

Machine Learning Finds Hidden Patterns for IIoT

By discovering patterns hidden in existing operational data, companies can significantly improve the performance, throughput, quality, and yield of industrial operations.

One approach to machine learning (ML) requires extensive training using captured data on high-powered servers. It allows for the creation of a deep neural network (DNN) that can then be deployed on systems requiring significantly less storage and performance. This method of training and identification is just one, but not the only way, to implement ML.

I talked with Nikunj Mehta, Falkonry CEO and Founder, about the company's technology that can be used for real-time analytics for the Industrial Internet of Things (IIoT).

Where do you see industrial operations and managed services evolve to in five years?

We're seeing two key trends in manufacturing and process operations that we expect to be commonplace within five years.

The first area is *analytics-driven digitized industrial operations*, where data from sensors, machines, and production systems isn't being collected to just enable a few, siloed use cases. Instead, the COOs and plant managers are taking an analytics-first approach to identify the next level

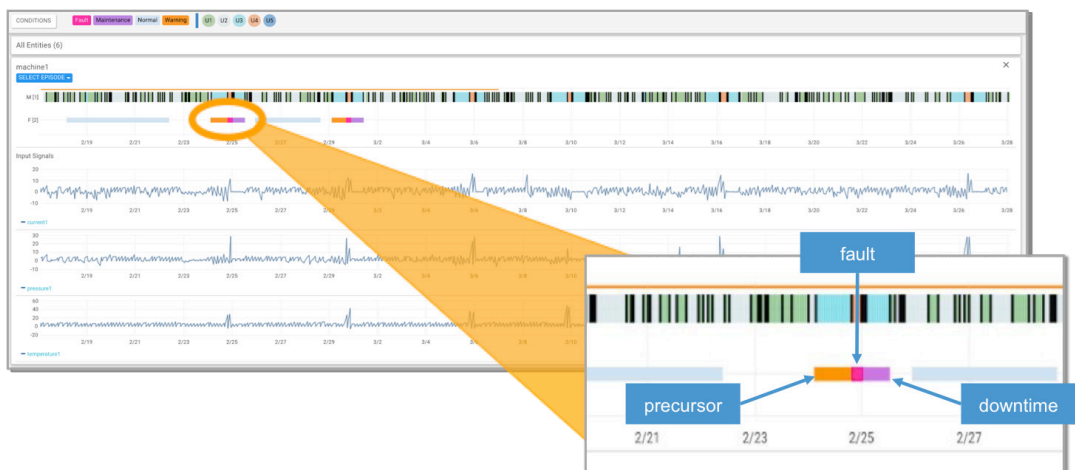
of improvement in production throughput, batch quality, and process safety. This aligns with where analysts such as McKinsey and Company are predicting the next 20% productivity rise in manufacturing will come from.

The second area is the revenue stream of traditional industrial equipment manufacturers evolving from a "product-driven" model of hardware with some managed services, to an "outcomes-driven" revenue model where "software + managed services" is a larger revenue contributor than product hardware. We have seen some projections where software + services will be more than 60% of revenue. Within that model, remotely managed services for assets deployed at a customer's site may be the largest recurring revenue contributor.

Falkonry has some heavy-hitter customers using its operational machine-learning system such as Toyota, Kawasaki, Ciner, and others. Tell me a little about your company and your technology.

Falkonry provides operational machine-learning and predictive analytics to Global 2000 industrial companies to help improve throughput, quality, and safety in their

1. Machine learning enables pattern discovery and early warning predictions.





Oil & Gas Operations

- Detect pre-shutdown patterns
- Real-time alerts for uncontrolled emissions



Mining & Metals Production

- Discover equipment downtime patterns
- Real-time production throughput adjustment



Power & Energy Operations

- Fault classification of power electronics
- Distributed asset monitoring



Automotive Manufacturing

- Detect deviations in discrete manufacturing
- Real-time quality estimation of welding



Semiconductor Manufacturing

- Predictive maintenance of equipment
- Optimize machine utilization



Chemical Manufacturing

- Quality estimation of batch process
- Monitor machine health

2. Machine learning is being deployed for a wide range of use cases and industries to improve operations.

operations.

Today, there's a tremendous amount of time series data that's generated by equipment and production systems in discrete manufacturing and continuous process operations. This data is rich in insights; however, most of it goes unutilized as it's too much data to be processed by traditional approaches. The insights in time series data comes from multivariate trends, also called patterns, which need to be learned from many signals across data sources.

Our "ready-to-use" machine-learning system, Falkonry LRS, helps improve operations by analyzing real-time data streams to discover patterns hidden in multivariate time series data, applying predictive analytics, and providing actionable insights in the form of alerts—without requiring data scientists (Fig. 1).

Falkonry LRS can be deployed both in the cloud, as well as on-premises, to meet the requirements of business-critical operations. It easily integrates with existing data collection and controls operations to deliver short time to value and low integration costs. In fact, many of our customers start seeing results in as little as three weeks.

And, finally, since the pattern discovery process is agnostic of industry use case, and the patterns discovered are unique to the operations being analyzed, Falkonry LRS has been successfully deployed across various industries and use cases (Fig. 2).

You claim your technology is a data scientist in a box? Can you explain this?

What's unique about our approach, and why we refer to it as a "Data Scientist in a Box" is that it can be used directly by industrial practitioners, such as operations engineers, process control engineers or manufacturing engineers, without requiring the assistance of data scientists.

We have built three key capabilities into the system to enable this:

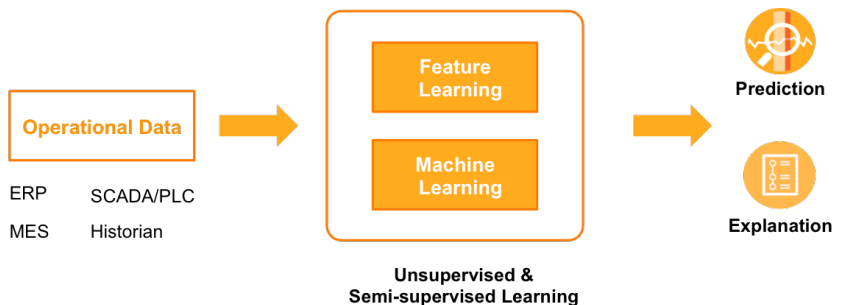
1. **Unsupervised Feature Learning** that autonomously discovers patterns in time series data, including those that would go undetected by human observation. Once the patterns have been discovered in the data, then machine learning can be applied.

2. **Machine Learning & Predictions** where the user selects a pattern that appears consistently prior to the event and labels that as a *precursor*. Based on the use case, the precursor could provide an early warning or prediction that ranges from several hours to several days ahead of the event occurring.

3. **Explanation** is key as most artificial-intelligence (AI) and deep-learning techniques don't provide an explanation as to how they arrived at their results. Our approach can determine which of the provided data signals were actually required to make the prediction, as well as how much each signal contributed to that result (Fig. 3).

Machine learning is hot right now. How do you differentiate from what other companies offer?

That's a good question. Our focus is on providing a software-based machine-learning system that easily integrates with existing data infrastructure and control loop. Unlike other approaches, there are no additional sensors, data management, or dashboards needed, so we don't disrupt existing workflow and can be up and running in a day and delivering results in as little as a few weeks.



3. Referred to as a "Data Scientist in a Box," Falkonry LRS is a self-service machine-learning system for engineers.

Also, with Falconry LRS, data stays within the environment and, since engineers are creating the models and using them, they retain full control over the data, models, and findings. Unlike other approaches, there's no need for data scientists or the vendor to handle the data.

And unlike other domain-specific machine-learning applications that apply to specific use cases like predictive maintenance, our approach applies across use cases and industries.

What markets are seeing the most adoption of machine learning and can you share any case studies showing real results of your system in some of these markets?

We're seeing strong adoption of machine learning in discrete manufacturing operations, specifically automotive, electronics, and semiconductor and industrial equipment, as well as in continuous process operations such as oil & gas, mining & metals, and chemicals.

In the case of discrete manufacturing, one of our customers builds motors for automotive applications. One of the failure modes requiring repair and replacement is burnout in the motor winding. They deployed Falconry LRS to predict the quality of the winding during the wire-bending process itself and several manufacturing steps before the motor assembly using those windings. The benefits are a 10% to 20% reduction in a motor being assembled with a suboptimal winding, and a significant cost reduction associated with rework, scrap, and manufacturing a replacement motor.

In the case of continuous process operations, one of our

customers operates a mining and processing plant. They were experiencing unexpected downtime due to "unknown" low-quality ore grade that would clog their vertical mill. Each time the mill clogged, it cost the company \$30,000 per hour of production, potentially leading to a loss of \$720,000 per day. Using Falconry LRS, their engineers were able to recognize the patterns three steps ahead of the mill that led to the blockage conditions. From there, an early-warning system was created to detect low-quality ore grade upstream of the mill and mitigate downtime events.

How easy is it for a company to unlock value in their data once they start using the Falconry system?

It's very easy for a company to get started and, depending on the use case, see results in as little as a few weeks. To help them achieve their objectives, our customer success team meets with them before the project begins to identify the use case, what approaches they have tried before, define the criteria for success, and ensure they have the data infrastructure in place to collect and archive the necessary time series data.

After that, product training is about half a day and we recommend customers create models and run assessments on different sets of data and iterate to optimize the most the value. The created models can then be deployed out at the edges of their environment or run centrally if they prefer. While we do provide guidance and best practice sharing, we have seen that users are often self-sufficient after the first week, and start to capture measurable results in a few weeks.