ei3

ConnectedAI

A Machine Learning Approach To Practical Predictive Insights

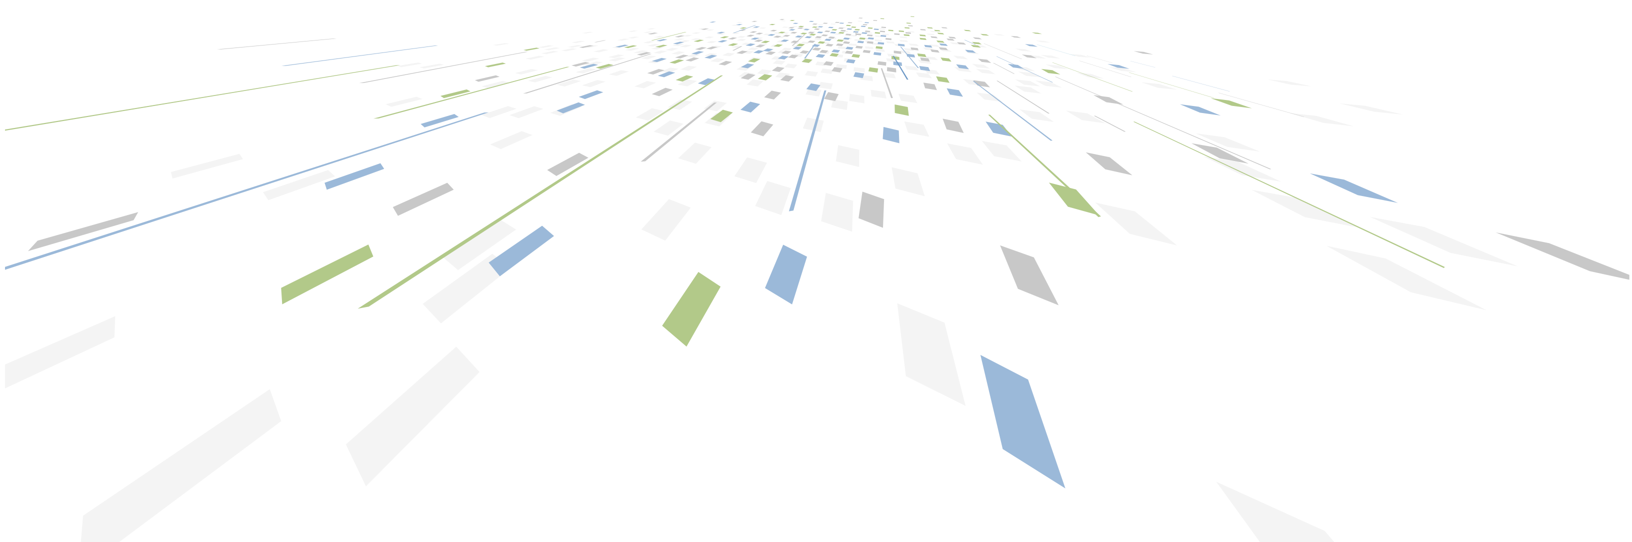

Contents

Revolutionizing Industrial Manufacturing with ConnectedAI

Artificial Intelligence is revolutionizing wear and tear detection on industrial machinery. Applying new machine learning methods on large amounts of process and low-cost sensor data to holistically assess the health of an industrial machine creates meaningful insights into future machine availability.

Over the last years, ei3 has developed "ConnectedAI", a version of centralized federated learning[1] tailored to the needs of the manufacturing industry, to create predictive insights for machine health. When combined with "Symphony", a smart, low-cost sensor capable of detecting and analyzing acceleration, vibration, and sound data emitted from industrial machines, ConnectedAI surpasses traditional vibration analysis while dramatically reducing cost and installation requirements.

This technical paper outlines the application scenarios for ConnectedAI, presents its technical foundation, and explains the physics of the wave field analysis used in the Symphony sensor.

ConnectedAI is a version of centralized federated learning tailored to the unique needs of industrial machinery, making AI a cost-effective and practical solution to predict machine health.

[1] Centralized Federated Learning
S. Abdulrahman, H. Tout, H. Ould-Slimane, A. Mourad, C. Talhi, and M. Guizani, "A Survey on Federated Learning: The Journey From Centralized to Distributed On-Site Learning and Beyond," in IEEE Internet of Things Journal, vol. 8, no. 7, pp. 5476-5497, April 1, 2021, doi: 10.1109/JIOT.2020.3030072.

1.0 I Introduction

Over the last decade, manufacturers in many industries have adopted "Industry 4.0" or "IoT" tools as a key method for improving production efficiency. These tools use process parameters and sometimes additional sensor data to create standard metrics, such as "Overall Equipment Efficiency" (OEE), to allow operators to assess and track machine performance. ei3's suite of IoT applications stands out as an early entrant into the market in 1999 and today as the leading provider in the plastics and converting industry. Machine owners, plant managers, and quality engineers widely use ei3's no-code applications to track performance and reduce machine downtime, waste, and energy consumption - and subsequently lower production costs.

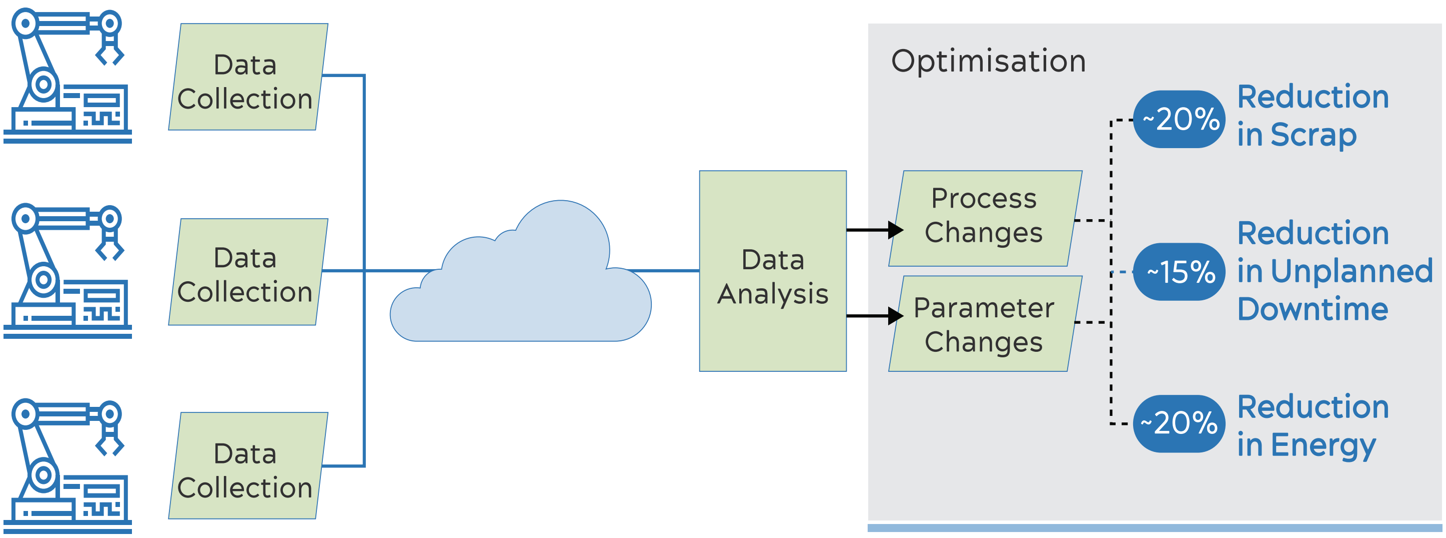

Figure 1: ei3's suite of IoT applications has proven to be easy to install and quick to deliver returns: Many users report reductions in scrap rates, unplanned downtime, and energy.

However, the pursuit of greater production efficiency is an ongoing battle as competitive pressures and new requirements for production sustainability continue to drive industry needs: Shifting into focus today is the need to judiciously use natural resources, the reduction of any environmental footprint, and also the lack of skilled labor in many geographies.

Given the hype - and maybe fear - created by recent advances in AI (such as "ChatGPT"), it is not surprising that approaches based on AI are seen as the next big opportunity to achieve a "step change" in production performance. In this technical paper, we explore ei3's unique approach to AI for industrial manufacturing, which differs from conventional approaches used for predicting machine performance and breakdowns, which are often based on a digital twin.

2.0 I Not all Digital Twins are Created Equal

The notion of the digital twin is popular due to its apparent simplicity and accessibility. The machine designer or engineer formalizes the expected behavior of a machine using mathematical notation, creating its digital twin. Sensors measure the machine's actual behavior, and any deviations of the measurements from those predicted by the digital twin are interpreted as signs of looming failure. Countless lab-bench experiments have been conducted and widely publicized to validate and popularize this approach.

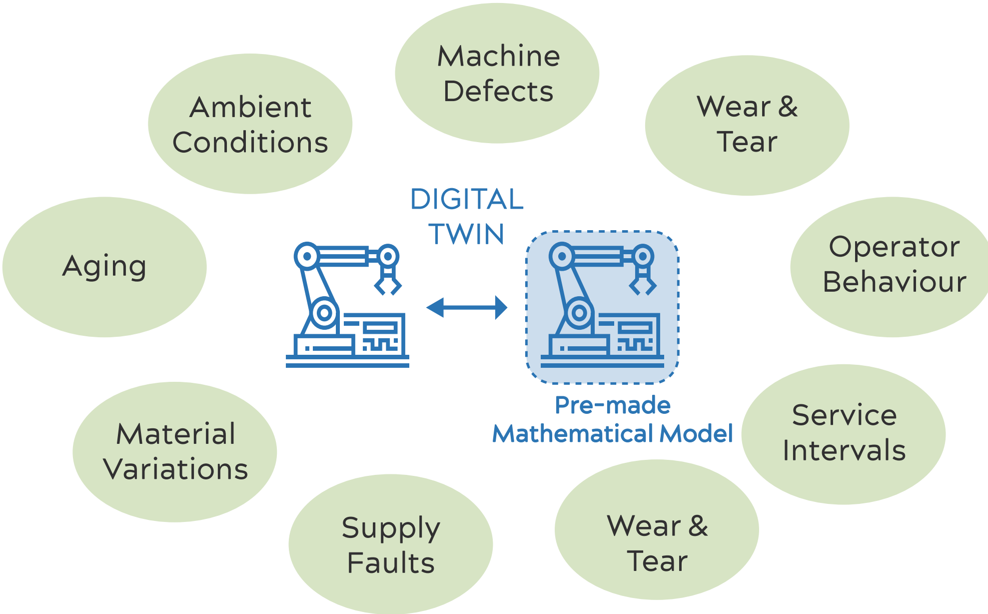

Figure 2: External factors impacting the quality of a digital twin.

However, real machines are not lab experiments: The noisy and erratic nature of real-world machines, compared to their lab-bench-bound cousins, makes this seemingly simple strategy remarkably hard to execute in reality. There are many factors external to the machine that have a dramatic impact on its performance. These are extremely hard to capture in a digital twin. Many digital twins therefore suffer from limited fidelity, rendering them unable to identify machine abnormalities.

Successful real-world examples of the digital twin approach are "health and usage monitoring" (HUMS) solutions, which are commonplace in sensitive applications like the high-powered gearboxes that connect the turbine engines of a helicopter to the rotor blades. The effectiveness of such systems comes from carefully developed mathematical models and precisely positioned sensors that are tuned to key frequencies that pose the highest risk to the gears. HUMS increase safety margins and supplement the stringent maintenance regime already commonplace in aviation, where any failure could lead to loss of life.

However, the complexity of the required sensor systems and upfront engineering makes such solutions cost-prohibitive for all but the most (life-) critical industrial applications. It is worth noting that gearboxes are the most prominent example of HUMS applications in the literature, and that, with all due respect to the intricate detailed engineering involved, gearboxes are comparatively easy to describe mathematically.

Unlike gearboxes, most industrial production machines exhibit behaviors that are less deterministic. Like a powered double pendulum (a “chaos machine”), even the slightest deviation or mechanical imprecision can lead to seemingly chaotic behavior.

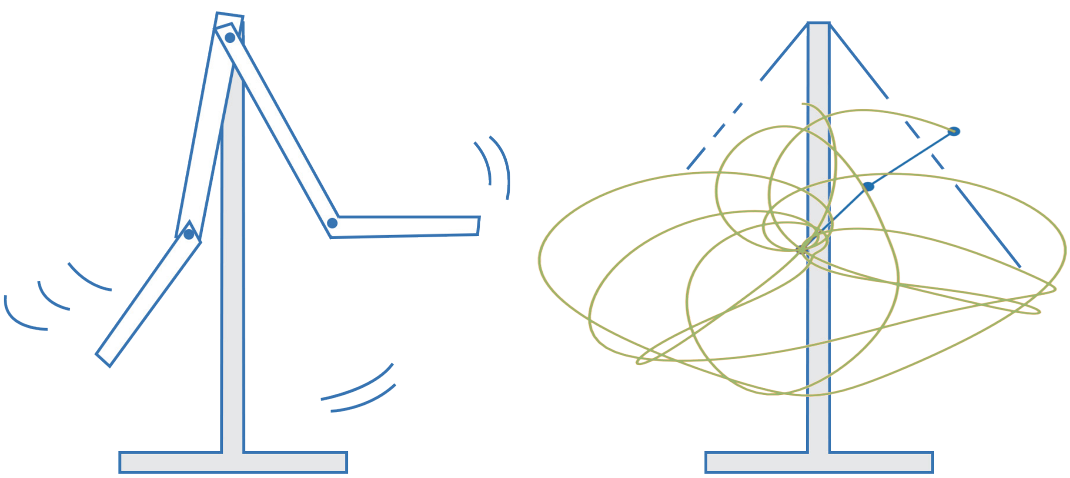

Figure 3: Although industrial machines are controlled by process logic and governed by stringent engineering, small imperfects and changes in usage and environmental conditions add noise and chaos that make them hard to predict accurately. The chaos machine - a double pendulum, is a popular example to illustrate that even simple mechanical assemblies can exhibit chaotic behavior.

Adding to this complexity is the fact that machines are operated by humans who deviate from standard procedures, having both positive and negative effects on machine performance: While experienced operators can sometimes improve performance by applying human ingenuity, mistakes are more common and can lead to part failures and machine issues. This “human impact” is by no means limited to small operations that might suffer from skill shortages or lack operator oversight. Thomas Zurbuchen, recently retired Head of Science at NASA (considered by many the epitome of engineering excellence), reports that 70% of all failures at NASA are due to human error, versus 30% that are caused by technical failures, a ratio that he himself had expected to be the other way around.[1]

[1] Zurbuchen: https://de.wikipedia.org/wiki/Thomas_Zurbuchen

As a result, mathematical models for larger and more complex assemblies - let alone entire machines! - are becoming increasingly challenging and unmanageable due to the multitude of internal and external factors that influence machine behavior, ranging from changes to machine settings or source material to environmental conditions such as temperature or humidity, to the operators applying the “human impact.

Therefore, while the idea of a digital twin for general industrial machinery is obvious and compelling, it quickly meets the harsh realities of a real-world deployment and requires large investments that make it impractical for many applications in the manufacturing environment.

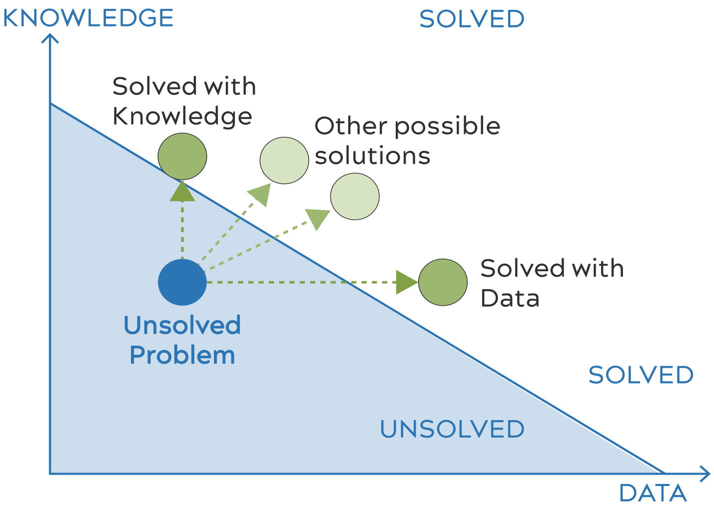

Figure 4: Knowledge and Data are interchangeable in finding a solution to a problem, even though not all solutions are equal. Machine learning has the ability to ingest vast quantities of data, find patterns, and find data-driven solutions, even when knowledge is not available.

Machine learning opens up alternative avenues. Professor Carl Henrik Ek, University of Cambridge’s Head of Machine Learning, illustrates the trade-off between “knowledge” and “data” using a simple graph, showing that complete knowledge (such as the elusive accurate “digital twin”) allows for solving any solvable problem, while on the other end of the spectrum, having access to all the data can also lead to a solution even if fundamental understanding is lacking or cannot be obtained.[1]

Professor Ek admits that combining both “knowledge” and “data” is required for most solutions, and that the quality of the solutions may vary. It, therefore, follows that the ability to collect and process large amounts of data can be used as a practical replacement for a complete and accurate digital twin. “ConnectedAI” is based on this premise.

[1] Knowledge vs Data: “Free Lunch”? How we can Learn from Data with Carl Henrik Ek”, https://www.youtube.com/watch?v=NLzUOnWnH1k

3.0 I Machine Learning and Federation: ConnectedAI Approach

ei3's ConnectedAI approach to industrial data science is a departure from the conventional digital twin approach and is unique in three areas:

INSIGHTS BASED ON ASSESSMENT	ei³ DIGITAL TWIN	FEDERATED LEARNING
Automatically generated on site to create a high-fidelity fingerprint of the machine in operation.	Enables efficient learning of rare events across fleet of machines without compromising data privacy.	Detects and interprets subtle changes to the machine behavior based on full view of machine data.

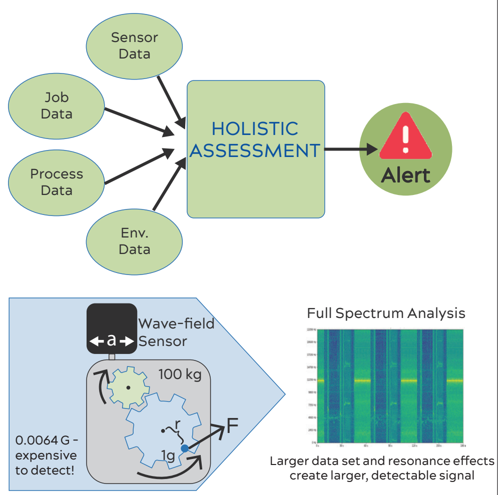

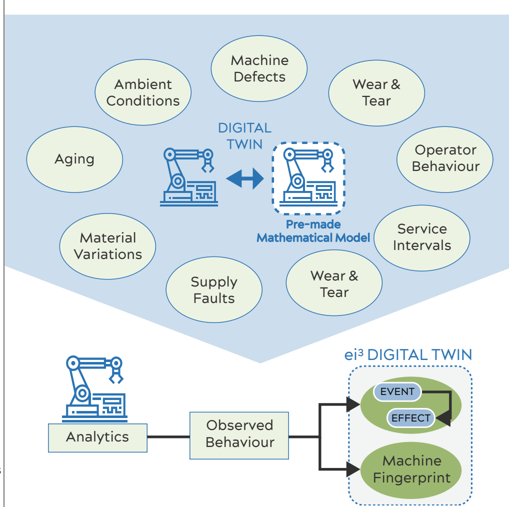

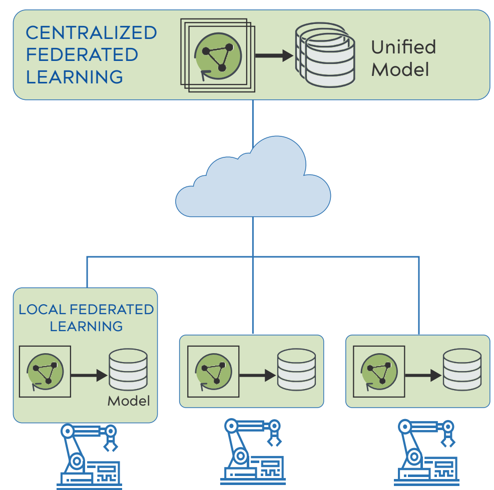

By using a machine learning approach that is widely known in academia as "Federated Learning", ConnectedAI creates results faster and more accurately, and ensures data privacy.

ConnectedAI is practical today because of the ubiquitous availability of cost-effective data collection and processing, which can be implemented in edge-devices embedded in each industrial machine. Together with secure network connectivity this allows complex machine learning algorithms to operate on large amounts of data in a distributed manner. The three architectural pillars of ConnectedAI are introduced and explored in more detail in the following sections.

3.1 I Machine Learning vs. Signal Processing

As described earlier in this paper, HUMS systems can detect failures in gearboxes but require extensive engineering and high-cost sensors to pick up tiny signals. This is because capturing the vibrations created by a failing gear directly is nearly impossible. In this section a typical industrial gear box, such as one used in an extrusion machine, is assumed to suffer from an imbalance defect from wear and tear.

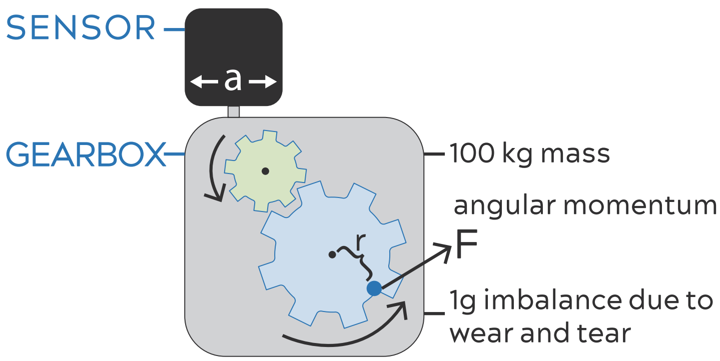

Figure 5: A gearbox failure is modeled as an imbalanced mass, creating a rotational force and acceleration on the gearbox assembly.

As a first approximation, the force created by a rotating off-center mass (a "rotating imbalance"), such as a gear that has experience wear-and-tear, can be calculated as the product of the rotating mass, its eccentricity, and the square of its rotational speed:

$$F = m * r * w^2$$

A typical gear may have a diameter of 50mm (r = 0.025m); assuming a mass imbalance of 1 gram (a high conservative number which would indicate an advanced stage of metallic ablation and a rotational rate of 500 RPM, the force created through the angular momentum is

$$F = 0.001\ \text{kg} * 0.025\ \text{m} * 500^2 \frac{1}{s^2} = 6.25\ \text{N}$$

6.25 N acting upon a 100 kg mechanical assembly (a low conservative estimate for the mass of the gearbox) results in an acceleration to be measured by the sensor of

$$a = F / m$$
$$a = 6.25\ \text{N} / 100\ \text{kg} = 0.0625 \frac{m}{s^2}\text{, or about } 0.0064\ g$$

Since no mechanical assembly is entirely free of vibrations, picking up this signal requires a sensor with suitable differential sensitivity. Few sensors on the market are capable of this type of sensitivity; those that are must be carefully located on the assembly to be effective and require specialized signal processing to remove noise. All of this adds cost and complexity.

Turning to machine learning allows the substitution of expensive signals with large amounts of data. ConnectedAI uses machine learning and pattern matching based on a multitude of sensors and observations, to assess the overall health of a process or machine. Instead of using a few expensive, high quality signals, ConnectedAI uses a much larger set of data to indirectly detect tiny signals and events, and tolerate larger error margins inherent in those lower cost signals.

ei3 has created a data sensing and processing device called "Symphony" to monitor the health of mechanical assemblies, such as rotational equipment, robots, or linear motion drives. Symphony is a wave field sensor, which uses multiple MEMS-based sensors to record a machine's wave field ranging from low-level accelerations starting from about 0.2 Hz, through "traditional" vibration from 2-20 Hz, sound from 20-20'000 Hz, to waves beyond sound up to about 40 kHz. The assessment of the wave-field information is accomplished within the Symphony device itself and can use the entire frequency spectrum.

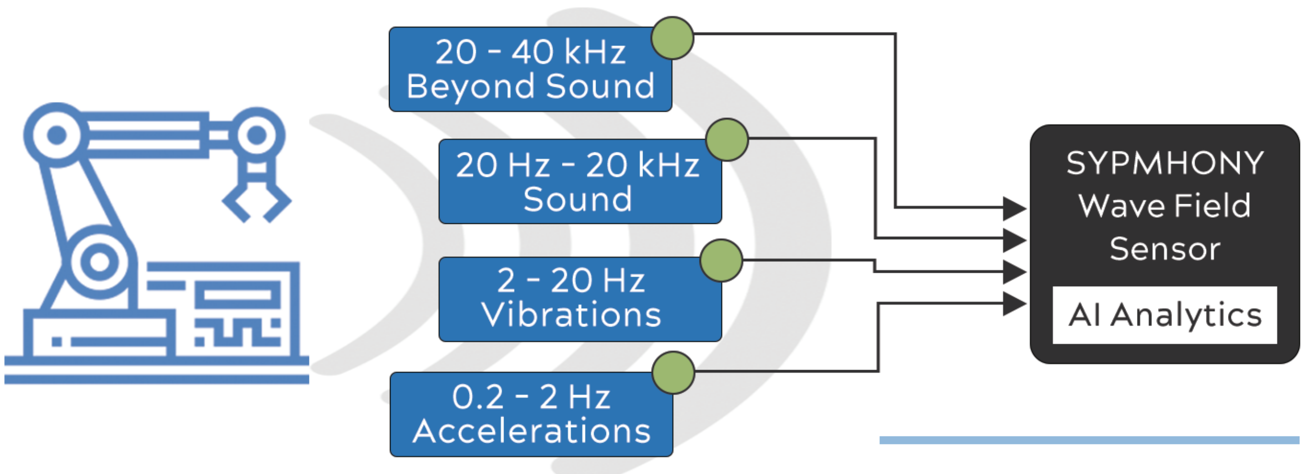

Figure 6: The Symphony device uses several sensors to record the wave field emitted from a machine. Careful analysis over time reveals subtle behavior changes that occur as a result of wear and tear.

Using the full spectrum enables the Symphony to detect changes to a dynamic system (such as a running machine). In the case of the gearbox, ei3's wave field analysis leverages the amplifying forces of resonance: While the vibration created by the failing gear is undetectable, its effects on the overall system are. By exhaustively and carefully recording, and analyzing, a wide spectrum of frequencies that capture the full behavior of the dynamic system, the signal becomes detectable through the resonance effects of its surrounding mechanical assembly.

Resonance effects in a mechanical assembly can lead to failures, sometimes catastrophic ones: The most well known example is probably the spectacular collapse of the Tacoma Narrows bridge in 1940, which occurred as a result of aeroelastic flutter caused by wind currents. The wind was only moderate and by itself caused no damage; mechanical resonance caused the collapse. Likewise, the wear and tear on rotating mechanical assemblies is often less detrimental than the consequences of the resonance that it creates.

It is important to note that this approach will not detect events that do not create resonances, such as mechanically well-dampened or isolated imbalances. By definition these imbalances have no impact on the overall machine and therefore their invisibility to the analysis may in fact be desirable.

Because of the unpredictable and chaotic nature of the resonance patterns created, the Symphony device can not use a-priori knowledge to detect an event. Even though the key frequency of the failing part can be calculated, it is unpredictable where the resonance effects will appear in the overall spectrum. Therefore, the Symphony device relies on machine learning to search and identify frequency patterns that indicate system changes.

The approach used by ConnectedAI is to use machine learning to create a Digital Twin by continuous observation, and use this twin as a baseline to detect changes to machine's behavior.

3.2 I Creating Digital Twins by Observation: The ei³ Digital Twin

ConnectedAI creates a Digital Twin of the machine behavior based on observations, instead of pre-determined mathematical models. This is achieved through the collection of sensor readings and process variables, which can be continuously collected easily and at a low cost, and their local analysis.

This data collection - the variables concerned and the read-frequency - is adapted to the nature of the machine and the process. Process variables are directly retrieved from the machine's control logic. Additional data may be collected from sensors that are placed on the machine, where necessary. The Symphony device is an example of such a sensor. Although the data collection is continuous, each collected data set represents discrete samples of the machine state at the particular time, creating a time-series.

Based on this time-series, the creation of the ei3 Digital Twin model is a two phase process that converts the collected data into "fingerprints" and "events".

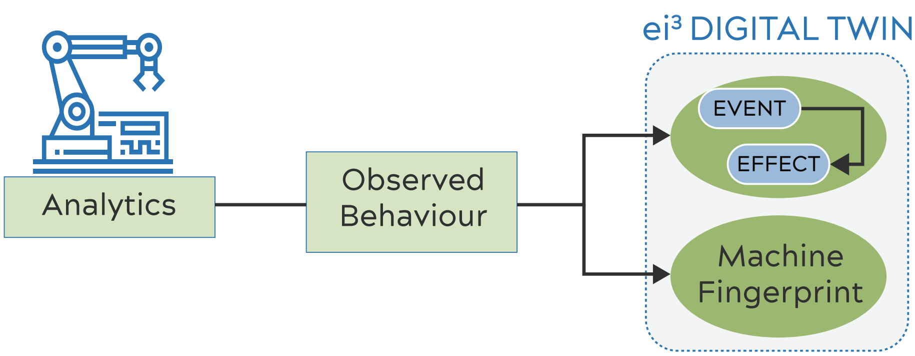

Figure 7: ei3 creates Digital Twins by observing the machine in operation, removing the need to create mathematical a-priori models.

In a first step, the time-series is processed through standard transforms to enrich the data. These transforms include the computation of classical statistical information for SPC (statistical process controls such as mean, average, standard deviation), a Fourier transform to assess the data's inherent frequency components, and a first derivative to track the rate of change. The resulting data is retained as the machine "fingerprint".

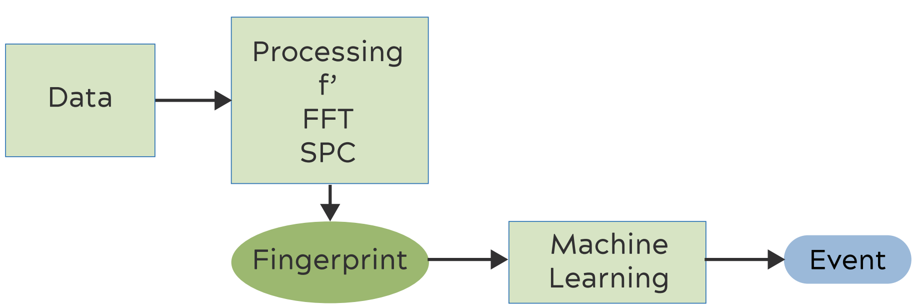

Figure 8: Machine fingerprints and machine events are created by analysis of recorded machine data.

In a second step the fingerprint is fed into a pattern detection algorithm that looks for changes to the fingerprints over time. These may include significant changes to SPC parameters which would flag variable outliers, emergent new frequency components within the FFT spectrum, or accelerating change to a variable sequence. Such occurrences are reported as events.

An example of events detected by the system are shown in the Figure 9. Multiple data sets of a single variable over time are superimposed to show the repetitive nature of a mechanical processing step. Drift is visible, as is typical given the non-perfect nature of real world machines; however an outlier is clearly in the center of the graph, highlighting a potential process failure.

Figure 9: A process failure is detected by observing a time series variable of a repeating mechanical process.

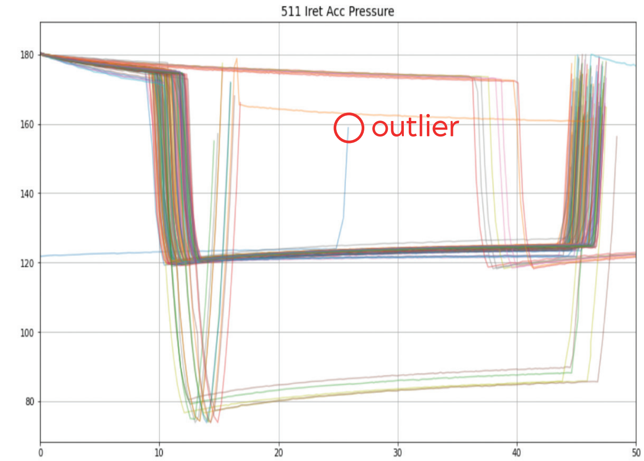

In Figure 10, two frequency spectra show significant differences as several new frequency components emerge between the data sets. The graphs show the amplitude of a frequency as a color, with dark blue showing low amplitudes, green and yellow progressively higher amplitudes, against the frequency on the Y-Axis. Time is running along the X-Axis.

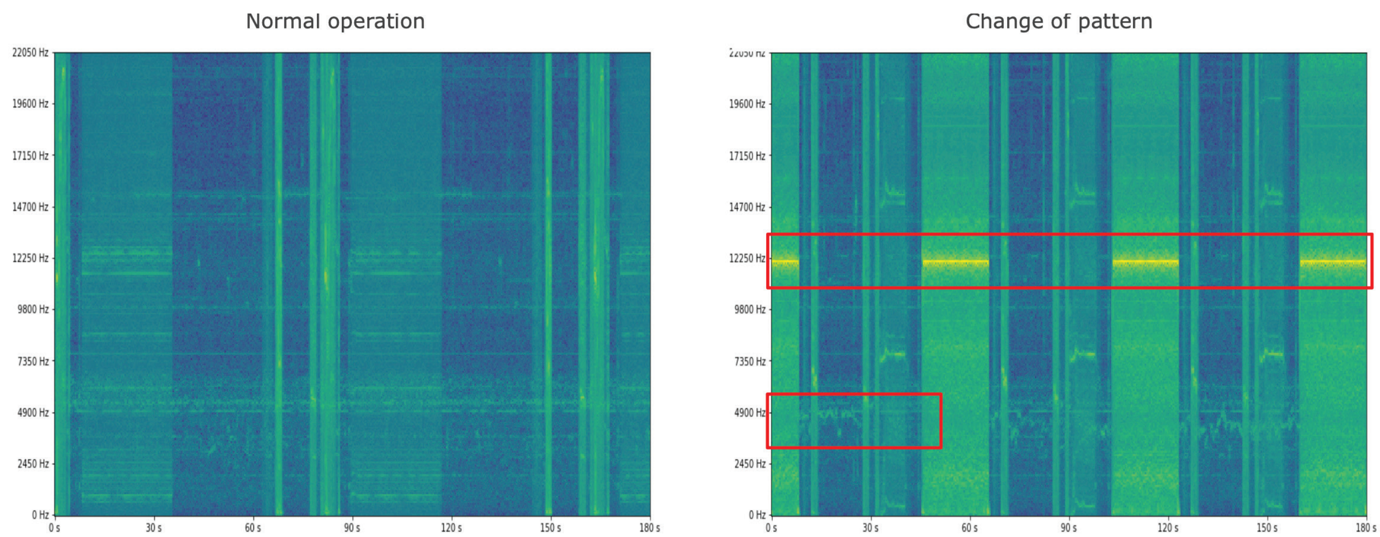

Changes to the model can be assessed to indicate an impending failure, gradual wear, or the result of a change to the machine's operating regime. This allows the model to be fluid and adapt to operational changes, ensuring optimal performance.

Figure 10: A process failure is detected by observing changes to the frequency spectrum of a time series

3.3 I Connected Assessment vs. Analytics in Isolation

Given the relative rareness of failure events and the slow progression of mechanical wear, the sharing of the insights of the individual machines across the fleet of machines is essential for the Digital Twin to emerge quickly enough to be useful.

ConnectedAI therefore shares data about events: observing events as described above is not done in the isolation of each individual machine, but in coordination across the owner's or manufacturer's fleet of machines. This ensures that events seen at one machine can be taken into account in the data analysis of other machines, and helps accelerate the detection of failure patterns.

This "Federated Learning" approach[1] allows many machines to contribute to a shared and common Digital Twin of a machine by sharing events without necessitating the sharing of the underlying data.

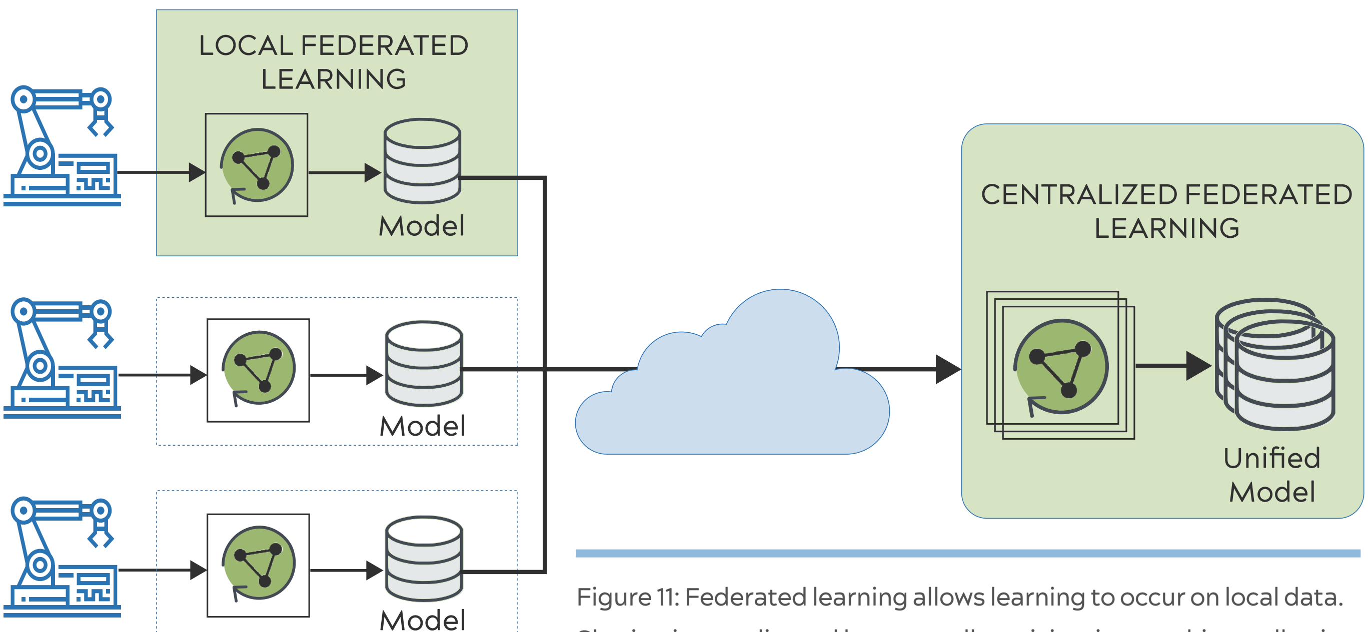

Figure 11: Federated learning allows learning to occur on local data. Sharing is coordinated between all participating machines, allowing each machine to benefit from the events detected by its siblings.

2 An excellent overview of federated learning approaches can be found in Federated learning: Beutel, Daniel J., Taner Topal, Akhil Mathur, Xinchi Qiu, Javier Fernandez-Marques, Yan Gao, Lorenzo Sani et al. "Flower: A friendly federated learning research framework." arXiv preprint arXiv:2007.14390 (2020).

This approach is in contrast to both a centralized and a completely isolated approach of learning. The centralized approach requires all events to be collected at a central location for in-depth analysis. This is the most common approach used today, however runs afoul of the privacy needs and data sharing policies of many machine owners and users.

Using the "centralized federated learning approach, ei³ provides a global twin that encapsulated collective knowledge and ensures data privacy.

On the other hand, the isolated approach is unsuitable because each individual machine typically does not experience enough events to provide efficient learning in useful time, and has the fatal flaw that a failure event must first be seen by a machine before it can be anticipated in the future.

ei3's ConnectedAI uses "centralized federated learning", where the ei3 data center acts as a central instance that coordinates learning sessions across multiple machines. Each individual machine inspects local events and event data to train a local Digital Twin. These local twins are then collected by the central coordinating instance in parametric form to construct, and share, a unified global twin, incorporating the insights and learnings from all machines. This global twin thus encapsulates the common, collective knowledge without sharing the events themselves, thus ensuring data privacy.

The common collective twin represents the learning and insights that are common to all machines across the fleet. Local deviations continue to reside in the Digital Twins at each machine. These local deviations may be due to different operating regimes and modes, or environmental conditions.

Using this approach, ei3 can provide fast and efficient predictions for even rare failure modes, and create insights that are both meaningful to individual machine owners, fleet owners, and machine builders: Events that are unique to individual machines most likely indicate part failures on those machines; events that are unique to the fleet of machines of an operator point to issues relating to this operator's SOPs, choice of materials, machine settings, or supply chain. Finally, events that are common across the entire fleet of machines will be of primary interest to the machine builder as they may indicate machine design issues.

4.0 I ConnectedAI in Action: Detecting Mechanical Issues in Large Scale Laminators

One of the first deployments of a predictive solution based on the principles of ConnectedAI has been put into operation with an operator of continuous web laminators. This particular operator creates plastic films for medical applications. A laminator is used to deposit chemicals on top of plastic substrate. The precise layer thickness of each deposition is critical and must be maintained exactly for the final product to be usable.

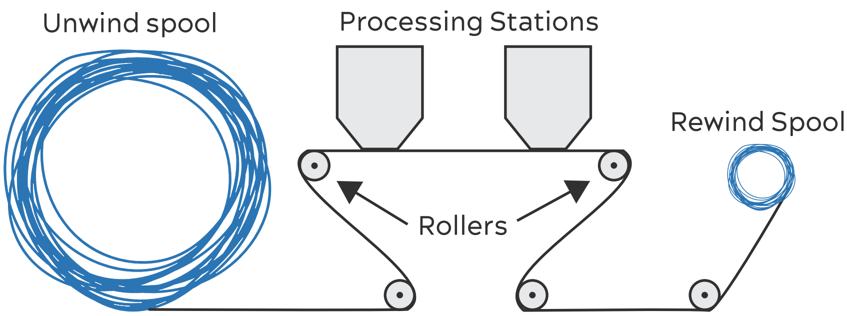

The lamination machine is a large and complex system, where the plastic moves along a convoluted pathway, guided by rollers, from an unwind spool, through multiple processing stations, to a rewind spool. In this application, even slight vibrations create ripples in the material and the deposited layers, resulting in defects and an unusable product. Failing bearings in the rollers are a key source of vibration, and are the primary target of the data analysis.

A set of ei3 Symphony devices, mounted at different places on the main frame of the laminator are used to capture the wave fields emanating from the laminator. The captured spectrum of data includes the vibrations and other important signals related to the health of roller bearings, spools, and motor operations, such as the higher frequency signals created by metal ablation effects that indicate early bearing wear – often audible as a high pitched "whine" to experienced service engineers.

The massive amount of data created by the Symphony sensors is impractical to aggregate it in a central database. The "Federated Learning" approach of ConnectedAI allows the Symphony sensors to collaborate and create a common behavioral model of the machine while it is in operation, utilizing shared insights on the datasets in lieu of a central data collection. (As the laminator belongs to a single operator, in this case there would not be any data privacy concerns related to sharing this data.)

During operation, the speed of the laminator, measured in RPMs (revolutions per minute) of its rollers, changes from time to time. In addition, the diameter of the unwind and rewind spools vary as the material travels through the machine from the unwind spool to the rewind spool. This background could present a formidable challenge to any conventional application of signal processing, looking for the vibration effect of a failing bearing.

As an example, consider the amplitude level for two different machine states in Figure 13. The blue and the orange line each are distinct recordings of amplitude levels on the Y-axis of the same, randomly selected frequency, against the RPM range of the machine on the X-axis. As to be expected, the amplitudes rise exponentially with machine speed, however resonance effects across the complex mechanical assembly add significant deviations, creating different number of peaks to occur at different speeds for different operating modes: slight changes in the machine's process state - such as different diameters of the unwind and rewind spools as the material is being processed by the machine, lead to different resonance behaviors.

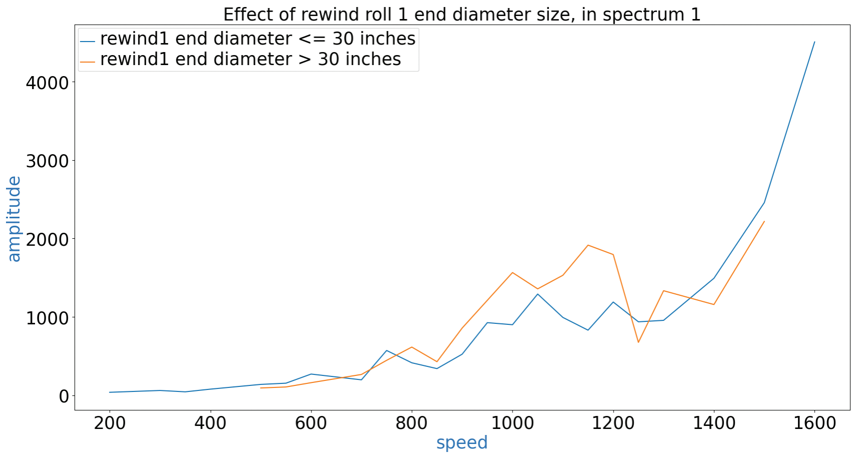

Figure 13: Amplitude response of a critical frequency over the speed-range of the machine, for two different operating modes.

ConnectedAI creates Digital Twin models based on the fingerprint of the wave field data, as the machine progresses through its operating states. These states are characterized by key machine operating parameters, such as machine RPM, spool diameters, and state information of deposition heads. Although the space described by these machine parameters in principle is infinite, data quantization effects lead the algorithm to identify approximately 100 distinct states, and for each a characteristic wave field data set is recorded.

Each new wave field data set is compared against the existing library of previously recorded fingerprints. A distance function between the new data set and the historic data determines if the new data set can be deemed "normal" or if an abnormality is indicated, in which case an alert is raised.

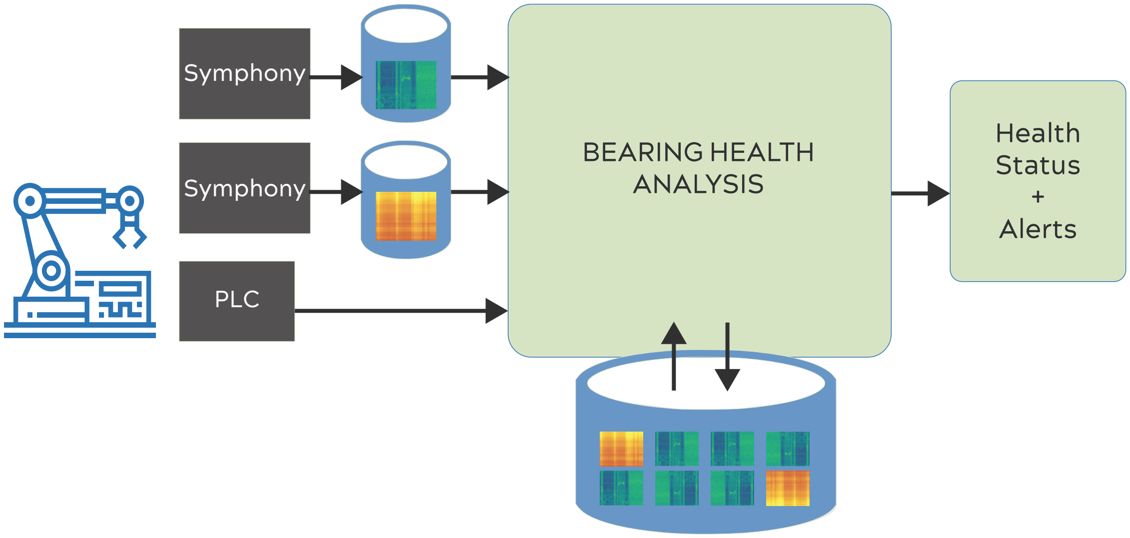

The effect of this process becomes apparent in the graph below, where all recorded wave fields are shown, each as a single dot representing a complete frequency spectrum. To simplify the visualization, the Y-position is based on the weighted average amplitude of the spectrum, causing similar spectra to appear at a similar position along the Y-axis.

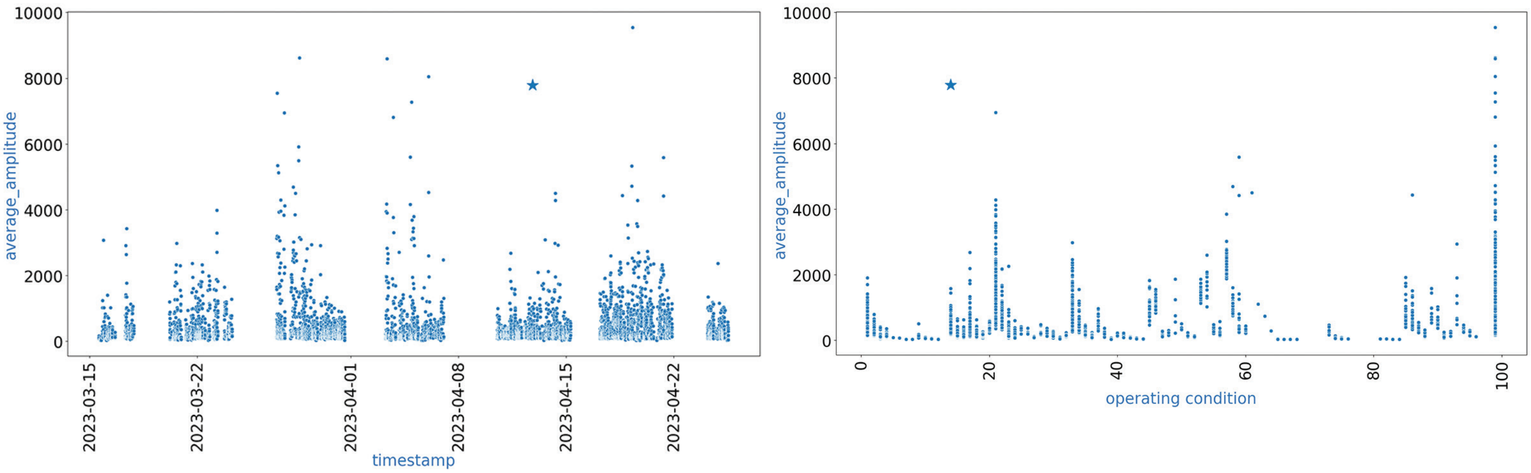

On the left hand side of the graph all recorded wave fields are shown with the X-axis representing the time of the recording. Periods of machine inactivity are visible as gaps in the graphs. The data was collected over a one month period, during which one failure event occurred, indicated by the blue star. On the left side of the graph, the failure event is not reliably distinguishable from normal operations.

On the right side of the graph, the same wave field data is sorted into the approximately 100 operating states identified within the Digital Twin model created by ConnectedAI. The failure event is again identified as the blue star, this time clearly identifiable as an outlier within its respective operating state or “column” of relevant historic fingerprints. The outliers are identified based on statistical analysis within each fingerprint, based on average and standard deviation. The event identified by the blue star is the only recorded event outside the 2x standard deviation band around the mean within its respective fingerprint group.

5.0 I A Look at What's Next

After many years of trying, recent advances have allowed AI-based chat-bots to surprise the popular mainstream with human-like prose, if not always factually correct content. Applying similar technology to industrial applications with their stringent correctness requirements, limitations imposed by data privacy, and the relative scarcity of data continues to present a formidable challenge.

ei3's ConnectedAI is addressing this challenge by applying a centrally orchestrated, federated learning approach to model the behavior of industrial machines without a-priori knowledge. This enables learning based on limited and disjoint data sets, without the need to share data, and can accommodate large amounts of operational variation between machines in different operating environments.

With its unwavering commitment to innovation and excellence, ei3 is poised to remain at the forefront of developing practical ConnectedAI solutions that delivers positive outcomes for the owners and builders of industrial machines.

ConnectedAI enables learning based on limited and disjoint data sets, without the need to share data. It can accommodate large amounts of operational variation between machines in different operating environments.

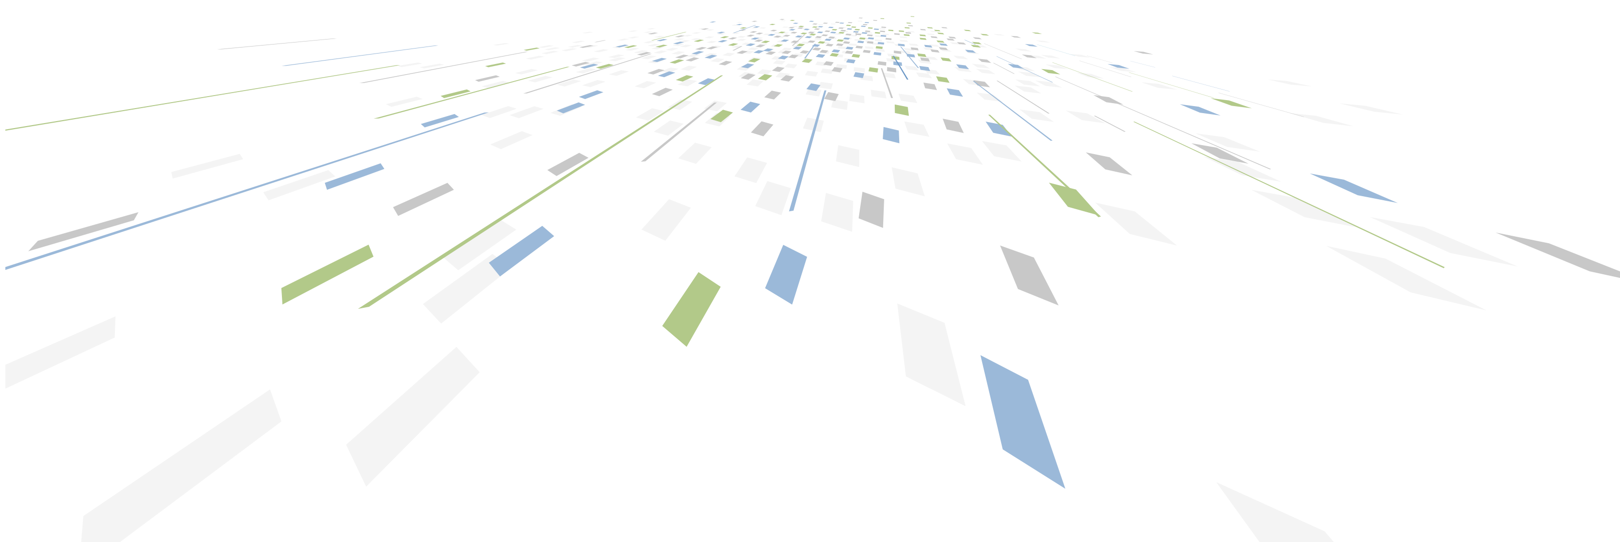

Visit www.ei3.com

Discover our growing list of client success stories, our products, and explore ways to stay ahead of the competition.

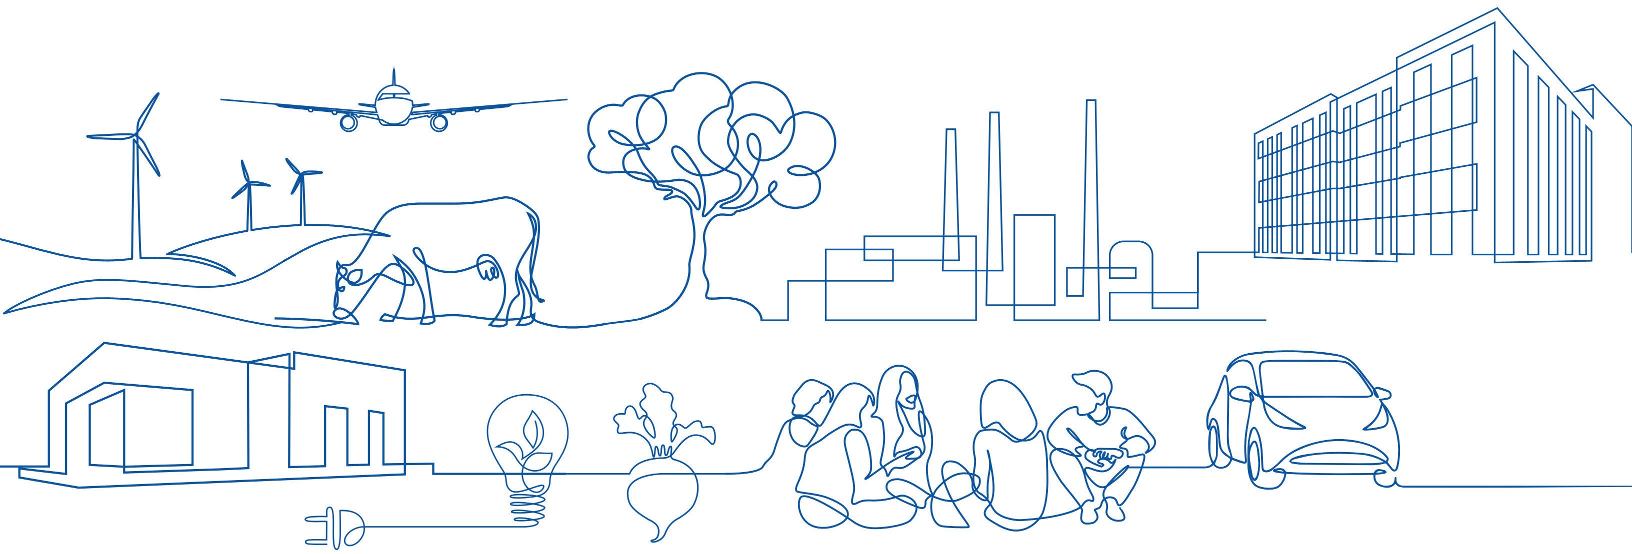

ei³ Corporation

Tel: +1 201 802 9080 I E-mail: contact@ei3.com