AISTech 2022 — Proceedings of the Iron & Steel Technology Conference
16–18 May 2022, Pittsburgh, Pa., USA
DOI 10.33313/386/156

Time Series AI for Anomaly Detection and Diagnosis in Steel Production

Crick Waters[1], Prerna Jain[1], Raj Talla[1], Keval Bhanushali[1]

[1]Falkonry Inc.
Cupertino, CA 95014 USA
Phone: +1 408 461 9286
Emails: crick@falkonry.com, prerna.jain@falkonry.com, raj.talla@falkonry.com, keval.bhanushali@falkonry.com

Keywords: Time Series AI, Machine Learning, Time Series Data, Predictive Operational Excellence, Predictive Maintenance, HRT Motor, Continuous Casting, Reliability, Condition-Based Maintenance

ABSTRACT

Artificial Intelligence (AI) and Machine Learning (ML) techniques have been used to solve complex manufacturing operations problems. Applying ML and AI to anomaly detection and diagnosis at scale, however, has been a significant challenge. This paper discusses how Falkonry's Time Series AI platform leverages ML/AI for automated detection and diagnosis of equipment in steelmaking, detecting precursor conditions to critical equipment failures 3-10 weeks in advance of breakdowns. This methodology is scalable across use cases without the need for data scientists. Precedent detection of novel equipment conditions provides insight into maintenance operations that would otherwise have been missed. Such insights lead to proactive maintenance interventions, thus avoiding loss of production due to unexpected downtime events.

INTRODUCTION

Advances in cloud computing, connectivity, bandwidth, communication standards, and storage have fueled the opportunity for steel manufacturers to extract latent information otherwise hidden in operational time series data. Terabytes of data collected contain latent information that is becoming accessible, and most importantly, interpretable using artificial intelligence (AI) and machine learning (ML) techniques. The continuing decline in the cost of data collection, management, and storage combined with advances in processing speed from the use of graphics processing units (GPUs) is further enabling the application of unsupervised learning at plant scale — effectively creating the ability for plant operations management to have "eyes and ears" everywhere at all times.

Until recently, achieving plant-scale detection and diagnosis of anomalies was not possible. Even the most visionary, well-funded, and technically sophisticated companies are only just beginning to use AI to silently and continuously observe operations[13] for emerging adverse throughput, maintenance, and quality conditions at scale. Until now, these companies have put data scientists to work using historical data for specific, targeted use cases. This data science approach yields small-scale near-term results and can be used to attain high precision and recall from hand-crafted models. Putting those models into operation is non-trivial as operational deployment requires time-series data ingestion, management, ML pipeline, results application layer, and integration into a plant-side solution.

Even more challenging is the continuously changing physics of plant equipment. Changes in the product being made or parts being replaced, not to mention normal entropy of systems, means that an ever-increasing complement of data science and IT application maintenance personnel must be deployed to overcome native change and entropy. It's like business churn but as applied to AI/ML. Leading manufacturers have recognized that a different approach[2][7] is necessary to achieve true scale of anomaly detection and diagnosis.[8]

TIME SERIES AI AND CHALLENGES

Time Series AI

A single line in a single casting mill can generate over 1,000,000,000,000 bytes of data in a month. Sampling rates of 10, 5, and 1 ms are common, and sampling rates as fast as 100 ns are used for some high-speed operations, making streaming this volume of data *for a single line* tremendously challenging. Consider scaling from a single line to many lines across a steel production facility, from the furnace to finishing, and one begins to comprehend the data volume challenge to applying AI/ML at scale.[12] A mill may have:

- 10 different sources of time series production-related data
- 350,000 tags
- 1,500 monitored asset- or process-specific models to create and maintain
- 36,000 critical assets
- 131,000 general assets

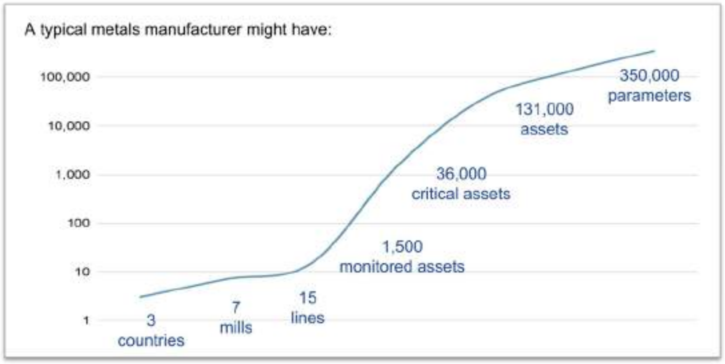

Figure 1. The scale of a steel plant's data generation.

It is evident from the scale of a steel plant's data generation that steel possesses special needs for AI in time series data. Why a special AI for time series data?

- Need for task-appropriate algorithms
- Time series data infrastructure
- Operational usability at the scale of thousands continuously operating, interconnected, and often co-dependent systems of assets

Challenges

Substantial challenges to effecting model creation and maintaining models becomes clear at this scale. Bespoke models and model revisions at this scale involve teams of people, including model monitors, diagnostic technicians, data managers, modelers, and model managers. A single large manufacturer may employ over 200 people dedicated to these tasks.

The top lessons learned by leading metals manufacturers regarding the application of AI/ML are:

1. Behaviors of interest evolve over time
2. Behaviors of interest are very difficult to describe in terms of sensor data
3. Making analytics solutions work at scale requires the operations team to be in the lead

Numerous challenges stand before the innovative steel manufacturer intent on increasing productivity without new capital investment by using existing capital assets and leveraging time-series operational data. One example involves the continuous casting of steel, a well-understood process. Once a caster is put into operation, the manufacturer can do little to change the physics of the caster that would improve production over rated throughput. Rated throughput, however, is not continuously attainable in real-world applications because of lost production time.

Despite the maturity of continuous casting equipment and casting operations, valuable production time is lost to unscheduled downtime events in real-world applications. A high-speed continuous caster producing 150 tonnes per hour can generate $2,500,000 per day in production revenue. Therefore, a single day of lost production is equivalent to $2.5M in top line loss. A manufacturer can release tremendous stored potential value by eliminating unscheduled production downtime.

Casting molten steel, not surprisingly, is hard on heavy equipment. Components wear under harsh conditions leading to failures or poor product quality. An opportunity exists to apply machine learning techniques to detect early evidence of conditions leading to equipment failure. Early detection of such conditions is a warning to maintenance and production scheduling managers that downtime needs to be scheduled for repair before failures occur. The first challenge is to develop a machine learning application that can detect those conditions preceding the equipment failure.

The next challenge is deploying ML applications into a production environment. Sensor and machine parameters are often available in SCADA, PLC, or data acquisition systems such as iba.[10] A Time Series AI platform uses these operational data to detect operating conditions observable to the AI. Time Series AI platforms may use recent or historical data to develop models that can detect specific failures or operational modes indicating pre-failure conditions. Operations practitioners need to be able to deploy those applications into continuous operation without costly and time-consuming software development and then evolve them quickly over long periods of usage. The Time Series AI platform must fit within the steel manufacturers' overall data and operational management software architecture.[4][9] It must receive raw data while also providing insights to manufacturing management systems such as asset performance management (APM) or computerized maintenance management systems (CMMS).

Creating and deploying ML/AI applications for detecting and predicting equipment maintenance requirements is further complicated by two main issues:

1. relatively infrequent actual problems and
2. changes to operations over time.

Industrial operations insist on maintaining low variance and long mean time before failure (MTBF). As a result, problem conditions tend to be rare and unique - meaning that what one believes today to be a "complete" design will invariably miss a future, previously unknown, failure mode.

Over the ordinary course of plant operations, maintenance will be performed, and operating conditions will change. Replacement parts may have different characteristics than original parts. Operating modes will evolve, and products will change as raw material sources shift to accommodate market, demand, and product dynamics. Furthermore, both of these challenges mean that ML/AI approaches must be self-regulating enough to accommodate change and agile enough to accommodate required changes. These features require Time Series AI to be capable of discovering conditions on its own while being managed by operational experts, not data scientists. This is one of the most challenging requirements for Time Series AI.

Time Series AI approaches can be overwhelmed by data and computational scale.[12] Though one can brute-force model development with custom data science methods, the more significant challenge is scaling models across the hundreds of diverse applications from casting to milling to finishing and including the dozens of failure modes of each, along with the different root causes of these failures. The computational scale of Time Series AI must accommodate thousands of parameters, terabytes of data, and hundreds of assets per production line. A scalable Time Series AI platform must provide a usable, consistent methodology that can be applied against use cases from rotating to static equipment such as winders, rollers, and caster components.

Achieving Scale

One of the more pernicious problems faced by manufacturers attempting to apply AI to their time series data is the focus on historical events as the basis for investment justification. Manufacturers are pragmatic in their approach to solution development. The problem statements at the genesis of investment into Time Series AI are rooted in historical fact even though the investment thesis is one of fact-based prognostication. Manufacturing leadership wishes to avoid repeating past, costly events. Classically, these events, such as failures that had unexpectedly taken a caster out of service, become the focus of the investment thesis. Though seeking to create models that will be applied to future operations based on past experience is a logical and technically supportable approach, as noted, this approach may not be scalable, and further, invariably, the approach is prone to the vagaries of the future, with either entropic or business-driven dynamics degrading retrospective AI models.

One method of achieving scale within the economic, technical, and human constraints of mode-specific modeling combines two components: Unsupervised Automated ML[1] and Starting in the Now.

Unsupervised Automated ML

Unsupervised machine learning uses algorithms to automatically analyze and cluster unlabeled data. These algorithms discover hidden patterns in data without the need for human intervention (hence, they are "unsupervised").[5]

Supervised learning trains algorithms to recognize these specific past events which requires accurate knowledge of the past events. If the past event reoccurs in the future just as it had in the past, a supervised ML model will recognize that future event.

In contrast, unsupervised learning does not require volumes of accurately labeled historical events. Instead, unsupervised machine learning applies algorithms to data as it is received. Unsupervised learning algorithms develop a model of equipment operation based on patterns (waveforms or waveshapes) that are observed in the sensor measurements of physical equipment as it operates. These models, applied to continuously arriving data, compare what is happening "now" with what has been previously learned by the algorithms.

Starting in the Now

A substantial challenge to applied AI at scale in steel manufacturing is the sheer volume of information. A hierarchy might appear as:

- Country
    - Region
        - Plant

- Mill
    - Line
        - System
            - Asset
                - Component
                    - Sensor

When one considers the millions of sensors contributing to global understanding of local components, assets, systems, and line performance, it becomes clear that an unsupervised approach to time series AI at scale is necessary. Furthermore, supervised ML necessitates both historical data *and accurate* historical records of temporally aligned events, failure modes, and root cause analyses and validation using a historical record. Records with sufficient accuracy rarely exist, and certainly do not exist at scale.

Starting "in the now" with unsupervised machine learning enables steel manufacturers to deploy AI at plant scale without the need for historical data, accurate records for labeling, or for teams of data scientists to craft the thousands of models that would be necessary for hand-crafted mill-wide AI insights into mill health.[6]

TIME SERIES AI METHODOLOGY

To address the aforementioned challenges, Falkonry's Time Series AI methodology includes two sections: a Machine Learning Pipeline and a Service Infrastructure.

Machine Learning Pipeline

In the machine learning pipeline, analysis (learning and inference) and application development occur in the context of a datastream. This concept unifies both historical and real-time data for analysis. A datastream combines signals and entities into an abstract structure representing a use case or a class of equipment. Each entity in a datastream represents a distinct object in that class (*e.g.* compressor) and the approach allows applications to target multiple similar equipment or use cases to be solved with a single application. The analysis of a datastream comprises three distinct stages of machine learning employed in a sequence, as shown in Figure 2, which is the internal structure of an ML application.

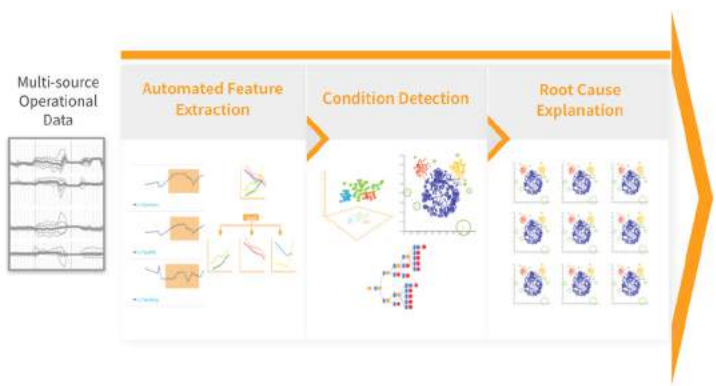

Figure 2. Analysis pipeline inside an ML application.

The output of a Time Series AI application is an "assessment." An assessment represents the current condition of an entity in that application and is produced at a desired time interval. The condition classification assessment provides a condition label, the confidence in that label, and the explanation scores for each of the signals that form a part of that datastream.

The three stages of the ML pipeline follow:

Automated Feature Extraction

Time Series AI extracts information from signals produced by the operation of systems and their effect on the resources consumed in the operation. This operation is observed in signals that are collected by the Time Series AI platform. The platform detects patterns present in such multivariate data. An operational pattern of an entity at time t comprises the trends in multiple signals of that entity as seen over a period of time leading to at a point t. We can treat this pattern as a condition.

Falkonry's Time Series AI discovers operational patterns in automatically extracted signal features. An unsupervised approach to learning features is used, with the goal of producing a feature vector at each required time t to represent the pattern at time t without knowledge of downstream clustering or classification. The feature vector should consider shapes and waveforms in each of the signals and minimize dimensionality for efficient processing. The feature extraction approach adapts to vagaries of industrial data production such as irregular sampling, multiple independent data sources, noise, gaps and compression as well as different data types - numeric and categorical.

This stage is based on a composition of standard signal processing techniques and leverages adaptive windowing to optimize the use of data history. The stage can take advantage of signal metadata provided by the user (such as data type or maximum allowed sampling interval) but does not require algorithm selection by the user. This design approach to automated feature extraction eliminates expensive data preparation efforts that can be repeated for each application.

Condition Detection

This stage aims to produce a condition value at every required time *t* based on the feature vector produced by the automated feature extraction stage for that time *t*. The reported condition at time *t* may be one of the following: 1) a user-label, 2) a system-generated label, or 3) unknown. To do so, it uses two additional pieces of information - labeled events and the desired degree of generalization. Each labeled event included for condition prediction, also called a fact, carries a start and end time as well as a condition label. Facts may be either ground truth or hypothesis to be tested. The desired degree of generalization controls the degree of tightness of match between supplied labeled events and patterns arising at other times and results in a choice between high reliability or high predictability. This approach is referred to as semi-supervised learning and produces a condition value even when no example data is provided.

During this stage, feature vectors are first clustered dynamically and a cluster identity is attached to every feature vector. This first step does not depend on the time order of feature vectors and can be called time-free. In the second step, feature vectors supplemented with cluster identity are combined with facts in a classifier to produce a percentage match for each feature vector to the feature vector for each of the fact labels. The result of this stage is further resolved down to a single condition label by using the generalization factor and to isolate anomalies into the unknown condition. The highest match percentage is used as the confidence in the label of the condition value.

The conditions predicted by Time Series AI can be selectively reviewed and labeled with the help of experts from the operations team as this stage is repeated. In this way, experts' domain knowledge is digitally recorded and available for future applications.

Explanation

The previous stage outputs most of the condition assessment for the datastream. For every assessment, the datastream also provides a measure of each signal's contribution to making that assessment's condition label (called an "explanation score"). Explanation scores range from 1 (highest contribution) to 0 (no contribution) to -1 (contradicts the assessment but was outweighed by other signals). Explanation scores are calculated per assessment using the condition label and the feature vector from each signal (internal to the condition models). For each signal, the feature vector ("assessment point") is compared against a sample of the feature vectors used during model learning ("sample points"). The algorithm draws a neighborhood boundary around the assessment point to find the nearby sample points. Among these nearby neighborhood sample points, the ratio of points with the same, vs. a different, condition as the assessment is found. This ratio is calculated for all the sample points (not just the neighborhood). These two ratios are used to calculate the explanation score.

Service Infrastructure

The Time Series AI platform is designed for easy deployment into IaaS (Infrastructure as a Service) environments like AWS or Azure or into privately managed compute and storage infrastructures running a container-compatible Linux operating system. Figure 3 presents the overall architecture of the application deployment.

Figure 3. Application deployment architecture.

A time series AI platform is composed of multiple architectural and functional elements. These components are orchestrated as a system. The system includes branches of functionality including:

- Data ingestion and signal processing
- Deployment interfaces from cloud to air gapped environments

- Hierarchical metadata including asset, event type, process, location, sensor type, and source system
- Integration agents to facilitate complete integration with plant workflow systems (e.g. CMMS, APM, etc.)
- Machine learning and AI technologies and associated pipeline management system
- Application layer including workflow management
- Reporting features

A representation of a time series AI platform composition is illustrated in Figure 4.

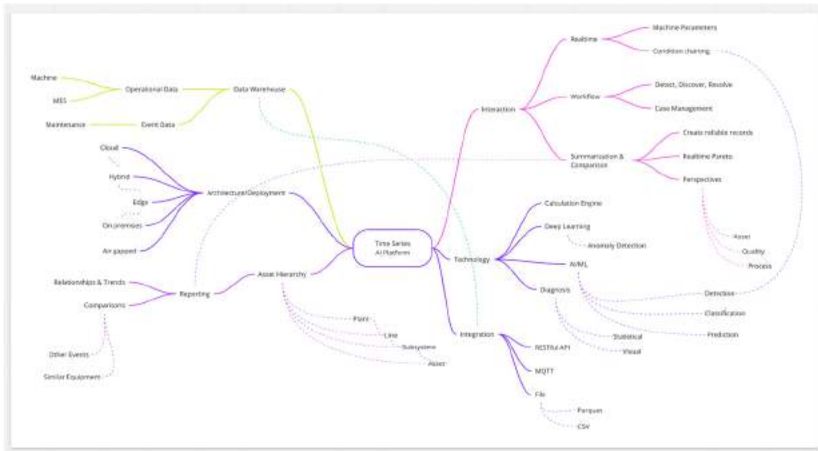

Figure 4. Features map for AI-based operational excellence.

HOT RUN TABLE (HRT) MOTOR ANOMALY

Case Study Description

Run out, hot run table, or cooling tables move steel plates and steel strips from one process operation to another. Hundreds of motor-driven rollers move steel plates and steel strips from point to point within the mill. These motors are driven by variable speed drives. All of the motors, couplings, and support bearings must perform correctly at all times because an incorrectly performing motor will introduce surface defects onto the steel plates and steel strip.

As a steel plate or steel strip moves, individual motors cycle on and off with the presence or absence of the material. Unit mass of a steel plate or steel strip and the roll-to-slab contact dimension, which depends on its width, affect motor current as it passes. The shape of the motor current waveform of the first motor to be in contact with the material will be different from that of the last motor to be in contact with the material. As the material passes, the changing motor currents create time-series waveforms.

The objective of using time series AI for HRT motors is to detect anomalies in the waveform before the motor-roller condition deteriorates to the degree that quality defects are introduced onto the steel strip or steel plate surface. With early detection of a pending quality-impacting motor condition, plant operations can take an errant motor out of service before quality defects can be introduced.

The challenge, in this case, is that the AI models, if hand-crafted using supervised learning methods, would require large historical data sets complete with motor-by-motor failure histories annotated with high temporal precision. Unfortunately, this level of information and data precision is neither available nor attainable in practice.

To solve the lack of data and failure history problem, a time series AI platform capable of unsupervised learning is used. One example of such a time series AI platform is that provided by Falkonry Inc. It quickly learns the normal patterns of motor current from small amounts of data as it becomes available from operating motors. In this way, a manufacturer can “start in the now” allowing the AI to first learn expected motor current patterns and then continuously compare these patterns with the ever-evolving “now.” When a motor’s behavior deviates from the expected, its behavior is anomalous — even if not detectably so by human observation. The time series AI platform will then detect abnormal equipment conditions as they emerge, leaving sufficient time between detection and quality-impacting failures for plant operations to avoid losses.

Precedent Detection

Operations teams need to be alerted of anomalous conditions before those conditions escalate to the point of failure. Detection of production-impacting failures is often obvious - such as motors tripping off line, strips breaking due to poor welds, or foldovers due to mill-sheet misalignments. The overall objective, therefore, is to detect precedent conditions to failures. Precedent detection typically occurs before humans can observe condition degradations. Due to both the complexity and subtlety of interaction between monitored parameters and because of the scale of production operations, it is impossible for humans to monitor the many hundreds of parameters that reflect numerically the physical condition of production HRT motors.

Hot run table motor and roller failures take several different forms. Some examples are:

- Guide rails improperly set causing motor overload and trip
- Coupling key shearing
- Motor trip due to overloading/overheating
- Stuck Rollers

Available data

One challenge to detecting conditions preceding roller failures is that limited data are available. Motor current and motor drive speed are commonly the only two available signals for HRT motor-roller sets. An HRT may have two hundred such rollers driven by multiple electric motor drive units with each drive unit powering a subset of the table's motors.

Table 1. Available HRT Motor Signals

Parameter	Sampling Rate
Motor current	1 second
Electric drive speed	100 ms, sample and hold

The waveforms for motor current and electric drive speed (DCSpeed) have a typical shape as depicted in Figure 5 below:

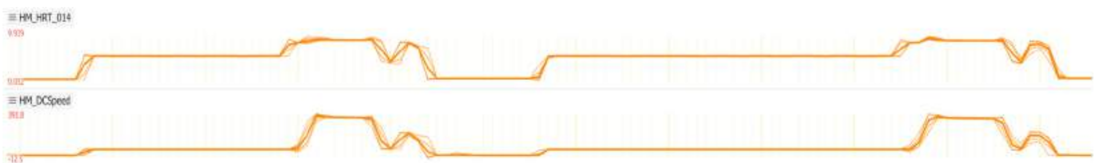

Figure 5. Two plates passing an HRT motor.

Workflow

The workflow to employ an automated time series AI platform begins with mapping available signals (or tags) to assets followed by streaming those signals to the AI platform. The sequence is:

3. Identify available signals
4. Establish a stream of real time data to the time series AI platform
5. Let the AI observe asset operation over a period of time
6. Observe that the AI "learns" the expected patterns of operation
7. Compare detected operating anomalies and observed equipment failure events
8. Explanation and Resolution
9. Monitor for anomalous events

Identify Available Signals

In the hot rolling table case, the only available signals are motor current and electric drive speed. In this application, the parameters are being collected in an *iba* system. Motor current signals (tags) are mapped to their respective drive motors. In this application, four electric drives control nearly 200 motors. Each motor will have one motor current signal and one electric drive speed signal.

Establish a Datastream

The Falkonry time series AI platform has the ability to ingest periodic files transferred automatically as "bursts" of streaming data, or continuously via an MQTT broker. Files may be in CSV, JSON, or Parquet formats. Falkonry automatically handles different sampling rates and missing values. As data is ingested, it is enhanced and prepared for historized and real-time visualization. Values, means, mins and maxes, and standard deviations are plotted as can be seen in darker and lighter lines on a parameter plot (Figure 6).

Figure 6. Visualizing signal information at a given time period in Falkonry.

Let the AI Observe

During this step, the time series AI automatically featurizes the incoming data stream then applies ML modeling algorithms to classify periods of time into operating conditions. During periods of time in which the shapes of signals are common, the AI will assign a system-generated classification label. This is unsupervised learning. Over a short period of time (days or a week), the AI will learn a sufficient number of operating patterns. This completes the bootstrapping period. The system then actively compares real time patterns to the patterns that it learned during the bootstrapping period.

Compare Detected Anomalies and Observed Events

Once the system has transitioned out of bootstrapping, it will announce the detection of anomalies. If the roller motor signal pairs exhibit waveforms that differ from what the system observed during bootstrapping, a “novel condition” alert is issued.

As plant operating conditions change, plant conditions are compared in real time with the system-learned and expected conditions of operation. When the AI system discerns a condition that has emerged that is distinctly different from system-learned conditions, a “novel” event alert is generated.

Explanation and Resolution

Novel events are detected when the signal shapes form patterns that are not represented in the learned waveforms. The AI platform “sees" these anomalous periods and provides notification of the occurrence of these novel events. Plant operations team members are provided a set of diagnostic reports explaining what the AI discovered that distinguishes the event period from other known (learned) conditions. Event reports include signal plot comparisons highlighting the signals that most explain the event, signal and frequency value distribution plots, and signal contribution comparison charts.

Figure 7. A HRT novel event signal trace.

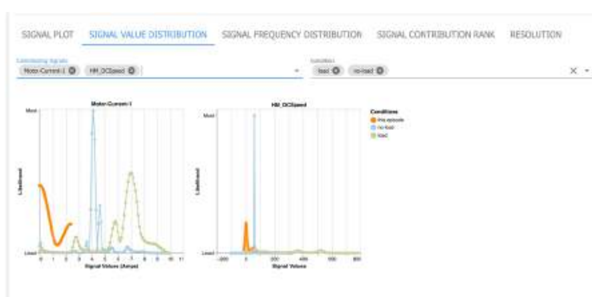

Figure 8. Signal Value Distribution plot of an HRT novel event.

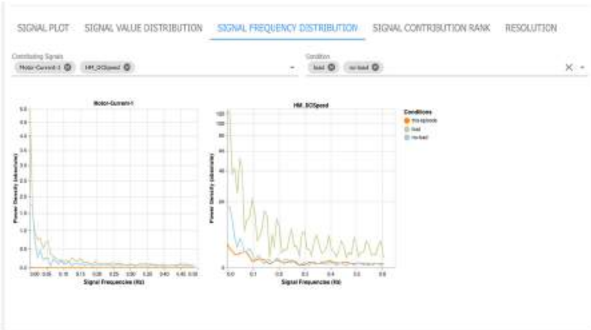

Figure 9. Frequency Value Distribution plot of an HRT novel event.

Figure 10. Signal Contribution Rank of an HRT novel event.

Monitor for Anomalous Events

After a short time, the time series AI platform will have learned the expected behavior of the operating equipment. The machine learning model is then applied live in real time to newly arriving data for each of the monitored HRT motors. The AI platform will constantly compare the patterns of operation evident in the "now" data with the learned, expected time series data patterns. When an unrecognized pattern emerges, it will be detected and reported as an anomalous event. Operations team members are then alerted to the occurrence of anomalous events the moment they are detected. When the HRT roller motor conditions are as expected, no alert is indicated. The blue marks in Figure 11 indicate periods of operation during which a motor exhibits a condition that has previously been seen by the time series AI platform, and therefore is "predicted" to be a "warning" condition.

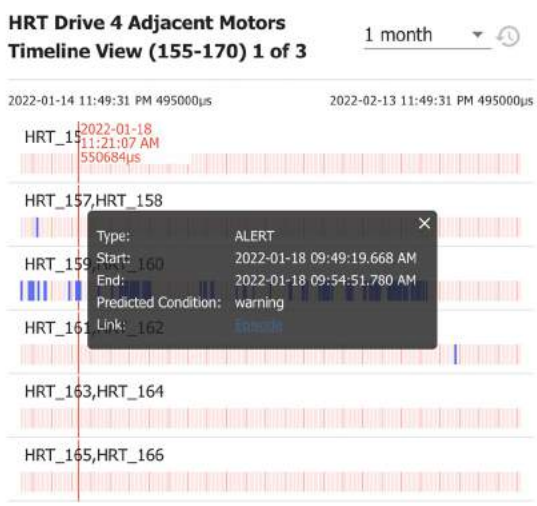

Figure 11. HRT motors timeline view.

Evolution of HRT Motor Monitoring

HRT motor monitoring presents a more challenging case than many other applications in metals manufacturing. The challenge comes in two forms:

1. There is only a single unique signal per motor: current. The only other signal available is a "shared" signal, electric drive speed. The electric drive speed signal is not unique, therefore, to any individual motor.
2. Motors are mechanically coupled when a steel strip or steel plate passes. As much as the motors drive rollers that move the material, the material's mass and rigidity mechanically couple adjacent rollers and can, therefore, influence the motor current waveforms of adjacent motors

To overcome the limitations of univariate motor current parametrics and near-neighbor mechanical coupling, a "paired" motor entity approach is taken. Instead of applying the time series AI to a single motor with one current parameter (and one shared speed signal), pairs of motors are modeled simultaneously. For each pair of motors, two current signals and one electric drive speed signal are used.

By evolving the AI from a single motor univariate analysis to a bi-motor multivariate approach, significantly more information is available to the classification engine. Features from both motor parameters better indicate the differential conditions between motors despite, or because of, mechanical coupling.

One example of bi-motor analysis is illustrated in Figure 12. In the zoomed out view on the left, the shaded green region identifies an AI-detected anomalous condition. In the zoomed in view on the right, the shaded green region reveals the signal differences between the two motor currents and reveals a trip condition in motor 2.

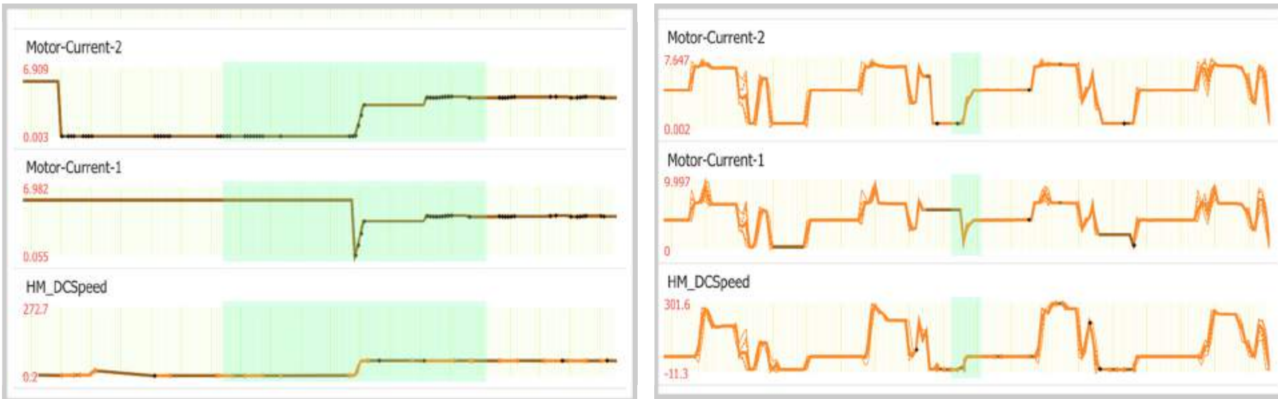

Figure 12. Bi-motor analysis signal plots.

Advantages

The advantages of using a time series AI platform applied to HRT motors were:

- Simple to use data integration architecture enabled connectivity and streaming of large data quantities from multiple sources with different sampling rates.
- Unsupervised automated machine learning of pattern recognition algorithms create and deploy into production ML models through bootstrapping against data as it is received (no historical data is needed).
- ML applications were created from multivariate data with various, irregular sampling rates, without data engineering effort, for equipment with dynamic operating modes. The automated, consistent, and repeatable AI methodology provided a platform and application layer that delivered to non-data scientist users a production-ready ML solution.
- Time Series AI incorporates both automated unsupervised methods and provides for semi-supervised ML techniques enabling both discovery of novel conditions and recognition of specific conditions. This method facilitates both learning the patterns of known operating conditions from previously captured operating data but also auto-identification of unexpected and novel conditions in real-time operational data.[11] Therefore, this learning system provides benefits for a wide range of use cases over time without needing data scientists to develop or maintain these applications.

The Time Series AI platform handled many different types of sensors, non-preprocessed and irregularly sampled data with multiple time scales, worked with minimal or even no ground truth data, and created condition assessments with cause explanations that were found interpretable to the degree that operators could make timely decisions to improve their production systems. These applications run in realistic environments and provide usable real-time insights that address business needs while overcoming the challenges associated with machine learning, ease of adoption, scalability, and deployment.

Results

Automated time series AI provides continuous monitoring of plant equipment providing visibility into equipment condition. The pattern detection and classification system can be applied to a wide variety of steelmaking equipment. By detecting the emergence of novel and known anomalous conditions, the times series AI platform delivers reports in the form of alerts to operations personnel. The onset of anomalous conditions prior to equipment failure enables operations personnel to rapidly diagnose equipment conditions leading to repair-or-replace decisions before unexpected equipment failures cause outage-induced losses.

Table 2 identifies a selection process and equipment on which automated condition classification based on patterns detected in time series data has been applied. The time between the detected emergence of anomalous patterns and either proactive actions taken or failures incurred varies with factors including the product being produced, equipment maintenance history, and the operating history of each equipment type influence the mean time to action/failure.

Table 2. Sample Applications, Parameters Used, and Results (Where Available)

Process	Equipment Type	Application	Parameters Available	Parameters Used by the AI	Digital Twin Development Time	Identification of Potential Failures
Continuous Casting	Slab Cutting Torch	• Detect anomalous cuts • Prevent iron powder blowback events • Maintain pressure in iron powder reservoir	12	6	12 iterations, 3 weeks	52 instances in 5 months
			12	6	14 iterations, 3 weeks	13 instances in 3 months
	Torch Carriage		13	4	12 iterations, 3 weeks	—
Finishing Mill	Hot Rolling Table Motors	• Detect motor events: decoupling, trip, stuck	200 (total across rollers)	189	4 live twins, 4-6 weeks, 11-49 iterations	31 instances in 1 month
	Finishing Mill Press	• Detect abnormal operations (e.g. high vibration) • Detect misalignment/ position deviation	38	10	3 weeks	—
	Entry Steckel	• Detect abnormal operation (temperature spikes, improper slotting)	30	4	13 iterations, 3 weeks	—
	Delivery Steckel		30	4		

CONCLUSION

In this work, we discussed the challenges of scaled AI deployments and how automated ML in a time series AI platform can be harnessed to overcome those challenges. The greatest challenge to scaling AI for steel manufacturers is the mills' sheer quantity of signals and assets. Custom development of AI models is expensive, and only possible at relatively small scale — and even at small scale, the overhead of creating and managing the application architecture and user interface needed in support of a machine learning model is not only expensive, but not core to the steel maker's business. For these reasons, automated unsupervised ML integrated into a time series AI platform that includes the entire architecture needed for ingestion at terabyte scale, model creation, model inference, and an application layer for non-data scientist operations personnel is a viable approach to overcoming the challenges of scaled AI in steelmaking.

- Automated unsupervised learning can deliver alerts of "novel" conditions as they emerge in operating equipment. Detection of novel conditions draws the attention of operations teams to development to potentially critical machine conditions.
- Unsupervised learning novel events can easily be labeled to create semi-supervised models. The time series AI platform becomes enhanced with the ability to recognize emergence of named conditions and therefore provide specific predictions of the nature of detected anomalous behavior. Further, the system will continue to identify unknown or "novel" behaviors.
- Automated unsupervised pattern recognition is well suited to diverse applications – those with few to those with many associated signals (tags). The flexibility of an automated multivariate pattern recognition approach to machine learning facilitates scaling deployment across processes and lines without the need for specialized operational or data science knowledge.
- A flexible time series AI platform allows for rapid interpretation of model results and flexibility to reconfigure approaches. One example of this flexibility is the ability to evolve models from univariate to paired-asset multivariate models within the same system with no coding or application development.

- Finally, real-time pattern recognition based on patterns emerging as equipment operates results in early detection of novel conditions. These conditions often precede the occurrence of adverse events such as equipment failures or the introduction of quality impacting conditions. The detection of precedent conditions can facilitate initiation of proactive maintenance interventions before substantial production losses are incurred.

REFERENCES

1. "Automated machine learning," [Online]. Available: https://en.wikipedia.org/wiki/Automated_machine_learning. [Accessed 10 January 2022].
2. C. Lee, "The Intelligence-First Path to Predictive Operations," Falkonry Inc., 6 July 2020. [Online]. Available: https://falkonry.com/blog/the-intelligence-first-path-to-predictive-operations/. [Accessed 10 January 2022].
3. C. Lee, "Database-First-vs Intelligence-First: The Cart Before the Horse," Falkonry Inc., 10 June 2020. [Online]. Available: https://falkonry.com/blog/database-first-vs-intelligence-first-the-cart-before-the-horse/. [Accessed 10 January 2022].
4. N. Mehta, "Operational AI: the bridge between Operational Technology and Operational management," Falkonry Inc., 2 December 2020. [Online]. Available: https://falkonry.com/blog/operational-ai-the-bridge-between-operational-technology-and-operational-management/. [Accessed 10 January 2022].
5. J Delua, "Supervised vs. Unsupervised Learning: What's the Difference?" SME, IBM Analytics, Data Science/Machine Learning 12 March 2021. [Online]. Available: https://www.ibm.com/cloud/blog/supervised-vs-unsupervised-learning/. [Accessed 10 January 2022]
6. J. Yoon, D. He and B. Van Hecke, "A PHM approach to additive manufacturing equipment health monitoring, fault diagnosis, and quality control," Nantes, 2014.
7. 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.
8. 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.
9. "Operational technology," [Online]. Available: https://en.wikipedia.org/wiki/Operational_technology. [Accessed 10 January 2022].
10. "iba System," [Online]. Available: https://www.iba-ag.com/en/iba-system. [Accessed 10 January 2022].
11. N. Mehta, "Rethinking Operational AI," Falkonry Inc., 12 May 2020. [Online]. Available: https://falkonry.com/blog/rethinking-operational-ai. [Accessed 10 January 2022].
12. C. Lee, "How intelligence first leads to faster learning and cost effective scaling," Falkonry Inc., 13 April 2021. [Online]. Available: https://falkonry.com/blog/how-intelligence-first-leads-to-faster-learning-and-cost-effective-scaling/. [Accessed 10 January 2022].
13. C. Lee, "Continuous Improvement of an Operational AI Deployment," Falkonry Inc., 18 June 2020. [Online]. Available: https://falkonry.com/blog/continuous-improvement-of-an-operational-ai-deployment/. [Accessed 10 January 2022].

