. k-means would initialize to have 100 centroids, which produced 100 clusters for segmentation. For real-world operations, the value of k is typically unknown, requiring an "elbow" procedure to find the value of k. This results in a computationally slow initial step prior to running the algorithm. The Nano autonomously tunes hyperparameters, automatically finding the optimal number of clusters. Table 3 shows the time measurements for this benchmark.

	The Nano	**k-Means**
Autoconfiguration Duration	3.7 seconds	Manual configuration required
Total Clustering Time	0.12 seconds	55.7 seconds
Per-vector Inference Time	1.2 microseconds	557 microseconds

Table 3. Accuracy benchmark speed measurements.

The details of the resulting accuracy measurements are beyond this paper. For those details, refer to the white paper Nano™ for Data Scientists. The following summarizes the accuracy findings of The Nano versus k-means.

	The Nano	**k-Means**
Overall Accuracy	99.9%	85.4%
Precision	100%	82.8%
Overall Recall	99.8%	88.2%

Table 4. Accuracy findings.

The Nano, with 1000x faster clustering than k-means and significantly higher accuracy, makes Amber a powerful anomaly detection engine for IIoT.

Case Studies

Boon Logic has deployed Amber across many industries with success in identifying anomalies faster and sooner than existing or competing solutions. A few case studies are described below. For more details, visit Boon Logic's website or contact Boon Logic.

Cost Avoidance: Specialty Gas Manufacturer Saves Up to $800,000 From Compressor Failure

Critical components in a manufacturing process can be expensive to repair and can shut down production for extended periods if unexpected failure occurs. In 2021, AIONT, an industrial pump and compressor service company, created a condition monitoring solution for a specialty gas manufacturer's Atlas Copco centrifugal compressors. A compressor supplies gas to a large semiconductor manufacturer in Taiwan. Any unplanned downtime would mean hundreds of thousands of dollars in lost revenue. The solution monitored X, Y, and Z axis accelerometer sensors on each stage of the compressor. Vibration thresholds for the monitoring system were set to indicate issues in the component.

AIONT wanted additional monitoring, so they added Amber to their existing system. Amber was connected to the compressor's sensor data streams and trained on vibration conditions. After a month of live training, Amber began inferencing the data.

For a month and a half, the Compliance Score for one of the compressors decreased, finally falling below 40 percent, indicating the compressor was in a new, never seen state. Yet, the existing condition monitoring system did not indicate any issues.

Two weeks later, the Compliance Score dropped to 24 percent, signifying a critical condition with the compressor. The associated Feature Significance values were evaluated, showing an issue with Stage 3 of the compressor. Upon inspection, maintenance staff found small cracks in the compressor's cooling bundle. Further investigation revealed several tears in the cooling fin.

The cooling bundle cost $60,000 to replace, however, had the condition gone undetected, replacement of the core unit would have resulted in over $800,000 in repairs and months of downtime for the semiconductor manufacturer.

Download the full case study here.

Prescriptive Maintenance: Predicting and Diagnosing a Future Failure of a Diesel Engine Fuel Injector

Large diesel engines are used throughout heavy industry, such as mining, defense, and transportation. These engines operate under heavy loads and typically suffer from many unplanned failure events. Such events can shut down or constrict operations and result in large production losses.

A large German engineering company added Amber for evaluation, monitoring sensors built into a large diesel engine to test Amber's failure predictability. The evaluation resulted in the following from Amber:

- First alarm was generated six months before failure
- Root Cause Analysis indicated the failure was related to the fuel system

To launch the evaluation, the customer identified ten key features that were most relevant for monitoring. The deployment required less than one hour of a Boon Logic data scientist's time to configure the model. With five months of data (250,000 samples), Amber trained and built a high-dimensional model, representing all the complexity of the engine's normal operating state.

For over two years, Amber's Compliance Score across the feature set remained nominal (100 percent). See Figure 5. Six months prior to failure, the Compliance Score dropped by nearly 50 percent, issuing an anomaly warning. Feature Significance showed an issue with the fuel system. Over the next six months, the Compliance Score dropped further. Six months after the first indications and warning, a fuel injector failed.

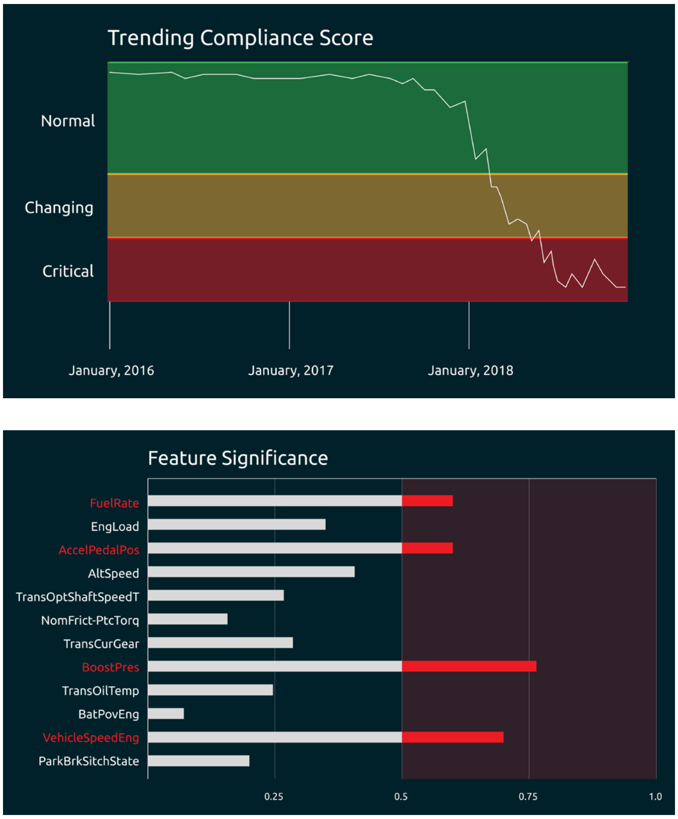

Figure 5. Compliance Score (upper chart) showed no issues with the engine for over two years. A warning was issued in March 2018. Feature Significance (lower chart) indicated an issue with the fuel system six months prior to injector failure.

For this engine application, Amber was deployed quickly, trained on 250,000 samples, and warned of an impending failure six months prior to occurrence.

Increased Reliability: Maximizing Uptime for the World's Largest Gold Mine

One of the world's largest gold mines in Indonesia operates a mature reliability program with good equipment instrumentation, offering valuable data about its equipment operations. Still, it suffered from alarm overload and occasionally missed failure events in its production lines. Amber was added to its monitoring processes for anomaly detection and failure prediction on a Ball Mill used in the production line. Amber was trained on two months of historical data (400,000 samples) across 32 features. The algorithm learned 67 clusters for normal operating states of the Ball Mill (Figure 6).

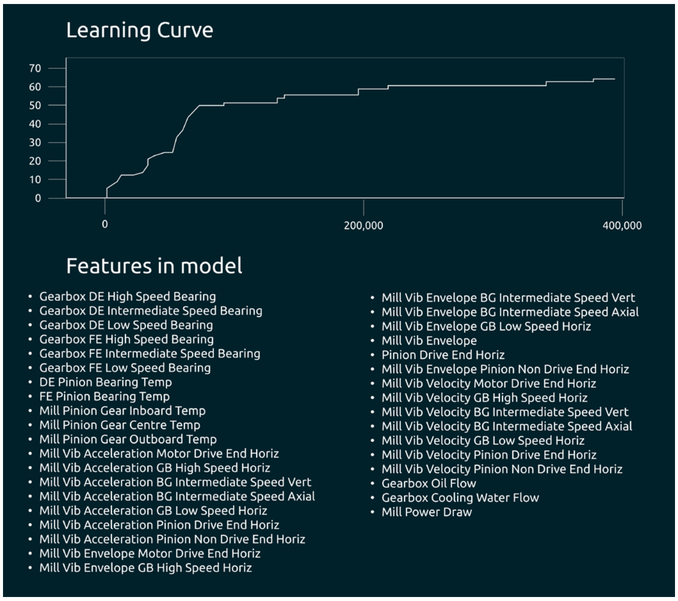

- Gearbox DE High Speed Bearing
- Gearbox DE Intermediate Speed Bearing
- Gearbox DE Low Speed Bearing
- Gearbox FE High Speed Bearing
- Gearbox FE Intermediate Speed Bearing
- Gearbox FE Low Speed Bearing
- DE Pinion Bearing Temp
- FE Pinion Bearing Temp
- Mill Pinion Gear Inboard Temp
- Mill Pinion Gear Centre Temp
- Mill Pinion Gear Outboard Temp
- Mill Vib Acceleration Motor Drive End Horiz
- Mill Vib Acceleration GB High Speed Horiz
- Mill Vib Acceleration BG Intermediate Speed Vert
- Mill Vib Acceleration BG Intermediate Speed Axial
- Mill Vib Acceleration GB Low Speed Horiz
- Mill Vib Acceleration Pinion Drive End Horiz
- Mill Vib Acceleration Pinion Non Drive End Horiz
- Mill Vib Envelope Motor Drive End Horiz
- Mill Vib Envelope GB High Speed Horiz
- Mill Vib Envelope BG Intermediate Speed Vert
- Mill Vib Envelope BG Intermediate Speed Axial
- Mill Vib Envelope GB Low Speed Horiz
- Mill Vib Envelope
- Pinion Drive End Horiz
- Mill Vib Envelope Pinion Non Drive End Horiz
- Mill Vib Velocity Motor Drive End Horiz
- Mill Vib Velocity GB High Speed Horiz
- Mill Vib Velocity BG Intermediate Speed Vert
- Mill Vib Velocity BG Intermediate Speed Axial
- Mill Vib Velocity GB Low Speed Horiz
- Mill Vib Velocity Pinion Drive End Horiz
- Mill Vib Velocity Pinion Non Drive End Horiz
- Gearbox Oil Flow
- Gearbox Cooling Water Flow
- Mill Power Draw

Figure 6. Amber training of the Ball Mill across 400,000 samples each of 32 features.

Using Amber resulted in the following:

- Statistically relevant detection of anomalies that indicated an issue with the high-speed shaft bearing six months before the shaft bearing needed replacement.
- Early warning of a failure six days before the onsite reliability teams detected it.
- Feature significance pointed to parameters located on the bearing of the high-speed shaft of a Ball Mill gearbox. This was confirmed to be the issue.

Figure 7 charts the timeline of Amber warnings. It issued an alarm six days prior to the monitoring team's detection of an issue.

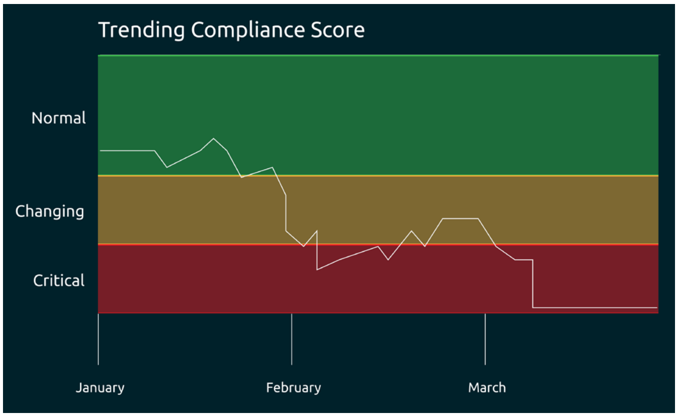

Figure 7. Amber issued early anomaly warnings for the Ball Mill, with root cause analysis showing issues with the high-speed shaft bearing.

Root cause analysis ranked the issues as follows:

1. Gearbox DE high-speed bearing
2. Horizontal vibration acceleration of the high-speed shaft
3. Horizontal vibration acceleration of the pinion drive end
4. Horizontal vibration envelope of the high-speed shaft

Amber was quick to deploy and resulted in early detection of anomalies and alarm conditions for the gold mine operations.

Keeping Fleets on Mission: Fairbanks Morse Defense

Fairbanks Morse Defense (FMD) was founded more than 100 years ago and has built a reputation as the world's premier supplier of leading marine technologies and naval equipment to the U.S. Navy, Coast Guard, and Military Sealift Command. Among FMD products, it manufactures massive shipboard diesel engines for U.S. naval fleets. These engines are critical to the operation and fulfillment of naval missions.

While built to the highest level of reliability, crews need to be able to respond sooner than later to potential problems that require maintenance or service—whether in port or at sea. Engine monitoring systems from FMD and the U.S. Navy provide reporting and alerts to potential issues.

With shipboard condition monitoring using AI, machine learning training of a model and inference must be done 100 percent on-premises. And because this is a highly sensitive environment, all computing must be done with robust security. FMD deploys Amber on ocean-going naval vessels for condition monitoring training and inference. Amber uses only Intel CPUs with Intel SGX to provide the highest levels of hardware-based security for analytics modules.

In several instances, Amber provided early warning through its Compliance Scores and alerted crews to anomalies with its Feature Significance ratings sooner than existing systems.

Summary and Call to Action

Amber enables fast and accurate anomaly detection using unsupervised machine learning based on the Boon Logic Nano algorithm. Amber uses Nano to learn what is the normal state of a process or machine, instead of learning what is an anomaly. Nano runs 1000x faster than the competing algorithm, k-means, and delivers higher accuracy—without accelerators—as shown by Boon Logic benchmarks.

Amber runs exclusively on Intel CPUs in the cloud or on-premises, providing both performance and security for data, code, and the customer's model. Amber is easily deployed, plugging into existing IIoT condition monitoring systems.

To learn more about Amber, visit the Boon Logic website or request a demo.

[1] See the white paper Nano™ for Data Scientists.

[2] See case studies in this paper for examples.

[3] See the sidebar on Boon Nano.

[4] Large scale K-means clustering using GPUs", Li, Frank, Pfahringer, Data Mining and Knowledge Discovery 2023 benchmarks a K-means segmentation of the KDD-Cup 1999 data set on a 3500 core GPU. We benchmark this same data set where we compared KDD-CUP 1999 with K-means vs Nano in the Nano for Data Scientists white paper.

[5] This "6x" benchmark is based on conversations we have had with engineers at Xilinx, not on any public benchmarks that they have provided. We could drop this particular sentence. We could say that Boon and Xilinx engineers working together implemented Amber on Alveo FPGAs, but it did not accelerate Amber beyond what is already possible on Intel CPUs.

Legal Notices and Disclaimers

Performance varies by use, configuration, and other factors. Learn more on the Performance Index site.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See configuration disclosure for details.

No product or component can be absolutely secure.

Your costs and results may vary.

Intel technologies may require enabled hardware, software, or service activation.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

© 2024 Intel Corporation. 0124/MT/HBD/PDF ♻ Please Recycle 358210-001US