



“SYMPHOS 2019”, 5<sup>th</sup> International Symposium on Innovation and Technology in the Phosphate Industry

## Merging AI and HI to boost plant OEE in the chemical industry

Chen Linchevski,

*Precognize, 5 Nahum Het Street, Tirat Carmel, 3903203, Israel*

---

### Abstract

Precognize focuses on predictive analytics for the process industry, including chemical processing plants, phosphate plants, and oil and gas and petrochemical refineries. We address this industry because its needs, circumstances, and requirements are significantly different from those of any other industry, handicapping standard predictive analytics from making an impact. The SAM GUARD solution overcomes the obstacles that prevent other predictive analytics systems from working effectively in a busy process manufacturing plant, by combining artificial intelligence with human domain knowledge to produce a solution that is greater than the sum of its parts.

*Keywords:* artificial intelligence; machine learning; predictive analytics; digital transformation.

### 1. The chemical industry faces growing challenges

The chemical industry in general is coping with increasing challenges which come from a number of different directions. These include:

**Fluctuations in commodity cost.** Energy prices are extremely volatile. While lower-cost energy sources have entered the market recently, price still changes greatly in different countries. The cost of raw materials can also fluctuate immensely around the year and around the globe; yet plants have to keep quality and price consistent.

**Regulation and compliance.** Governments and regulatory bodies worldwide have raised pressure on process plants to maintain a consistently high-quality product. Recalls have a serious impact on a company's reputation and revenue. Plants need to provide in-depth product details and retain strict quality management records, so that they can provide information in the case of an audit.

**Safety and environmental regulations.** These have increased in recent years, adding to the strain on plants to deliver high-quality, low-cost product without compromising employee safety and health, affecting the health and wellbeing of those living near the plant, or harming the greater environment.

**Balancing competing demands.** Manufacturers of all types need to balance the demands of shareholders to increase their revenue and profits with the demands of customers to deliver high quality for lower prices, at the same time as meeting safety, sustainability, environmental, and quality compliance regulations.

**Maintaining a consistent quality with inconsistent raw materials.** Seasonality, purchase cost, and greater socio-economic factors like wars, droughts, and rebellions can affect the quality and availability of raw materials. Chemical plants need to optimize their formula and resource mix in order to maintain a consistent high quality product without varying the product price.

**Rising cost of raw materials.** For many chemical refineries, raw materials make up as much as 50% of their production costs. In the face of tariff changes and an unsettled global economy, prices are continuing to rise, making it imperative for plants to increase productivity so as to make the most of expensive raw materials.

### 2. What chemical plants need from predictive analytics

These challenges impact chemical processing plants in a number of different ways, but they all underscore the importance of maximizing plant productivity. When budgets have been tight, chemical plants have primarily focused on reducing maintenance costs and prolonging equipment lifetimes, to prevent large, unexpected expenses that damaged their bottom line. Today, in the face of the rising cost and varying quality of raw materials, increased regulation around quality, safety, and the environment, together with competing demands from various



stakeholders, it's no surprise that chemical plants are focusing on increasing production, avoiding failures and downtime across the entire plant, and improving product quality. It's an issue that is generally referred to as OEE, or Overall Equipment Effectiveness. Equipment maintenance is still a concern, but it has been subsumed into the greater issue of OEE.

Focusing on predicting maintenance for specific, high cost pieces of equipment is not enough to improve overall OEE, which requires a holistic approach to plant performance. Problems that affect OEE can arise anywhere, from the expensive processors and condensers to the valves and catalysers. Predictive analytics for OEE means looking at the entire plant, managing data from all the sensors in order to improve the bottom line.

Handling such a wealth of data requires machine learning, a form of artificial intelligence, but the chemical industry places a number of unique obstacles in its path. Given the limitations of machine learning and the specific circumstances of the process industry, how can we proceed to bring advanced AI-powered predictive analytics to bear on the challenges of improving overall productivity and efficiency? We will explore these unique stressors and situations, discuss the potential and the drawbacks of machine learning, and share the ways that combining artificial intelligence (AI) with human intelligence (HI) can together provide an effective solution for chemical plants and refineries to boost their OEE.

### **3. The drawbacks of machine learning in the chemical industry**

Machine learning and the chemical industry face a difficult relationship. On the one hand, the complex predictive analytics needs of the chemical industry to improve OEE demands the input of machine learning. Every plant has thousands, if not tens of thousands, of sensors, covering every asset. Only AI can make sense out of this onslaught of data; it is far too much for a human analyst to process. But on the other hand, chemical plants hold several pitfalls for machine learning. Both supervised and unsupervised learning methods for machine learning models are insufficient for plants, in specific and different ways.

#### *3.1. The weakness of supervised learning*

The success of supervised learning depends upon historical examples, and that's something that is sorely lacking in chemical plants even though they are awash with data. Historical examples mean instances when the same piece of equipment failed in the same way, under the same operating conditions, multiple times. If a machine learning model has enough of these historical examples, it can predict when and how the next instance will occur. In general, you need around 30 examples of failures for supervised learning to succeed.

However, it's exceedingly rare for a problem to occur under the same operating conditions multiple times in a plant. If a failure occurs in a particular part or machine, it's likely to have been only once or twice, and each time under completely different conditions. For example, if the same valve fails first in the summer, and then again in the winter, the conditions are different. If you use these historical examples to try to predict failure in a chemical plant with supervised learning, it will simply not work.

What's more, in a plant we find that the failures of the past are rarely the failures of the future. If a piece of equipment fails, personnel will fix it so that if it fails again, it will be for an entirely different reason. Once a problem is fixed, it tends not to recur, so you can't even use those few historical examples that you possess.

It's vital to remember that while supervised machine learning is very smart, it's only as smart as the training data that you feed it. If your examples are biased or lacking in some way, your results will also be biased, lacking, or somehow unreliable, no matter what tool you use.

#### *3.2. The fallibility of unsupervised learning*

Since supervised learning lets plants down, the next option is to turn to unsupervised learning. Unsupervised learning involves defining the normal state for a plant, and then setting your machine learning model to alert you when there's an abnormality. It rests on the assumption that you can set an algorithm to spot anomalies in your plant data.

Here too, the unique circumstances of chemical processing plants make unsupervised learning inappropriate for your needs. A plant has thousands of sensors sending measurements every minute. To gain a holistic picture of the state of the plant and be able to spot anomalies, you need to assess the states of all your sensors together with each other. Looking at the data from only a few sensors at a time creates an unbalanced picture that can hide a growing and serious issue.

Consider the simple example of a leaky bathtub. If you have a stream of hot and cold water coming into the tub, and water leaking out the bottom, the temperature and water height may remain at a steady state, which could indicate that all is well. You would not know that there is a leak unless you measure the incoming stream or the drain, or both.



Fig. 1. Even in a leaky bathtub with water coming in, the temperature and height can be at a steady state, hiding the fact that it is leaking

Unlike in a bathtub, the number of sensors and states in a plant is vast. Every plant has thousands of sensors. Every sensor could be in one of several different states at any given moment. The number of states in a plant is written as  $2^n$  (2 to the power of  $n$ ), where  $n$  is the number of sensors you have multiplied by the number of states of each one of the sensors. It's a massive number, far bigger than the number of all the atoms in the universe, which is  $2^{86}$  (2 to the power of 86). This creates an overpowering amount of noise.

What's more, chemical plants also have to contend with a finite amount of historical data. The finite selection of data combined with the enormous number of states results in states that are so sparse as to be at zero for all practical purposes. The probability that the plant is in a state that has occurred before is essentially nil, which means that there is effectively no such thing as a "normal" state for the plant. In fact, across the entire plant, one could say that "normal" is an anomaly.

Some plants attempt to use unsupervised learning by taking all this data and cleaning off the noise. It is possible to take only a few of the plant assets and identify exactly which tags concern you, and then clean off the noise. It requires a great deal of manual work, and takes up a lot of time for the Operations team, but it can be done. However, it is not scalable. It cannot then be rolled out to the hundreds of other assets in the plant; each asset needs to be cleaned and prepared individually, which is not a practical option in busy plants.

#### **4. Combining AI with HI: the unique Precognize solution**

Chemical plants need machine learning to apply predictive analytics to the complex data of the processing plant, but they cannot use supervised learning because of the lack of historical examples, and they cannot use unsupervised learning because of the high noise level and sparse states. Both supervised and unsupervised machine learning are fallible, but if the results of the machine learning algorithm are put into context, it can give them meaning.

One way to provide this context is through human domain knowledge, and this is the approach that SAM GUARD takes. SAM GUARD draws on the domain knowledge of someone who understands the plant thoroughly,

like the process engineer, to create a domain model. It's simple and fast to create a domain model, and doesn't require any advanced conceptual modeling skills.

The domain model is turned into a mathematical graph, so that SAM GUARD can automatically apply it to the plant sensors and data. Using this domain model graph, SAM GUARD can identify separate regions within the plant, to divide it up into smaller, more relevant groups of sensors. Operations teams can now focus their attention on only those regions that are relevant. With fewer sensors in each region, it becomes possible to apply unsupervised machine learning techniques to the data and sensor states to define "normal" for this area of the plant, enabling SAM GUARD to spot real issues, instead of getting lost in the fog of constant anomalies. And users benefit from receiving only the relevant alerts that relate to those real issues.

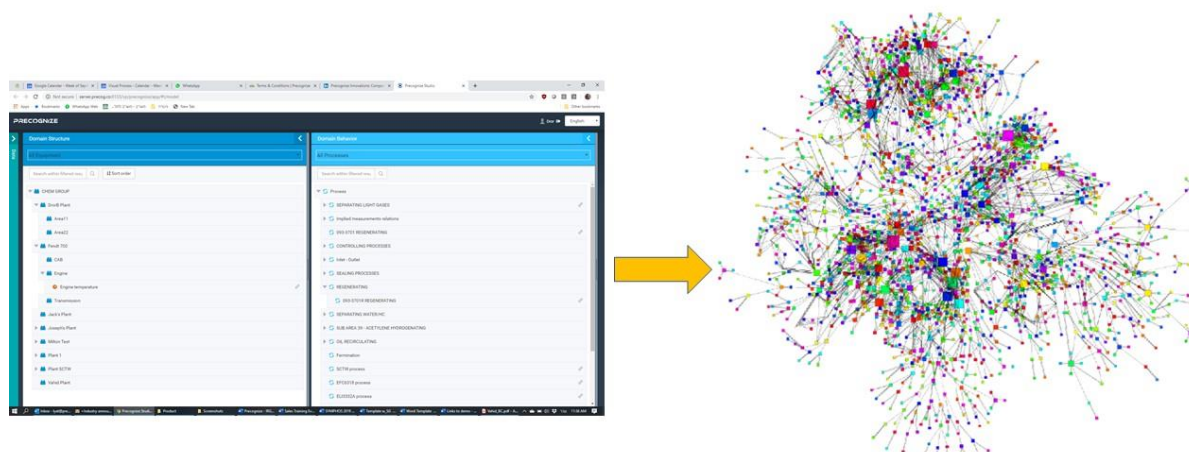


Fig. 2. SAM GUARD turns the domain model into a mathematical graph of the entire plant. Anomalies are then projected onto that graph to determine their relevance

SAM GUARD automatically identifies the relevance of each sensor, and then monitors only those which are in appropriate regions. Problems are aggregated through both time and space, so users only receive an alert if a problem is developing in a relevant location.

SAM GUARD's approach to predictive analytics combines Human Intelligence (HI) with Artificial Intelligence (AI) to produce a new machine learning approach that is neither supervised learning nor unsupervised learning, based on the understanding that the chemical process industry needs a different strategy. It uses domain knowledge to enhance machine learning, guiding users to the right areas of the plant to search for potential issues that could affect OEE.

What's more, SAM GUARD adds a feedback layer to make the AI-HI combination even more effective. When SAM GUARD sends an alert, it begins by presenting a very small number of tags and asking the human analyst or Ops team member to send a response about which tags were helpful, which were of no use, or if none of them were of any use. Then it will show a new set of tags, and then another one, until the problem is entirely cleared up and resolved. The system learns from the human response and remembers it for the next time, so that it learns through each iteration which are the most useful tags to show alongside the alert.

## 5. SAM GUARD Case Studies in the Phosphate Industry

Let's explore some examples of SAM GUARD in action.

### 5.1. Use case #1: High temperatures in a pump oil tank

In a phosphate plant, the SAM GUARD system issued an alert that the temperature of oil in a pump oil bearing tank was alarmingly high. It flagged a temperature 10 °C higher than normal. SAM GUARD presented the duty engineer with a unique second warning about a higher-than-usual oil level in the secondary oil reservoir as well, thanks to the domain model graph, which saw the connection between the two anomalies. The relation between the two alerts enabled the engineer to immediately understand the problem and the action that was required.

The engineer realized that the secondary oil tank had been filled with an excessive amount of oil, making it impossible for the pump bearing to remove extra oil back to the reservoir. If the high temperatures had continued, it could have caused a bearing failure that would have damaged the pump.



Fig. 3. SAM GUARD sent two alerts that it identified as related to each other: higher-than-usual temperature and higher-than-usual oil level. The two alerts together enabled the engineer to immediately understand the problem and the required action

### 5.2. Use case #2: Increased pump vibrations

An engineer at a phosphate plant received an alert from SAM GUARD's predictive analytics solution regarding increased vibrations in a pump. No other system in the plant flagged this rise in vibrations, since they were still within acceptable values, but SAM GUARD noticed the deviation from standard vibration levels while performance and other significant parameters -- such as power -- remained the same. This unusual confluence of two measurements, one out of usual range and one within usual range, triggered an alert.

The engineer investigated the alert and verified the problem using an external vibration measuring device that identified the source of the problem in pump #3. She found signs that the pump was imbalanced and suspected that materials were crystallizing on the pump impeller.

The decision was taken to stop the pump for cleaning, thereby preventing the high vibrations from damaging the structural integrity of the concrete pump station and averting pump failure.

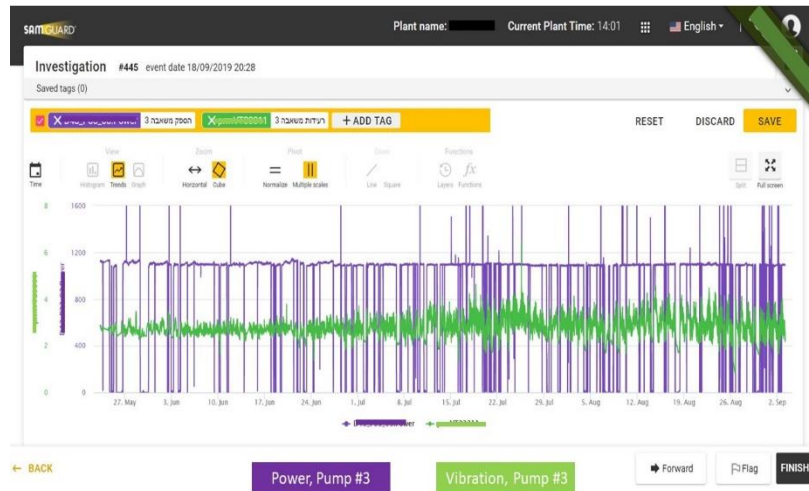


Fig. 4. SAM GUARD alerted about the pump vibrations far before it reached the static threshold that would have indicated a problem, thus identifying the problem early enough to prevent damage or even failure

## 6. Conclusion: Predictive analytics is a spotlight, not a magic wand

It's important to understand that machine learning and predictive analytics are not a magic wand for chemical processing plants. It's not possible to apply the magic spell of AI and expect the entire system to improve. Instead, the SAM GUARD solution serves like a spotlight or a laser pointer, guiding you as to where to focus your attention. There still needs to be a human element to respond to the alert and identify what needs to be done next, but thanks to SAM GUARD, that person receives a relevant alert that directs them to investigate the right issue.

It's clear that typical supervised or unsupervised machine learning alone will never be able to produce effective improvements in OEE for process plants. There's not likely to be any change that increases the number of historical examples to enable supervised learning, or rationalizes the sparse number of states for unsupervised learning. However, as the SAM GUARD approach of domain knowledge together with machine learning is upheld, the system feedback will continually improve to produce an ever-more efficient solution that enables plants to quickly identify potential problems and optimize OEE.

## Acknowledgements

Special thanks to Dr. Hila Nachlieli, Head of Research at Precognize, for her invaluable contribution to this paper.