AISTech 2023 — Proceedings of the Iron & Steel Technology Conference
8–11 May 2023, Detroit, Mich., USA
DOI: 10.33313/387/237

"Hands-Free" Fully Autonomous, Plant-Scale, Anomaly Detection AI

Joe Porter, Keval Bhanushali, Dan Kearns, Nikunj Mehta[1]

Falkonry Inc.
10020 N De Anza Blvd #200
Cupertino, CA 95014 USA
Phone:+1 408 461 9286
Emails: {joe.porter|keval.bhanushali|dan.kearns|nikunj}@falkonry.com

ABSTRACT

This paper presents a self-supervised autonomous "plant scale" AI — capable of monitoring every PLC and IIoT parameter of a steel plant — automatically detecting and accelerating diagnosis of anomalies. Automatic anomaly detection proactively informs plant operations of conditions that otherwise would go undetected — leading to informed production and maintenance decision-making. Self-supervised AI overcomes the challenges of constant equipment, environment, and product changes that thwart classical machine learning approaches. Normalized severity scoring of the AI results further enable prioritization of the anomalies for investigation and action. We describe several use cases of this new AI in commercial operation along with the corresponding user workflow.

Keywords: Generative AI, Time Series AI, Deep Learning, Time Series Analytics, Maintenance Productivity, Electrical Maintenance, HRT Motor, Continuous Casting, Hot Strip Mill, Condition-Based Maintenance

INTRODUCTION

Steel manufacturing processes are heavily automated using PLCs and generate large volumes of industrial automation data. This data is in the form of time series, which could rise to over 5 million data points per second per plant. The time series data precisely represent the state of a physical system and production process at any given point in time. However, this data is hard to use for troubleshooting and improving operational productivity. The advent of cloud connectivity and digitalization has exponentially grown the volume of data collected . For instance, a typical integrated steel plant collects over 30,000 parameters from automation systems and sensors. Operating conditions change from time to time as programming parameters change and that makes it imperative to update baseline behavior and corresponding configuration of alarms. As a result, only 1-2% of these parameters are monitored to provide alarms to operators. This has the consequence that operational problems are frequent, troubleshooting takes a long time, and expertise is not easily transferred from expert to novice. Substantially increasing the analysis of automation and sensor data is therefore essential to achieving smart manufacturing objectives for the industry.

In recent years, various applications of artificial intelligence and machine learning have demonstrated the potential to harness some of the automation data being acquired for analysis by surfacing patterns of interest in the data. In order to fulfill the smart manufacturing mandate, organizations need an analytics approach that does not overwhelm the plant personnel with set up and maintenance. It is also ideal that such an approach exploits all available data and exposes important patterns in it. In this paper, we present a novel 'Unattended AI' that *automatically* analyzes all available parameters and presents the anomalous behaviors to the operations and maintenance users. This novel deep learning approach eliminates manual data analysis efforts and considerably reduces the time taken to diagnose production issues in the mill. As a result, users prioritize maintenance actions whether they have to be immediate intervention or targeted for a later action such as planned downtime. This paper presents several real-world applications of this novel 'Unattended AI" to minimize equipment downtime and scrap occurrences, while reducing the manual analytical effort to raise the productivity of the people involved.

© 2023 by the Association for Iron & Steel Technology.

Shreebhooshan Badrinarayan / Falkonry, Inc.
Order # 132176 / Order Date: 6/1/2023
Copyright Association for Iron & Steel Technology

CHALLENGES

Finding a Needle in a Steel Haystack

A modern steel mill operates at high speed and a staggering scale and produces hundreds of thousands of primary data points each second. The expectation is that such instrumentation will allow operators, technicians, and decision-makers to reduce operating costs for the mill and increase production throughput. Even as this automation substantially increases operational productivity, the increased complexity resulting from automation makes operational problems such as stoppages, scrap, and safety issues harder to troubleshoot. When such issues occur, it takes a significant amount of human effort to ascertain the cause of such issues. Often plants employ process and reliability engineers primarily to study data acquired from automation systems and correlate it with operational problems. Such engineers are hard to recruit and train, and their analytical capacity is usually well below organizational needs. Finding evidence of the causes of operational problems from plant data is time-consuming, and as the mill continues to run at its breakneck pace, the work of investigating problems can become an exercise in managing a backlog of issues while fighting fires. Automated anomaly detection would significantly ease the burden of addressing and avoiding issues in a steel mill, but there are several factors that make this challenging to realize:

1. The difficulty of specifying and maintaining operational waveforms of a “golden curve” given how much they vary from time to time and the need to manage continuous operations as opposed to short batches
2. The difficulty of matching operational waveforms to the right “golden curve” using rules tends to make golden curve-based analytics fragile and error-prone
3. Speed of operations in the mill, and the vast number of signals that must be analyzed mean that most available data is not used effectively, and causes are not well understood from data
4. Bookkeeping — reliably tracking innumerable details in measurements from mill components with limited personnel presents a grand challenge
5. Maintaining models — a steel mill presents immense complexity as components interact and conditions change. It’s essential to be able to retrain models, both to capture additional behaviors of normal operation when initial training is not adequate (incremental retraining), or to completely retrain a model when mill components or operating conditions change more dramatically

Monitoring: Golden Curves and Rules

Due to the scale of production in steel plants, anomalous behaviors are likely to induce operational productivity losses. As a result, it is very useful to monitor for anomalous behavior to be detected, characterized, and encoded into a form that automates the activation of meaningful alarms when anomalous behavior recurs. These alarms may be triggered by simple trends, such as when a particular signal exceeds a baseline level. Other conditions are more subtle, and detecting their recurrence amounts to measuring whether signal data has deviated too far from a “golden curve” — a previous characterization of acceptable operating states. Rules and golden curves are prone to error and fragile in operation — error-prone because of the manual effort needed to specify alarms, and fragile because signal data is inherently messy. In a true anomaly, the measured data may not reliably deviate from the representative data due to inaccuracies in timespan, shape, or consistency (e.g., due to persistent noise, transient spikes, or missing samples).

The more significant problem with these monitoring approaches is that they are labor-intensive: If a critical signal has numerous discrete states, then rules or golden curves must address all of the possible valid operating states for the components being monitored. Creating alarms requires time and expertise. To date there has been no simple, generic, robust, automatic way to compare time periods to determine similar behavior across long operating times without significant manual work.

Data Scale in an Instrumented Steel Mill

In order to maintain the enormous throughput of a modern steel mill, speed and accuracy of measurements and controls are essential. Scan times must be kept tight to ensure adequate control, particularly for electrical components. The extreme speed of sampling and control, together with the number of signals available means that many signals and events will remain completely unmonitored. The simple act of identifying operating states and specifying configuration for them is prohibitive, except for the most crucial situations. Unfortunately, crucial situations are only determined in hindsight. When fighting fires, it is very hard to also investigate causes, which makes the exercise of analytics even harder. The challenge is to create a monitoring system that reaches the scale of the entire mill — all components at all times, and in all of their modes of operation.

Promise of AI and Automated Monitoring

Imagine having a data scientist working on every signal in the mill — capturing subtle variations, characterizing nominal and outlier behavior, and identifying which outliers lead to problems and which are benign. This is the grand challenge of automated anomaly detection. In our current mode of operation, we can drill down into any chosen area of operation, but we are still

Shreebhooshan Badrinarayan / Falkonry, Inc.
Order # 132176 / Order Date: 6/1/2023
Copyright Association for Iron & Steel Technology

limited by the availability of expertise. Truly automated monitoring must be capable of automatically doing the extensive bookkeeping required to ensure full coverage of signals and system behaviors, without requiring an army of data scientists to operate it.

Automated Model Training

Short, fixed segments of training data are useful to capture the initial behavior of a signal in order to create a model. However, unless we're very lucky, the first training will not capture all of the important and acceptable variations that the signal exhibits operating over an extended period of time. Further, components may be changed out — which have similar but not identical behavior. Decisions may be made that alter nominal operating conditions, in order to improve the quality of the steel or the efficiency of the mill. In all of these cases, the automated monitoring system must be able to adapt to changes — if we realize the promise of full automation, the AI should be able to detect when additional training is needed for an existing model and schedule itself for training. If plant conditions have changed enough to invalidate the model, automatic retraining should start from scratch to produce a new model that captures the new operating regime.

AUTONOMOUS TIME-SERIES AI

Self-Supervised AI (CVAE)

Our approach to "hands-free" autonomous, plant-scale anomaly detection AI is based on a self-supervised deep learning model called a convolutional variational autoencoder (CVAE). An autoencoder encodes its current input data in a mathematically compressed form (an embedding), and then attempts to reconstitute the original input from that compressed representation. The autoencoder approach uses the fact that during training, the model learns to faithfully reproduce signals that it has seen, so that at inference time, novel signals induce reconstruction error — the autoencoder hasn't yet learned how to represent that novel input. Therefore, we use a measure based on reconstruction error to characterize the novelty of any given input, indicating a potential anomaly. This measure uses specially designed normalized statistics in order to be independent of the signal characteristics — this allows the error to be easily compared from one signal to another.

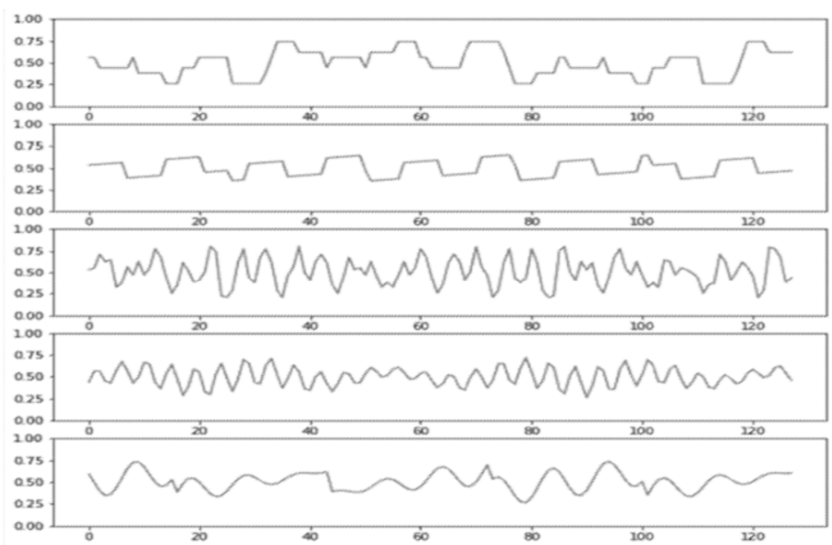

Figure 1. Examples of synthetic training signals.

We're able to avoid the large training data requirements typical of deep learning models because we use synthetic data to bootstrap the model, followed by a much shorter fine-tuning training step on actual plant data that adapts the autoencoder to the characteristics of specific signals. Time series ML is unique, in that it is very amenable to the use of synthetic training data — time series signals encode information temporally, in the value domain, in the time domain and in the frequency domain. We can use standard signal waveforms (sawtooth, square wave, sinusoid, etc…) in random combinations to bootstrap a generic CVAE model for time series data. This generic model is then used repeatedly to initialize signal-specific models. In our experience working with tens of thousands of signals, we have seen that the generic model can be fine-tuned to create a usable signal-specific model with as few as 1,000 samples, regardless of the sampling rate of the data.

Shreebhooshan Badrinarayan / Falkonry, Inc.
Order # 132176 / Order Date: 6/1/2023
Copyright Association for Iron & Steel Technology

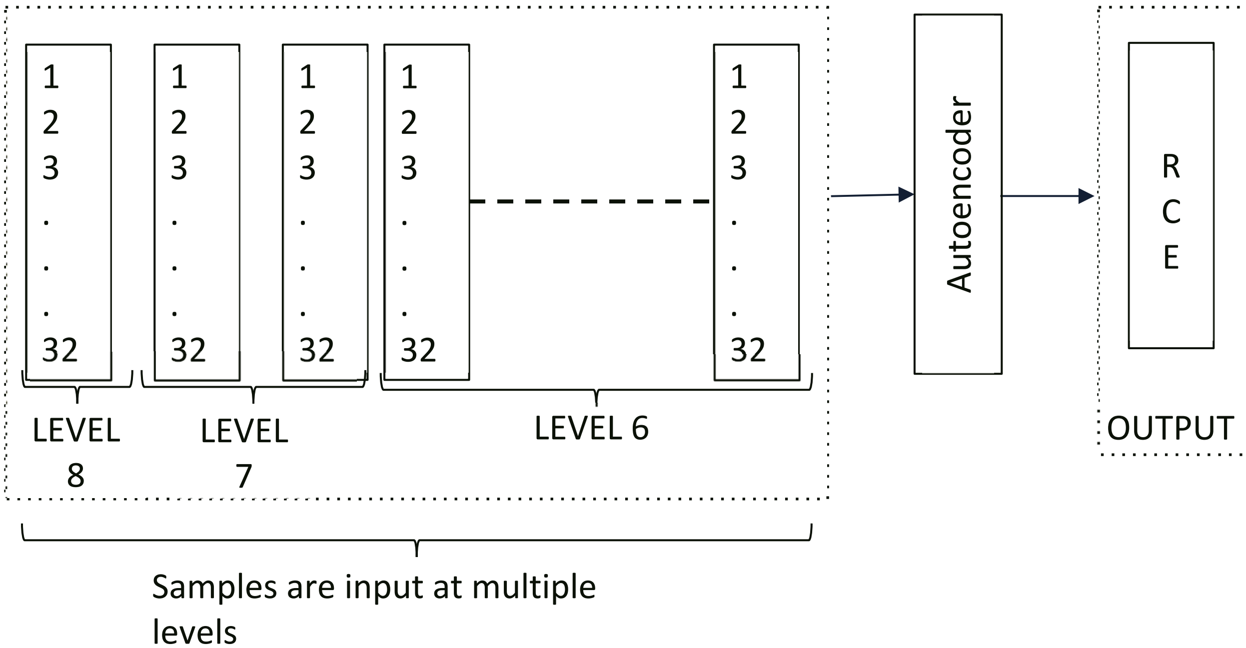

Figure 2. Diagram depicting training data organization for the CVAE. ‘Level’ indicates a preset sampling rate, and high level corresponds to low sampling rate (and vice versa). A single autoencoder instance will be trained with data from multiple levels.

For times series data, the CVAE architecture requires that each training sample contains a large amount of context. We feed multiple values at once to the autoencoder in a data structure we call a cell. Furthermore, we train with data at different sampling rates, using many more high resolution cells than cells at lower resolution, in order to capture subtle details of the training signal. We model sampling resolution as levels, which correspond to preset sampling rates. These levels follow an exponential arrangement. Level 6 represents the millisecond resolution and would be used when data is sampled faster than 100 Hz. Likewise, Level 9 represents the second resolution and will hold data if there are any samples taken at a one second level or faster. As Table 1 shows, our CVAE architecture will consume data at every level which is available for the given signal. There will be roughly 10 times more data at any level compared to a level that is just 1 higher. This arrangement of levels, which we call tiles, is also very helpful for storage efficiency and visualization.

Table 1. Data required for Learning at Different Time Resolutions

Approximate extent of Data from learning	**Sampling frequency (resolution)**	**Level**
A few seconds	1kHz (1 ms)	6
A few minutes	100Hz (10 ms)	7
Many minutes	10Hz (100 ms)	8
A few hours	1Hz (1 s)	13

For training and inference, we start with a block of time series samples. The samples are then divided into cells, which are further divided into episodes. An episode is a representation of a single window of samples and is the input unit for a single training or inference step of the machine learning model. For a particular signal, the data acquisition can be configured to distribute data within the cells. Configuration controls the quantity of data represented in each episode, as well as the distribution of data summarized within the episode. Multiple episodes are consumed in parallel in a batch of training or inference data for a single model, and multiple models can be run efficiently on the same GPU, allowing us to scale the compute resources to reach the scale of the signal data to be covered in the mill.

The final requirement of the hands-free autonomous architecture is automated improvement, i.e., incremental learning. For any given signal, we expect the anomaly scores to remain low most of the time. If a model has not been trained sufficiently, or if the distribution of the signal being monitored has drifted from the distribution that the model originally learned, then we expect to see lots of spuriously high anomaly outputs for that signal over an extended time period. It’s safe to assume that consistently high anomaly outputs over a long time do not correspond to catastrophic recurring failures, as those would be evident in the

Shreebhooshan Badrinarayan / Falkonry, Inc.
Order # 132176 / Order Date: 6/1/2023
Copyright Association for Iron & Steel Technology

operations of the mill. Therefore, we flag that signal for retraining, which will automatically occur as computing resources become available. This automated improvement is a key element of “Unattended AI” as the AI improves itself without the need for labeling or manual intervention. “Unattended AI” is the only viable approach when dealing with thousands of signals, as it becomes difficult to coordinate different parties to conduct incremental learning.

Advantages

Our machine learning architecture has several significant features that allow us to make the claim of ‘hands-free’ fully autonomous, plant-scale anomaly detection AI:

- Self-supervised learning — no need to annotate examples, as we need only provide adequate data samples for each monitored signal. A CVAE model captures all of the subtle variations and operating modes represented in each signal, implicitly representing the information that a data scientist would study in order to develop models for monitoring.
- Low sample requirements for training — because of synthetic pretraining, a few thousand samples can suffice to create a model for any signal.
- Robust data representation — avoid problems with extreme and missing values, and differing sample rates while eliminating the need to prepare data ahead of machine learning.
- Automatic retraining — unattended self-improvement of each model via retraining criteria.
- Normalized anomaly scores — by normalizing the anomaly scores, anomaly measures can be compared across signals. This allows us to automate the aggregation of anomaly data within components to facilitate user navigation of signals exhibiting unusual behavior.

These criteria enable full-scale monitoring of signals in the mill without the manual effort of data collection, analysis, and model-building.

Deployment of Plant-Scale AI

Beyond the usual operational considerations required to deploy machine learning models in production (e.g., cloud vs. on-premise, managing data storage, connectivity, security, and availability), we focus on a few important capabilities that support the ability to scale to the level of operations needed to ensure hands-free autonomy while ensuring cost controls — elastic compute clusters, and custom job scheduling.

Kubernetes allows us to control costs through the use of autoscaling and resource limits. Compute resources are costly, so they are only used as needed. Kubernetes automatically spins up new compute resources as they are required, up to configured limits. This allows decision-makers to budget for the needed compute resources and ensure that those cost limits are met.

Custom job scheduling allows us to manage the number of training and inference jobs that run on GPU nodes. CVAE models are not GPU memory-intensive, so several models can be run concurrently in a single GPU. Our custom job scheduling controls the loading and unloading of CVAE models, and the efficient streaming of sample data from online storage during both training and inference phases. The episode construction is done with the GPU, further increasing the efficiency of data movement and computation.

APPLICATION AND USE CASES

To demonstrate the capabilities of plant-scale automated anomaly detection, Falkonry Insight, a proprietary software system, was deployed to the hot mill process data of a US steel manufacturing plant. A hot mill process includes multiple stages such as roughing mill, crop shear, finishing mill, slab grinder, plate shop etc. where high temperature steel is converted into desired shape and dimensions. The entire portion of the hot strip is highly instrumented with approximately 1,500 PLC tags producing time series data at different rates 1Hz to 10Hz.

For this paper, we will focus on the finishing mill process of the hot-strip mill. The various sub-processes involved in the finishing mill are designed to achieve the desired final product specifications and any deviations from the normal process behavior can result in product defects and quality issues. Any significant deviations at this stage essentially results in bad product quality or an unplanned production downtime. Hence monitoring is critical to ensure the efficiency and smooth operation of the manufacturing process. This is where no-setup, automated time series AI from Falkonry is helpful at identifying anomalous behavior in an entire line, grouping relevant anomalies together and streamline root cause analysis of such anomalies — all for the benefit of making timely and targeted intervention actions.

Data Sources and Signals

The finishing mill sub-processes have around 400 signals that allow operations and maintenance teams to monitor flows, pressures, currents, voltages, temperatures, positions, forces, torques, tensions etc. These signals are available from the Finishing mill Steckels, downcoilers, cooling tables etc. where the downstream operations may be impacted by quality of

Shreebhooshan Badrinarayan / Falkonry, Inc.
Order # 132176 / Order Date: 6/1/2023
Copyright Association for Iron & Steel Technology

operations from the upstream processes. The signal information was made available for a period of around 3 months but Falkonry Insight required as little as a week to gather sufficient information about the operational patterns in time and frequency domain. Users have the ability to get a view of the anomaly state across a process / sub-processes of their finishing mill lines for a given time period. Further users also can quickly compare signals for different time periods and develop a quantitative understanding of root causes of the anomalous patterns reported by the Falkonry Insight's automated anomaly detection. Next, we describe the workflow by which users can learn about anomalies that need their attention and the likely factors affecting the anomaly to understanding what the causes may be. During this process, no analytics expertise is needed nor is it required to set up or maintain the automated analytics.

A. Using Dashboard to Prioritize Anomalous Events Requiring Action

Falkonry Insight enables the discovery of process and machine anomalies via the Anomaly dashboard. The dashboard enables users to focus their attention to signals or signal groups that deviate from the standard operational behaviors. The dashboard organizes anomalies requiring attention and action. This view allows users to prioritize critical sub-processes/processes that would have the maximum impact on plant productivity and product quality. This dashboard is updated continuously as the automated anomaly detection infers over the continuously arriving signal data and surfaces the anomalies for the user to review and act upon.

By filtering the list of anomalies for the previous 24 hours (or another relative time range) and sorting them by latest anomalies seen, the dashboard conveniently depicts all anomalies that were observed across different sub-processes of the Finishing mill. This potentially indicates to the user that there may have been an issue with a process/sub-process on Jan 24 that likely impacted the overall Finishing mill output. The user can then determine which anomalous equipment to analyze first based on the equipment location in the process (upstream/downstream) or using the severity score. From the image below it seems the Delivery-side Steckels and FM Cyclometer are the ones that require more immediate attention.

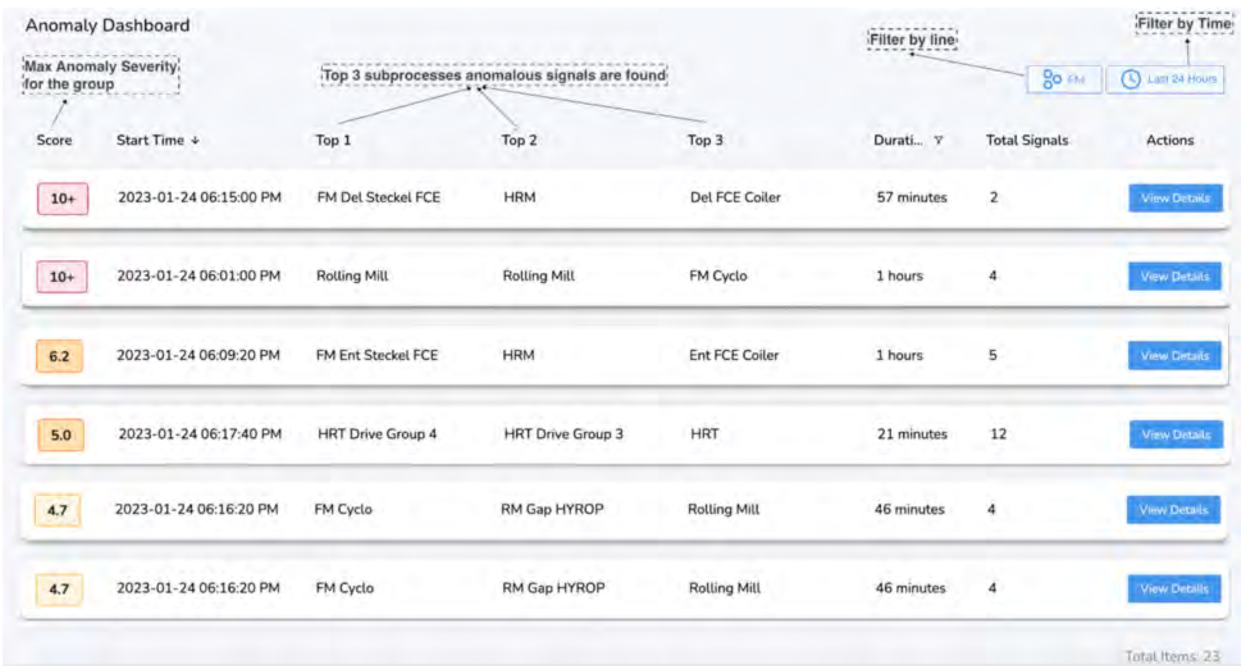

Figure 3. Anomaly Dashboard for a finishing mill process.

Figure 3 shows the Falkonry Insight Anomaly Dashboard for the Finishing Mill. It shows anomalies over the last one month. Each column header is explained below:

1. **Score** — the severity or degree of anomalousness of a signal. The score can vary from 3 to 10+, with 3 indicating the signal drifting into an anomalous zone and 10+ indicating a highly anomalous pattern.
2. **Start Time** — indicates the time when the behavior degradation was observed.
3. **Top1 , Top2, Top3 —** The top 3 groups/sub-groups in the plant to which anomalous signals belong
4. **Duration** — Length of the anomalous period detected by AI
5. **Total Signals** — Total signals identified anomalous and are potentially related

© 2023 by the Association for Iron & Steel Technology.

Shreebhooshan Badrinarayan / Falkonry, Inc.
Order # 132176 / Order Date: 6/1/2023
Copyright Association for Iron & Steel Technology

6. **Actions** — Expand the row and learn additional details about the specific row group

Assuming the user prioritizes based on severity alone, then they would investigate the Delivery Steckel anomaly. In the anomaly detail view, each signal is given a severity score for duration of the anomaly based on the maximum value of severity during that time. For this Delivery Steckel anomaly, the detail view highlights two signals — gas flow and recoup temperature that had an unusual behavior. The signal charts contain two visualizations for each anomalous signal — a line chart indicating the signal's value trend and a heatmap with colors ranging from light red to dark red highlighting the severity of anomalous behavior in the signal.

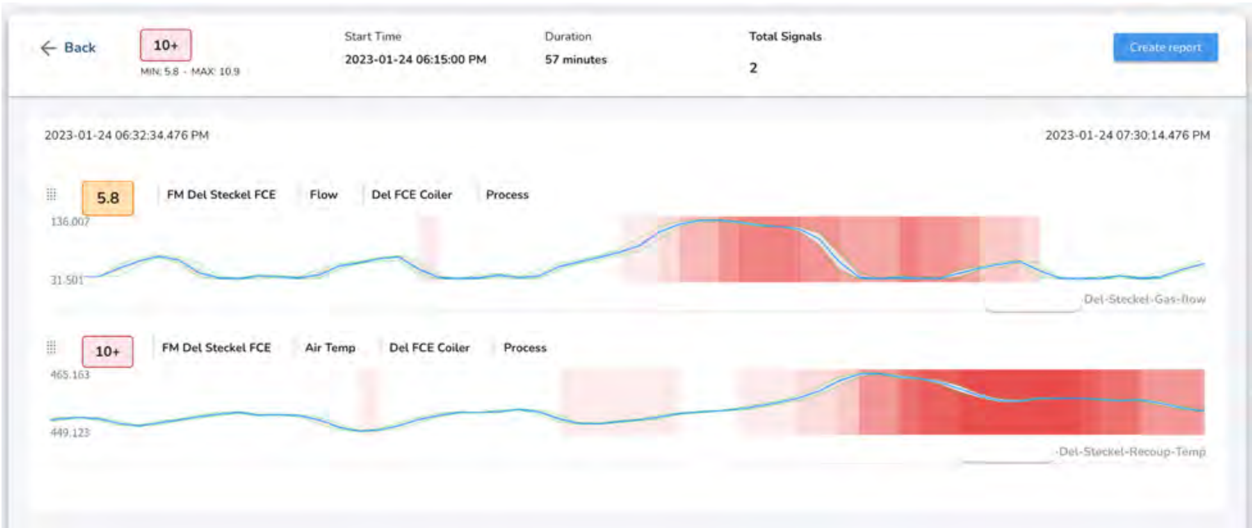

Figure 4. Details of anomalous signal behavior for a high severity anomaly in Delivery Steckel anomaly.

The anomaly detail shows that the gas flow in the Delivery Steckel jumped up 5 times of its typical operational value that may have led to an increase in the recoup temperature of the Steckel, which was beyond its normal range as well. This could potentially lower the yield of finished sheets of steel which may lead to the remainder of downstream processes being flagged as anomalous for that day. What is noteworthy is that no guidance was provided about normal operations nor thresholds of normal behavior identified.

We show some additional examples of anomaly detail view which demonstrate different kinds of anomalies that Falkonry Insight can detect in a finishing mill. Figure 5 shows pressure fluctuations that are seen as severe anomalies because the pressure stays high and generally changes more gradually. Figure 6 shows slopes that are detected as anomalies well before the levels exceed/drop below normal.

Figure 5. Example 1: Pressure fluctuations in a downcoiler bander.

© 2023 by the Association for Iron & Steel Technology.

Shreebhooshan Badrinarayan / Falkonry, Inc.
Order # 132176 / Order Date: 6/1/2023
Copyright Association for Iron & Steel Technology

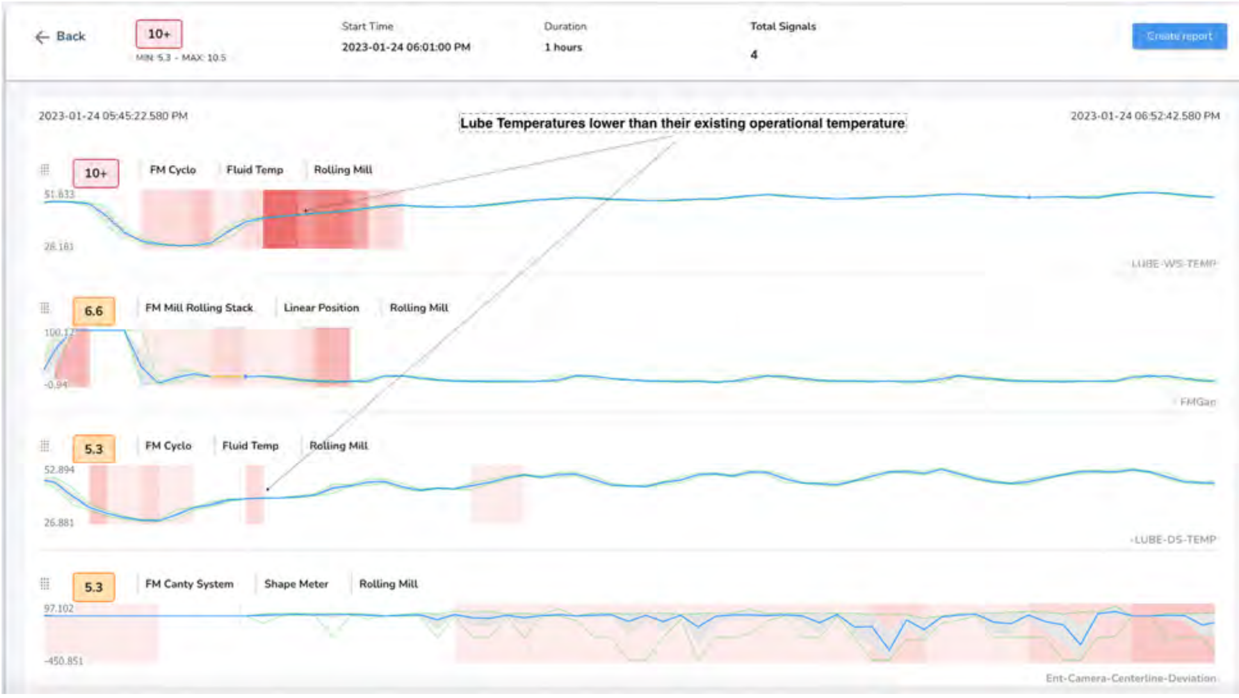

Figure 6. Example 2: Lube temperature falling and then not rising back to the desired operating temperature.

B. Monitoring Signals at the Process/Equipment Level

While the dashboard view provides the user an ability to identify anomalies and react, the Falkonry platform also provides the user a workflow to monitor signals over time to make priority decisions about planned maintenance. For this, a heatmap view is provided for different subsystems and is used to identify if there were more or less anomalies from one period to another as well as where they have occurred. A heatmap like this makes it easy to consume a lot of data in a single view and simplifies condition-based maintenance prioritization. Conversely, it also helps ascertain the impact of recent maintenance or equipment replacement. Besides this a heatmap allows users to see an overall picture that may get lost in discrete anomalies and also identify cause effect relations from the lead/lag behavior of anomalies. This can further accelerate root cause analysis.

The heatmap view paints the color gradient from deep blue to bright yellow as increasing anomaly severity with shades of orange — yellow being the highest. Each row here represents a different signal and time flows from left to right in the same way as for time charts. Seen in Figure 7 is an anomaly heatmap of the finishing mill over a two week period — from January 15th to January 31st. At the top of the heatmap is a dynamic caption that names the signal, the time period and mean anomaly score in the selected grid cell. From the heatmap, it's evident that two signals — Delivery Steckel recoup temperature and gas flow have been consistently anomalous for a long time and so it will be useful to investigate the Delivery Steckel processes.

Shreebhooshan Badrinarayan / Falkonry, Inc.
Order # 132176 / Order Date: 6/1/2023
Copyright Association for Iron & Steel Technology

Figure 7. Anomaly heatmap of the finishing mill process.

C. Analyzing the Signal Distributions to Perform Root Cause Analysis

Falkonry Insight provides tools to speed up the analysis of anomalies with its reporting interface. This reporting interface is capable of generating comparative and trending visualizations for any given set of signals and time periods. Trending visualizations display time plots to help users understand the operational context along with desired min-max ranges or and 3-sigma values or set custom thresholds to understand the signal behavior. Comparative visualizations can be used to compare the value distribution of different signals for a given time period to other time periods or to the value distribution of other signals. User feedback suggests that comparative understanding of signals is often the most important tool in the quiver of maintenance engineers. A detailed analysis for any anomaly can be generated with a single user action. This one-click report generation can be accessed through the anomaly detail view page. This automated report captures information about the anomaly in the coversheet which acts as a note-taking document for the report. Secondly, a timeline chart is populated with the line chart depicting signal mean values for the anomaly period. Thirdly, a comparative value distribution view is also populated for the top 4 signals where each plot area compares behavior of that signal between the anomaly period and one week before the anomaly period.

In Figure 4, one can see the “Create Report” button, which can be used to create a report specific to the anomaly to analyze the Delivery Steckel signals. One of the options after creating the report is to include any contextual signals to the timeline view to understand what was going on during the anomaly. This is a very common need as production recipes, lot numbers, control settings, and other such parameters help understand cause and impact of an anomaly. In Figure 8, while there’s some deviation in air flow and thermocouple temperatures for the Delivery Steckel, they don’t stand out as clearly as the gas flow and recoup temperature identified in the anomalous period.

Shreebhooshan Badrinarayan / Falkonry, Inc.
Order # 132176 / Order Date: 6/1/2023
Copyright Association for Iron & Steel Technology

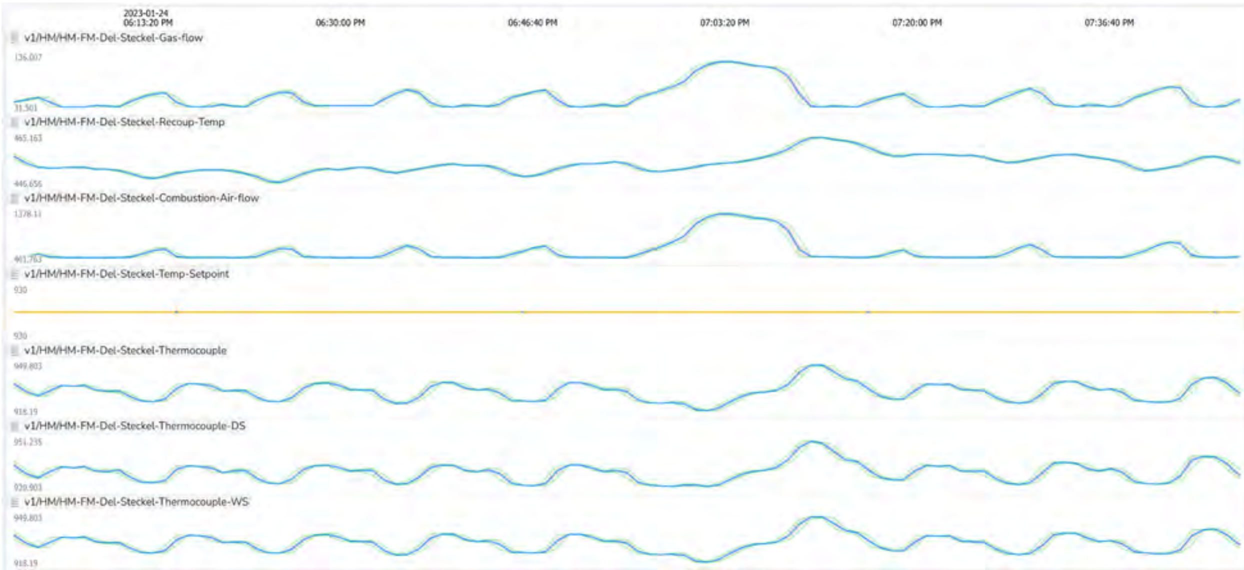

Figure 8. Anomalous and contextual signals visualized in a timeline view through the reporting interface.

The signal value distribution (SVD) analysis is in the form of a violin plot that demonstrates the range of signal values on x-axis and density of the distribution via y-axis. While an auto generated analysis automatically identifies useful SVDs, users can add or make changes to the signals being compared and the time periods being compared. After determining that gas flow and Recoup temperature have an unusual behavior, the violin plots help users visualize the difference between the typical operational range distribution and the anomaly period distribution.

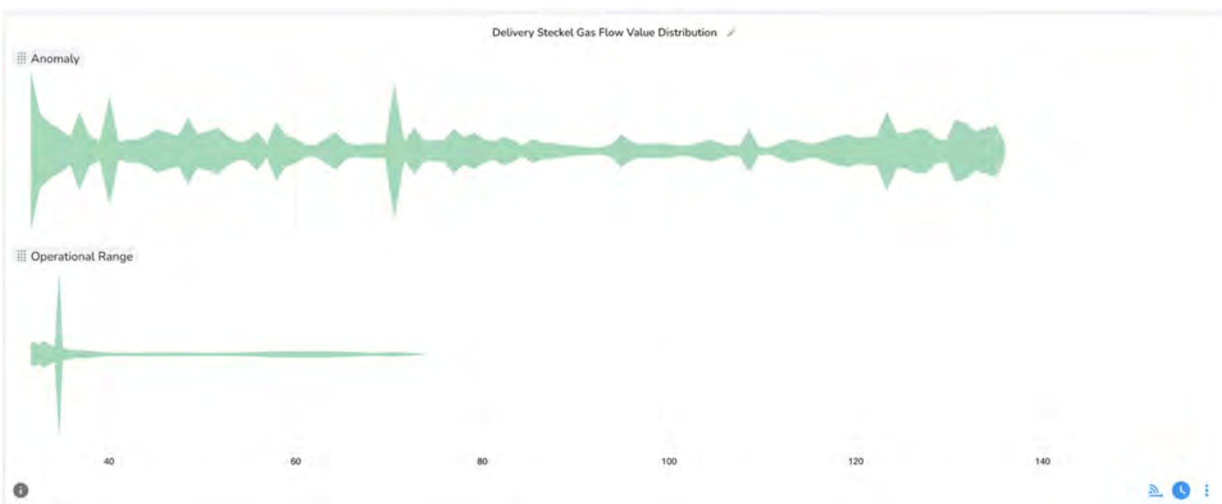

Figure 9. SVD of Delivery Steckel gas flow in anomalous period compared against the operational range.

© 2023 by the Association for Iron & Steel Technology.

Shreebhooshan Badrinarayan / Falkonry, Inc.
Order # 132176 / Order Date: 6/1/2023
Copyright Association for Iron & Steel Technology

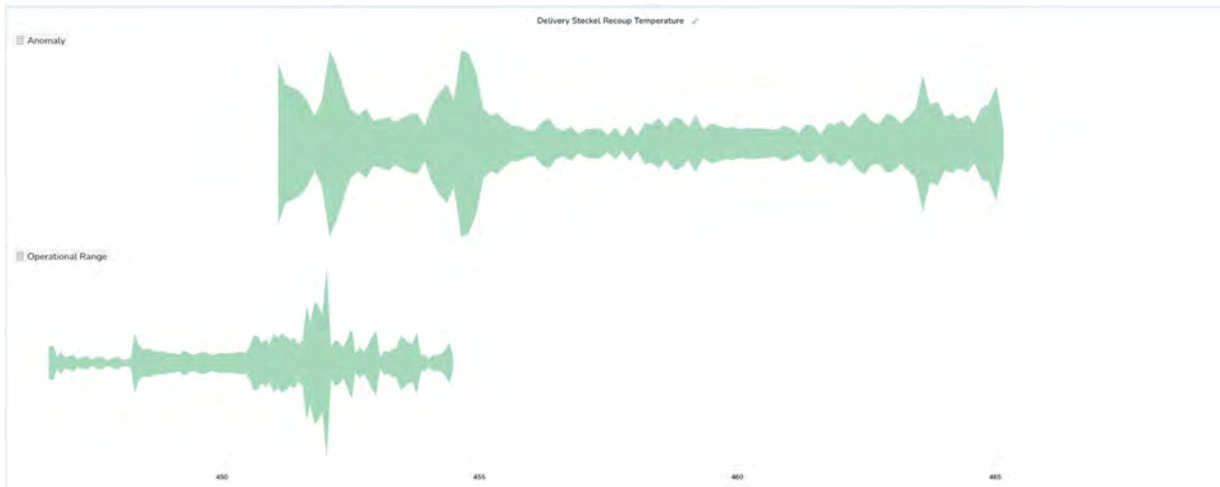

Figure 10. SVD of Delivery Steckel recoup temperature in anomalous period compared against the operational range.

As Figure 9 shows, the value distribution for the gas flow goes up 1.5X in the anomaly period. Similarly, for the recoup temperature, Figure 10 shows that there is a nearly 10 degree increase in the recoup temperature than its operational period. Besides these conclusions, users find SVDs to provide additional insight from the value distribution patterns of different operational states.

CONCLUSION

In this paper, we discussed the challenges in utilizing all available industrial automation data for the proactive discovery and diagnosis of operational problems occurring in the mill. We presented a novel self-supervised approach for the automated analysis of the time series data that equips the operations and maintenance team with insights required for proactive troubleshooting and root cause analysis of issues. This approach works with existing data and technology infrastructure prevalent in steel manufacturing and eliminates the upfront setup and training efforts required to process time series data at the required scale. The presented example from the finishing mill demonstrates the following features and benefits of the automated time series AI:

- The automated no-setup approach enables manufacturers to leverage existing mill data for smart analytics, without requiring costly IoT investments and time-consuming data set-up efforts.
- This approach is not just limited to detecting known excursions and failures, but also alerting the operations team of unknown and novel asset behaviors that could potentially lead to critical operating conditions associated with productivity losses
- Automated AI provides the operational and maintenance users with a detailed view of anomalies, their severity and causal factors, enabling proactive diagnosis and intervention across the plant
- The automated and rapid learning of baseline operating behavior of assets and processes allows robust anomaly monitoring workflow as the equipment and operating conditions change over time, without requiring any data science and model tuning efforts
- Finally, Automated AI is well-suited for diverse applications in the steel manufacturing process that involve time series data — from those with a few signals to those with hundreds of fast-moving time series signals. This enables the plant-scale deployment of automated anomaly detection without requiring specialized operational or data science knowledge

REFERENCES

1. K. Zope, K. Singh , S. H. Nistala , A. Basak, P. Rathore and V. Runkana, "Anomaly Detection and Diagnosis In Manufacturing Systems: A Comparative Study Of Statistical, Machine Learning And Deep Learning Techniques," in *Annual Conference of the PHM Society*, 2019.
2. H. Ding, R. X. Gao, A. J. Isaksson, R. G. Landers , T. Parisini and Y. Yuan, "State of AI-based monitoring in smart manufacturing and introduction to focused section," *IEEE/ASME Transactions on Mechatronics, pp. 2143-2154, 2020*.

Shreebhooshan Badrinarayan / Falkonry, Inc.
Order # 132176 / Order Date: 6/1/2023
Copyright Association for Iron & Steel Technology

3. C. Waters, R. Talla and K. Bhanushali, "Time Series AI for Anomaly Detection and Diagnosis in Steel Production" in *AISTech Annual Conference, 2022.*

4. N. Mehta, C. Singh, V. Toroman, and D. Kearns, "Automated, fast multi-timescale, time series anomaly detection for industrial data with Time Series AI" in *Advanced Process Control Smart Manufacturing of Semi.org, Abstract #22025, 2022*.

5. C. Waters, "AI implementation challenges in steel manufacturing", Falkonry Inc. 2022 [Online] Available: https://falkonry.com/blog/ai-implementation-challenges-in-steel-manufacturing/

6. N. Mehta, "AI is key to broad manufacturing productivity increase", Falkonry Inc. 2022 [Online] Available: https://falkonry.com/blog/ai-is-key-to-broad-manufacturing-productivity-increase/

7. D. Kearns, "Deepdive into Falkonry's unique tech capabilities", Falkonry Inc. 2022 [Online] Available:https://falkonry.com/blog/deep-dive-into-falkonrys-unique-tech-capabilities-with-cto-dan-kearns/

8. E. LaMalfa and G. LaMalfa, "Unsupervised Anomaly Detection in Time Series with Convolutional-VAE" in Journal of Mathematics and Statistical Science, pp 103-108, 2020.

© 2023 by the Association for Iron & Steel Technology.

Shreebhooshan Badrinarayan / Falkonry, Inc.
Order # 132176 / Order Date: 6/1/2023
Copyright Association for Iron & Steel Technology