Learning to Spot Trouble Before It Starts

Modernizing Plant Monitoring with Machine Intelligence


Ahmad Al Rashdan, Jacob A. Farber
Plant Modernization Pathway

The control room operators in a nuclear power plant are well prepared to handle any issues that arise, but they generally only act when they are alerted that something is wrong. By the time an alarm is triggered, and the control room operators respond to it, the nuclear power plant may have initiated an automatic shutdown, resulting in the loss of valuable energy production. Additionally, a piece of equipment may have been damaged, complicating the repair process.

The solution is to not have to wait for alarms to happen. With assistance from a computer system that monitors the numerous sensors in the nuclear power plant, the operators could spot anomalies before they lead to problems.

The idea that historical data patterns could be used to detect anomalies early enough to prevent shutdowns or damage is not new. Engineers have long used computers to detect when a group of sensors gives readings that depart from a known pattern. But it is hard to execute at a nuclear plant, because of its immense size. With thousands of embedded sensors, it produces a plethora of data, which can be hard to analyze.

To overcome this, machine learning is used to analyze historical data and detect subtle changes. It determines what is a significant change in plant conditions to alert the operators accordingly. LWRS Program researchers at INL call it the ALARM software, for Automated Latent Anomaly Recognition Method. It is essentially an algorithm that establishes correlations between sensor readings and off-normal conditions, in a way that smooths out the operations of a plant and saves money.

An example of how ALARM can smooth out operations and save money is detecting cooling fan degradation before the onset of failure. Cooling fans are used in boiling water reactors (BWR) to help limit temperature in the drywell, a structure that dissipates excess heat from the reactor. In a BWR, two cooling fans failed within a four-week timeframe, forcing the reactor to be shut down. One of the challenges was that there were no vibration sensors or other readings on the cooling fans that could be used to detect their degradation more directly. However, if fans were beginning to fail, places in the plant near the fans would show odd temperature readings. ALARM would monitor something broader than the fans themselves: the whole environment, from which the condition of the fans can be inferred. Two fans failed within four weeks, and the reactor had to be shut down for repairs. No vibration sensors or other readings could be used to detect this ahead of time. But if the fans were beginning to fail, places nearby in the plant would show odd temperature readings. The innovation is to monitor something broader than the fans themselves: the whole environment, from which the condition of the fans can be inferred.

Figure 8. The ALARM software consists of several anomaly detection modules that significantly reduce the number of false alarms (from the alarm list), ensuring effective use of anomaly detection.

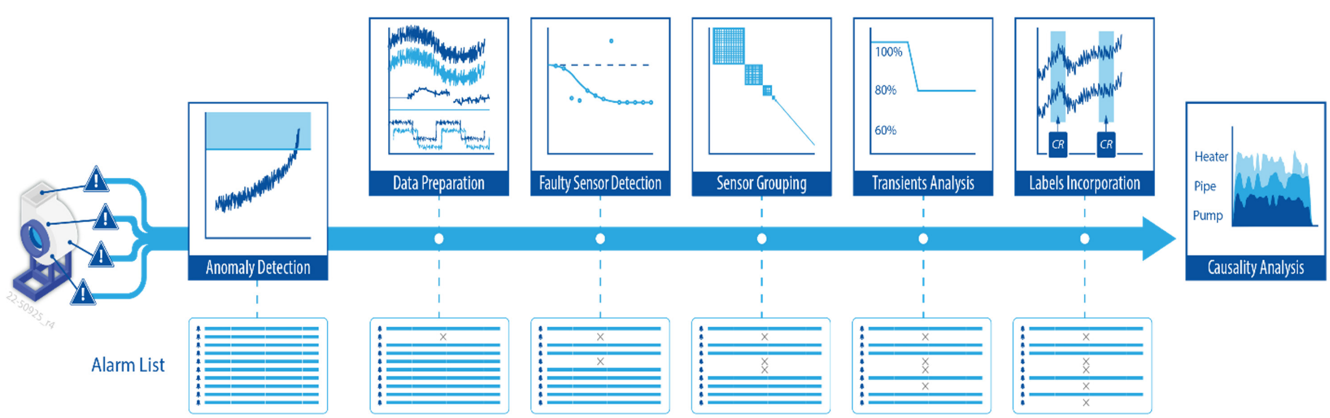

Two main categories of errors make any machine intelligence less useful. The first one is when it fails to detect an anomaly while the second involves triggering unnecessary alerts, alerting the control room operators so often that they become desensitized to the warning.

To reduce those errors, the ALARM software consists of several modules (Figure 8) that enhance anomaly detection. The first module detects deficiencies in sensor data and focuses on identifying indicators such as prolonged periods of unchanging data or apparent delays in data transmission. By identifying weaknesses in sensor readings, the software effectively minimizes false alarms.

Some of those readings will come from sensors giving incorrect data. Their readings may drift, or show spikes or other noise in the data, or they may simply fail. A second module looks for those sensors because fixing a sensor is easier than fixing an actual mechanical problem in a plant.

Scaling anomaly detection for a whole plant requires analyzing clusters of related sensors simultaneously, and then drawing inferences. The thousands of sensors within the plant can be divided into groups, treating each group as a single system or process, which is the focus of the next module. In addition to improving anomaly detection performance, this technique reduces computational requirements, because correlations need to be established only among a smaller group of sensors.

The next module looks for changes that will occur when the reactor operates at less than full power. A sensor reporting reduced temperature or pressure might trigger an alarm, but the reading might be normal for a reactor operating in that mode.

As alarms are generated in cases where they were justified, human operators can provide feedback to the automatic software as part of an ALARM module. By labeling the data, humans will help the machine logic get smarter.

In the complex and tightly coupled systems of a nuclear power plant, when multiple alarms sound simultaneously, human operators may require time to determine the root cause. Learning to recognize the source of an anomaly is the focus of the final module.

The goal of all these modules is to sort anomalies that represent problems from anomalies that do not. Software that can reduce the number of times it issues false alarms can increase the quality of information going to the operators, and the operators' confidence in the software. Ultimately this results in a lower number of unexpected outage and enormous cost savings with minimal distraction of operators.

Diego Mandelli honored with International Safety System Society Scientific Achievement Award

Diego Mandelli was recently recognized at the annual International System Safety Summit and Training Conference as a distinguished expert in the field of Risk-Informed Asset Management. He received the prestigious 2023 International System Safety Society Scientific Achievement Award. This esteemed award recognizes individuals or groups who have significantly advanced the realm of system safety by dedicating their efforts to pioneering research and development programs.

Mandelli received this award in recognition of his outstanding accomplishments in the development of data analytical methods for equipment reliability analysis and advancing model-based system engineering tools.

Diego Mandelli

Mandelli's dedication to improving system safety through innovative R&D is a major asset to the Light Water Reactor Sustainability Program. His contributions continue to shape the landscape of Risk-Informed Asset Management, providing invaluable insights and tools for extending the lifetime of the existing fleet of nuclear reactors.

INTERNATIONAL SYSTEM SAFETY SOCIETY
Professionals Dedicated to the Safety of Systems, Products & Services
Organized 1962
Incorporated 1973

SCIENTIFIC ACHIEVEMENT AWARD

Presented to

DIEGO MANDELLI

In Recognition of
Outstanding Development of
System Safety Analysis Methodology

Conferred on this day, the 30th of August 2023

David West
President

Liz Anderson
Executive Director