ODEN TECHNOLOGIES

Whitepaper

Introducing Mímir

Driving the Future of Industrial Automation: Machine Learning and Artificial Intelligence for Manufacturing

Make things. Better.

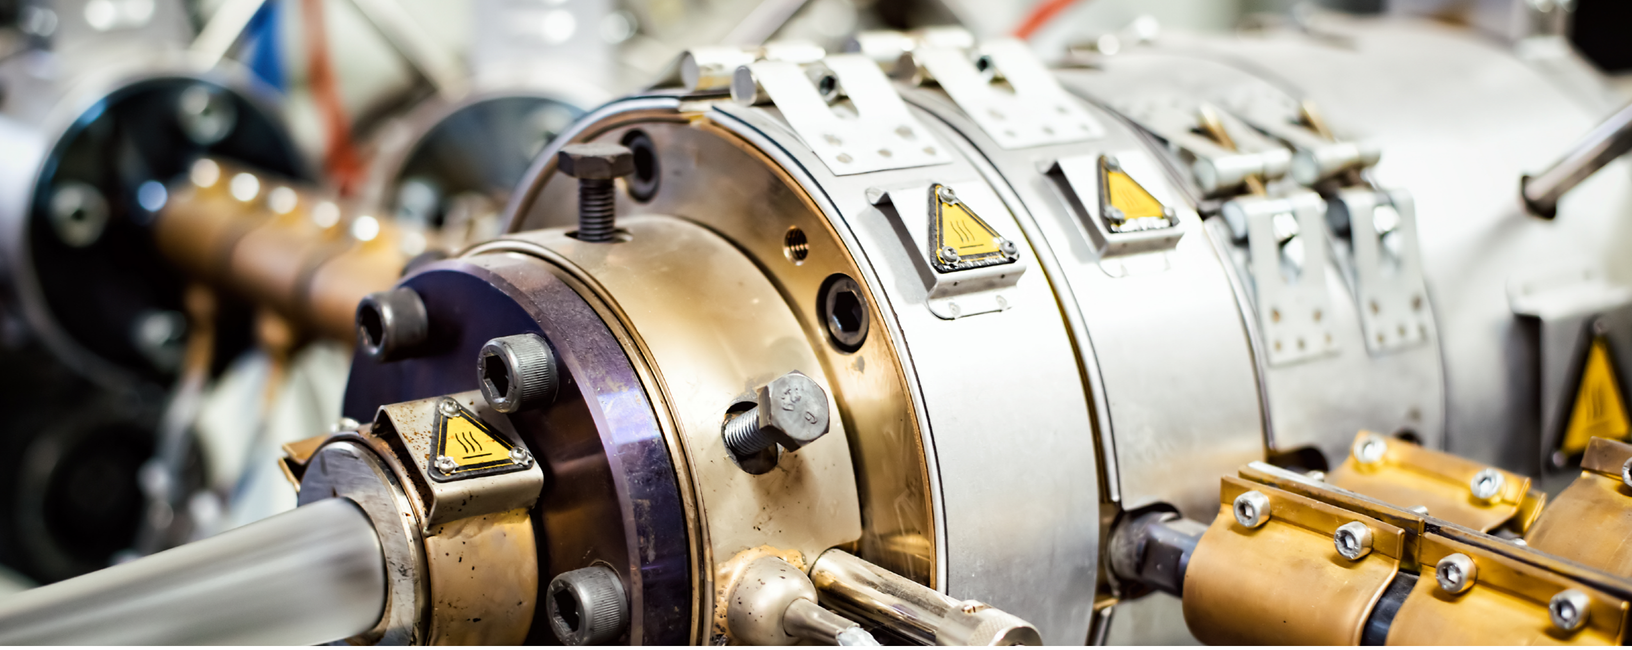

Table of Contents

Executive overview

There is a huge amount of excitement around the promise of Data Science (DS) and Machine Learning (ML) for Industry 4.0 and the Industrial Internet of Things (IIoT). It is widely accepted that DS and ML-based applications will be used to assist operators, process engineers, managers, and executives to optimize the manufacturing process end-to-end.

To be effective, an Industry 4.0 platform must do three things:

1. Deliver high-quality data
2. Support DS and ML application innovation to answer different types of questions, including:
   - → **Diagnostic:** Providing insights into the current and prior state of operations, and tackling questions such as "How am I doing relative to how I have done previously," "Show me something like this," and "What led to the current state?"
   - → **Predictive:** Making predictions about the future state of operations given the current state and prior history, and answering questions in the form "How is my quality," (Predictive Quality) "When do I need maintenance," (Predictive Maintenance), "Alert me before failure," and "Let me know if something is abnormal?" (Outlier Detection)
   - → **Prescriptive:** Providing recommendations on the best course of action, accounting for the potential constraints imposed by the business workflow (processes and resources), and handling the scale and diversity of the environment
3. Operationalize ML applications and integrate them into the business process workflow.

Users will expect DS and ML applications running on an Industry 4.0 platform to be always-on and self-learning, improving their capabilities over time from observation of the processes, user actions, and their corresponding outcomes. Users will also expect them to be interactive, providing multiple modes of interaction, learning from these interactions, as well as supporting the discovery of previously unknown or hidden behaviors that affect the process.

This requires fundamental advances to DS and ML algorithms, tailoring them to the specific characteristics of the manufacturing processes and their associated data. It also requires deploying these algorithms in an operational mode within production-grade solutions that can operate reliably, continuously, and scale up and down as needed within the manufacturing environment. However, operationalizing Machine Learning-based solutions is one of the greatest challenges that industrial customers face today[1,2,3,4].

Oden's approach to Machine Learning and AI

Mímir is named after a figure in Norse mythology renowned for his knowledge and wisdom. The Mímir framework within the Oden platform is the first-of-its-kind framework built from the ground up to provide operational Machine Learning-based solutions for Industry 4.0 customers.

It does this by addressing three core requirements:

1. **Tight integration with data:** both streaming and stored data from all sources in the manufacturing environment are continuously gathered, curated, cleaned, and made available.

2. **Advanced IIoT specific ML and DS tools and algorithms:**
   - → Open, modular, and extensible tooling that exploits the best of open source as well as customer specific tools for ML and DS.
   - → Best in breed DS and ML algorithms adapted to the data and requirements of IIoT.
   - → ML model lifecycle management that supports experimentation, continuous performance measurement, and validation of data, algorithms, and models.

3. **Production-grade deployment infrastructure:**
   - → End-to-end ML capabilities both in the cloud or on edge devices located on the manufacturing floor.
   - → Integration into the business workflow by leveraging the Oden IIoT platform.

The framework is designed to support ML innovation by significantly streamlining the ML solution development process – including iterative data gathering, algorithm experimentation (including what-if analysis on real data in an isolated development environment), and validation.

What this means is that the ML solution development process that currently takes weeks to months in the most advanced organizations, can be achieved over a matter of days with appropriate repeatability, and verifiability. The framework is designed to be open, allowing our own data science teams – as well as those of our customers – to experiment rapidly with novel algorithms.

The framework leverages the Oden IIoT platform to provide a robust production grade deployment infrastructure spanning the cloud and the edge. It utilizes the cloud for reliability and scalability, and the edge for low-latency and disconnected operations. It can also dynamically shift processing and resources between them – a capability unique to Oden.

This deployment infrastructure also provides scaling on demand and continuous performance monitoring of the deployed ML solution. This empowers DS teams to determine if models need to be replaced, or modified, and closes the train-validate-deploy loop for self-learning and adaptation.

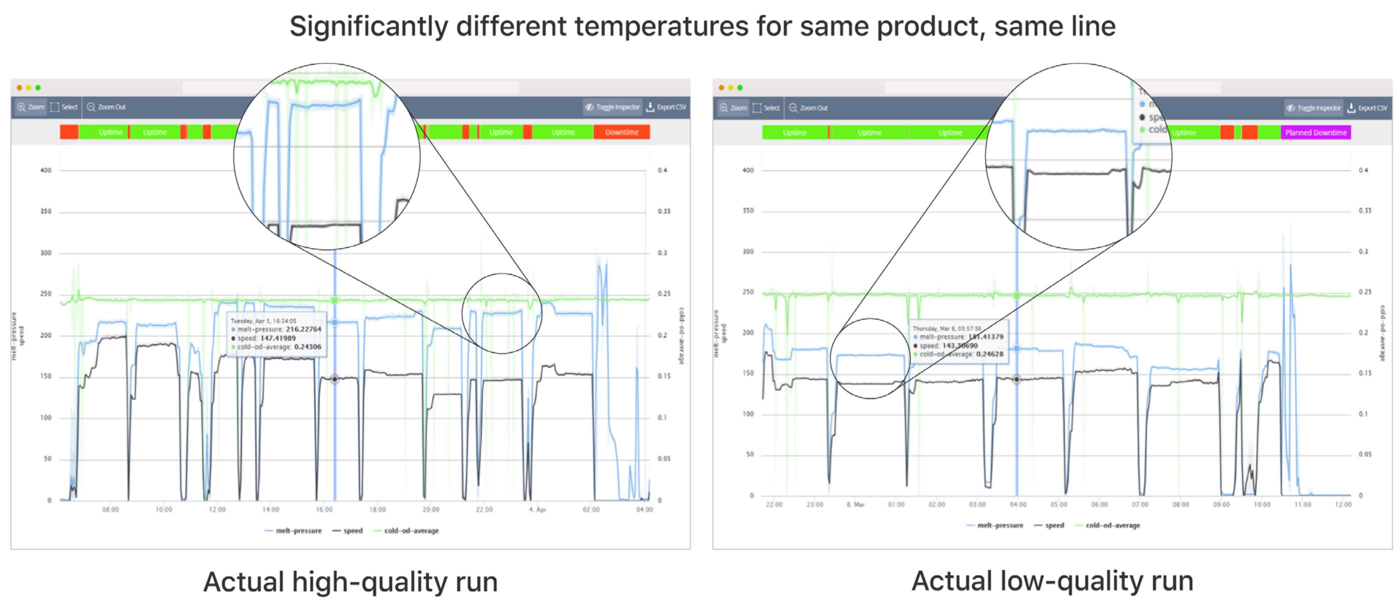

Figure 1. Example of discriminative attribute (Temperature) identified within the Oden console

The Mímir ML and AI framework advantage

The ML tools provided by the framework deployed on the Oden IIoT platform enable the optimization of the manufacturing process in multiple ways.

The following two examples highlight the diagnostic and predictive capabilities of the platform and illustrate their use in real manufacturing settings.

Diagnostics with discriminative patterns
The framework can be used to highlight attributes, and their patterns and trends, that vary significantly across different process outcomes (e.g. low quality and high quality) using a set of Discriminative Pattern Mining tools. These DS techniques can be used to comb through vast amounts of time series data efficiently to extract such patterns that can then be visualized to allow engineers to perform root cause analysis.

For instance, in Figure 1 it is clear that there was a significantly different temperature for the same product in two different runs resulting in the difference between a High-Quality Run and a Low-Quality Run.

Such patterns are exposed to the engineers who then use them to tighten and improve the process further.

These features are already being used by plastics extrusion manufacturers to increase steady-state throughput, reduce ramp-up times, reduce changeover and setup times, and to decrease compound waste.

Prediction of metrics and production quality
The Mímir framework can also be used to generate live predictions about metrics of interest so that their potential future impact on the process can be identified with sufficient lead time. This enables process engineers to modify and control the process in order to improve quality, reduce maintenance costs, and avoid any unplanned downtime.

Consider an example from injection molding shown in Figure 2. The framework ingests these metrics in real time and applies pre-trained predictive models to generate potential future values that can be then be quantized (into alerts of different sorts) and visualized within an interactive display. Engineers then use these alerts to control the process in real time. The system can then continue to learn based on the intervention as well as the resulting outcomes.

There are multiple other examples of the types of platform capabilities and their usage in manufacturing that we will publish as additional case studies. Using the Oden IIoT platform, our current customers have

Figure 2. Example of predictive alerts within the Oden console

already achieved reduced defects, increased engineer productivity, decreased defect resolution time, reduced waste and improved throughputs, resulting in savings and gains of $100,000s per line per year.

Using the Oden IIoT platform, our current customers have already achieved:

- → Reduced defects
- → Increased engineer productivity
- → Decreased defect resolution time
- → Reduced waste
- → Improved throughputs

How does the Mímir framework work?

The remainder of this whitepaper will describe the features of the framework in more detail, and discuss how the core Oden IIoT Platform is leveraged to deploy our ML and AI solutions.

The Mímir Machine Learning and AI framework is designed to provide the following capabilities:

The first end-to-end, open ML framework

The Oden platform provides ML algorithms access to both historical as well as streaming operations, resources, and performance data that is continuously collected in real time. It uses a first-of-its-kind architecture spanning the cloud platform as well as edge computing.

Historical data is stored in our large-scale cloud storage subsystem and made available via APIs that provide integration with several ML tools while streaming data is continuously available via subscription APIs. Our framework uses these to provide consistent views of the data across the cloud and edge and also enables ML algorithms to operate on the data for offline training or for online scoring (and potentially learning).

Results are also easily surfaced to process teams and executives through the Oden platform and to the factory floor via the Oden Operator Interface. This includes several tools for visualization of and interaction with data and results so that they can be accurately interpreted.

Open, modular and extensible ML tooling
The framework is designed to leverage the best open tools, algorithms, and libraries for ML capabilities. These tools include:

- → Data analysis and standard ML libraries[5] (scikit-learn, NumPy, SciPy statistics) from Python, Gradient Boosted Tree Libraries (XGBoost[6] and LightGBM[7]),
- → Deep Learning tools (Keras[8] and Tensorflow[9]).

This allows Oden to incorporate the latest advances in open literature within our ML pipelines. It is also uniquely designed to allow our customers and their data science teams to build models using these open tools, and deploy within the platform, so that targeted algorithms can be used as necessary.

Best in breed ML algorithms for manufacturing time series data
The framework is also designed to be extensible with custom and novel algorithms for time series data analysis in manufacturing. Oden recommends taking a two-tiered approach, in which our DS and ML teams develop algorithms that are generally applicable across multiple customer problems and validated against multiple customer datasets, while our customer's teams can either specialize in these algorithms or develop their own proprietary algorithms, specific to their manufacturing environments.

Customers can deploy combinations of these general and specialized algorithms to monitor, control, and optimize their processes.

Within the Oden ML team, we are building several novel and generalized algorithms for:

- → Univariate and multivariate time series analysis
- → Pattern mining (discriminative and frequent) and causal modeling
- → Clustering and graph mining
- → Predictive modeling and forecasting
- → Outlier detection

We are also actively investigating the use of deep learning models, such as Long Short-Term Memory (LSTM) neural networks and their applications to the manufacturing environment, for time series data. These custom algorithms and trained models will be made available via the framework.

Models can be seamlessly deployed into production in the cloud, or at the edge using the Oden IIoT platform infrastructure.

Model lifecycle management
The framework includes support to gather process metrics continuously to allow measurement of the effectiveness of the ML algorithms. This includes capturing labels/ground truth derived from the ingestion of online and offline quality subsystems and reports, feedback from users, as well as metrics computed on live process data, translated into measures of model performance, e.g. Receiver Operating Characteristics (RoC) or Precision-Recall (P-R) curves, or Process Capability Indices (e.g. Cpk) — in both experimentation as well as production mode. This allows the data science teams of both Oden and our customers, as well as their operations teams, to gauge the impact of any changes or updates, and allows for rapid iterative development. It also allows for A/B testing of models and inference of performance improvements in a statistically reliable manner.

Once validated, models can be seamlessly deployed into production in the cloud, or at the edge using the Oden IIoT platform infrastructure, as required.

Also included within the framework is support for ML model lifecycle management. This includes the storage of the model and associated metadata, versioning of the

model, metadata and links to data used to train the model, and ability to reinstate and replay data through models for auditing, experimentation, and compliance requirements.

End-to-end ML capabilities
The framework is designed to provide support to train ML models at scale, from small datasets to terabytes of data, using cloud scaling, and leveraging specialized hardware such as graphics processing units (GPUs) where needed. Using a novel Oden-created specification for metadata in addition to custom wrappers for ML pipelines, the platform then allows data scientists to rapidly close the loop by deploying these trained models against live streaming data in both an experimentation mode as well as in production.

Another unique capability of the platform includes distributed deployment of ML algorithms, within the cloud or at the edge (using Oden edge devices on the factory floor) to support the desired latency, robustness, and privacy concerns.

Production grade ML integrated into the operational path
The framework leverages the Oden platform — the de facto operating system for factory processes to translate ML solutions into production deployments — and includes native support for scaling, both in the cloud as well at the edge, reliability, and security and privacy protection. In addition, the core Oden platform enables tight integration with business processes — bringing in the appropriate data and outcomes, distributes results to the appropriate consumers, and provides an interactive environment to the end-users, all while supporting the latency, throughput, and performance requirements of the manufacturing floor.

Benefits and applications in the manufacturing industry

Oden's customers see improvements with the Oden IIoT platform ranging from 20% increases in monthly output to 50% decreases in total scrap, resulting in millions of dollars in savings and additional revenue each year.

Specifically for our customers, this improved time to value of ML-based solutions for diagnostic, predictive, and prescriptive applications will result in:

Improved quality and decreased compound waste/giveaway
With Machine Learning, we continuously learn the top contributors to product quality. This can be used in both a diagnostic manner, to identify what caused the quality to degrade, as well as — more interestingly — in a predictive manner by continuously forecasting quality even before it is measured.

This allows us to proactively alert engineers about previously undetected issues, thereby avoiding costly re-runs or scraps. Additionally, by combining such ML-based prediction with appropriate AI reasoning and optimization techniques, the system can recommend changes to the operating process parameters to optimize product quality — moving towards dynamic process control.

Product traceability
In cases where the data capture is not complete, we can use learning based techniques to infer some

of the lineages of a product, by comparing against other products to find those that are most 'similar'. Combinations of clustering, frequent pattern mining and similarity analysis are used to achieve this.

Increased output and reduced ramp-up and changeover times

Predictive maintenance is a great use case for ML. We use learning to identify the core contributors to machine wear and tear, and predict the likelihood of machine breakdown at a future horizon, given current and past operations.

Additionally, since these events are 'infrequent' compared to the granularity with which data is collected, we can also learn more robust general purpose models by aggregating across manufacturers (while ensuring data privacy), and specializing these general models to each individual customer. We can also identify abnormal states within the operations, such as unusually long ramp-up times, that can adversely impact machine health.

Increased steady-state throughput

ML can be used to learn the optimal equipment operating parameters for each product, and we can couple this with machine control to change the equipment operations dynamically in real time. The operations can be optimized across product quality, energy consumption, and equipment health. Furthermore, we can also use outlier detection techniques to identify when machine behavior is not as expected and use this to alert appropriate process control engineers.

Retaining process know-how

With ML, we can learn by observing the operations in the factory and correlating against product quality, equipment health, energy usage etc., to identify best practices in operation. This does not require explicit knowledge extraction from the workforce and can be performed with implicit observation.

We can also explicitly capture knowledge from the workforce using Active Learning-based approaches, where the system asks for clarifications only when it is unsure. This can be used to make sure that the system learns interactively from users, but does not add significantly to users' overheads.

In addition to learning from observation, the system also needs to be able to ingest structured and unstructured knowledge - from textbooks and user notes, operation logs and product manuals, for example — and be able to place current observations in the context of this knowledge, Here, by leveraging NLU techniques and AI Reasoning, the system can be used to make inferences about the process even without necessarily observing all possible conditions, as needed for a pure learning-only based approach.

Increased operational efficiency

The Oden platform provides continuous visibility into the process and factory operations. Our Machine Learning framework streamlines the ML solution development process — such as iterative data gathering, algorithm experimentation (including what-if analysis on real data in an isolated development environment), and validation.

This allows users to rapidly prototype and test new ideas, and — if successful — deploy them into their operations with little or no overhead. The platform also ensures that deployed ML-based solutions operate in real time, scaling them as needed, and keeping them resilient to failures in the network or infrastructure.

Artificial Intelligence (AI) Capabilities for IIoT

The Mímir framework is being designed to accommodate modalities and techniques in addition to ML that form the basis of realizing AI-based solutions.[10] Specifically, in the context of IIoT the platform is being designed to support capabilities for:

- → **Domain knowledge capture:** collection and ingestion of the vast amount of domain knowledge about the process, the equipment, and materials, acquired and curated by operators, engineers, and designers. This includes ingestion of structured and unstructured sources of domain knowledge.
- → **Techniques for reasoning**,[11,12] about domain knowledge, so that we can supplement ML and provide it with some guardrails, especially when there is not enough data or ground truth available.

- → **Techniques for planning,**[13,14] so that the platform will be able to optimally sequence and orchestrate a set of potentially complex actions to accomplish the desired goal of process improvement.

- → **Natural modes of interaction:** The framework is also being designed to support extensions for interaction with people through appropriate speech, video, and haptic interfaces.

The ability to integrate these different capabilities within one platform, build solutions that utilize combinations of them to solve the most complex problems, and deploy them seamlessly within the manufacturing environment will eventually bring Oden closer to having a true AI assistant for Industry 4.0.

Summary

The Mímir Machine Learning and AI framework is being designed as both modular and extensible with the best of breed open source, customer-specific, and Oden-developed algorithms. This provides tight coupling with data and operational processes through the core Oden platform such that we can formalize the ML application development methodology, significantly reduce the time to value of ML, and realize the best such operational solutions for IIoT.

For further information about our Machine Learning and AI framework Mímir and the Oden platform, please contact: mimir@oden.io

ODEN TECHNOLOGIES

oden.io
+1 (800) 230-9046

Sources

1. D. Sculley, T. Phillips, et al, “Machine Learning: The High Interest Credit Card of Technical Debt,” NIPS Workshop on Software Engineering for Machine Learning, 2014.
2. M. Zinkevich, “Rules of Machine Learning: Best practices for Machine Learning Engineering,” ICML Workshop on Reliable Machine Learning in the Wild Workshop, 2017.
3. https://www.quora.com/How-do-you-take-a-machine-learning-model-to-production; https://www.itproportal.com/features/overcoming-the-challenges-of-machine-learning-model-deployment/
4. https://conferences.oreilly.com/strata/strata-ny/public/schedule/detail/69255
5. http://scikit-learn.org/stable
6. https://xgboost.readthedocs.io/en/latest
7. https://lightgbm.readthedocs.io/en/latest
8. https://keras.io
9. https://www.tensorflow.org
10. With the significant gains achieved recently in multiple domains through the use of neural nets and Deep Learning techniques, there is a tendency to conflate AI with Deep Learning. AI requires other capabilities that leverage knowledge, reasoning, planning and interaction.
11. https://en.wikipedia.org/wiki/Knowledge_representation_and_reasoning
12. L Wos, R Overbeck, E Lusk, J Boyle, “Automated Reasoning: Introductions and Applications,” 1984.
13. https://en.wikipedia.org/wiki/Automated_planning_and_scheduling
14. S. Russell, P. Norvig, “Artificial Intelligence: A Modern Approach,” Prentice Hall, 2009.