// 2025

SPARK BEYOND

AI for **‘Always Optimized’** Industrial Operations

SPARK BEYOND

About us

Established in 2013 to accelerate AI-powered problem-solving.

Since then we have **delivered $Bns in tangible ROI** for our customers **across 100s of use cases.**

Mission

Unlock AI-driven **'Always Optimized' KPIs** for any organization

Global Footprint

Presence across Asia, Europe and US with employees spread across **8 countries**

Industry Validated

100s of success stories across within **Fortune 500 companies globally**

Partner first DnA

Partner-first organisation with global reach with GSIs

Enterprise Ecosystem

Azure aws Google Cloud

ChatGPT Claude Gemini

SPARK BEYOND

Industrial IoT Use Cases

Select top-line and bottom-line impact generating use cases

MINING

Fuel/Resource Efficiency
Detect shifts in fuel use and adjusts controls to minimize excess

Safety Incident Root Cause Analysis
Predict scenarios of safety incidents from historical data

DISCRETE

OEE Root Cause Analysis
Track OEE issues continuously to address inefficiencies early

Operator Copilot
Guide operators to act effectively with automated root causes and SOPs

CONTINUOUS

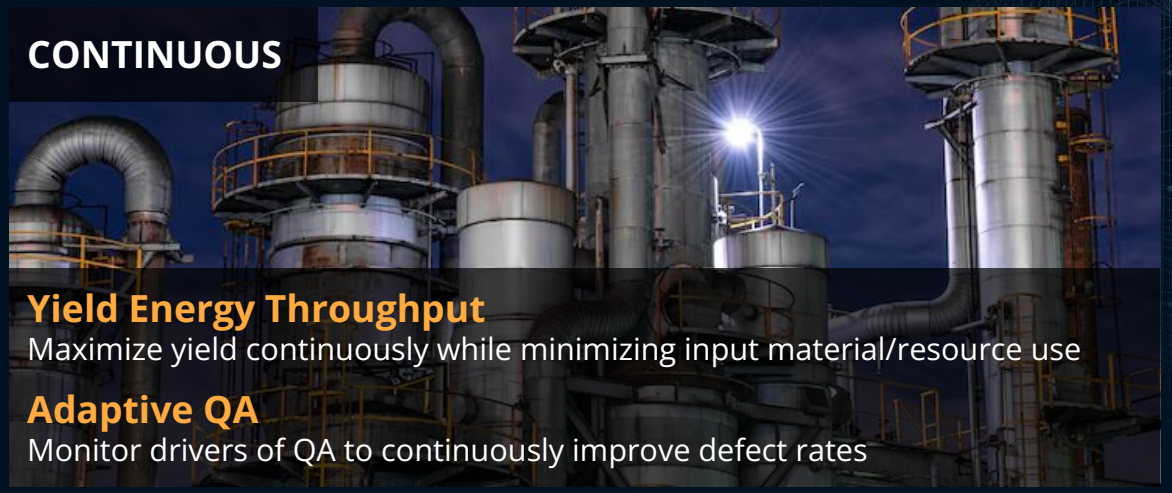

Yield Energy Throughput
Maximize yield continuously while minimizing input material/resource use

Adaptive QA
Monitor drivers of QA to continuously improve defect rates

ENERGY

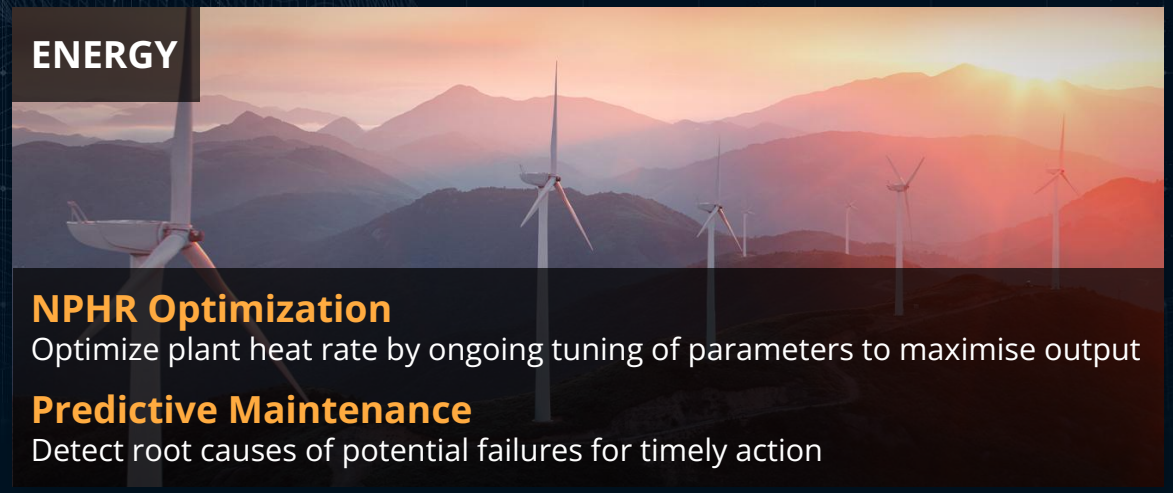

NPHR Optimization
Optimize plant heat rate by ongoing tuning of parameters to maximise output

Predictive Maintenance
Detect root causes of potential failures for timely action

Our Technology

SPARK BEYOND

Generative AI doesn't understand YOUR business.

For KPI optimization, AI must leverage knowledge from operational data

Challenges of LLMs

- Limited in understanding patterns hidden in complex operational data
- Unable to ground business reasoning in data.

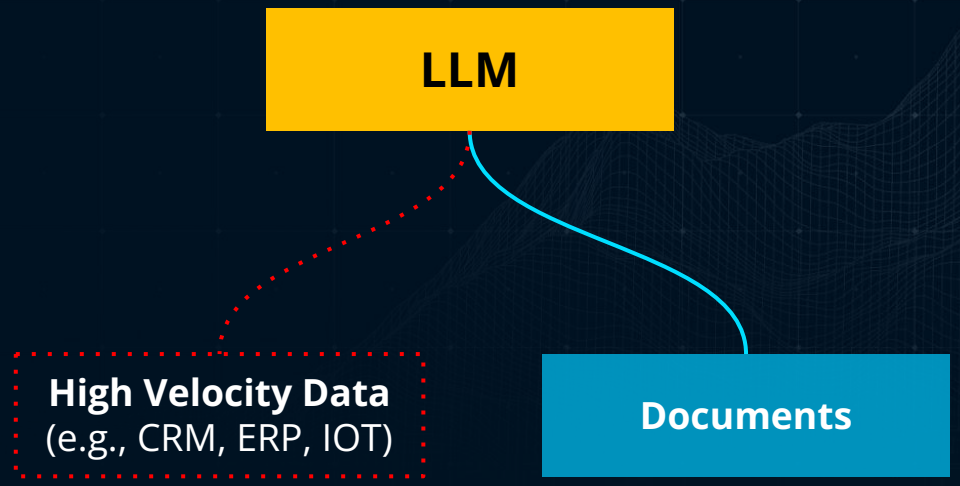

Unlocking LLM-powered KPI-optimization for solution-builders

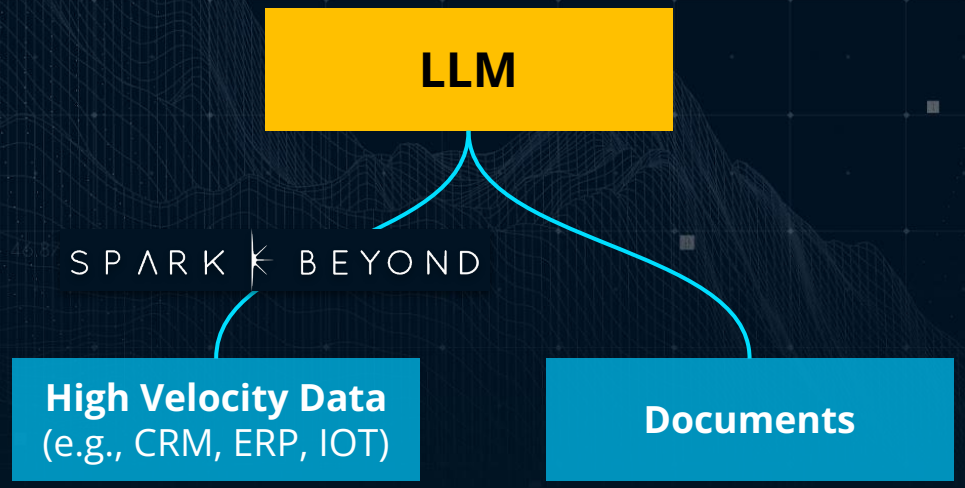

SPARK BEYOND

‘Always Optimized’ KPI Architecture

Continuous feedback loop creating impact from enterprise structured data

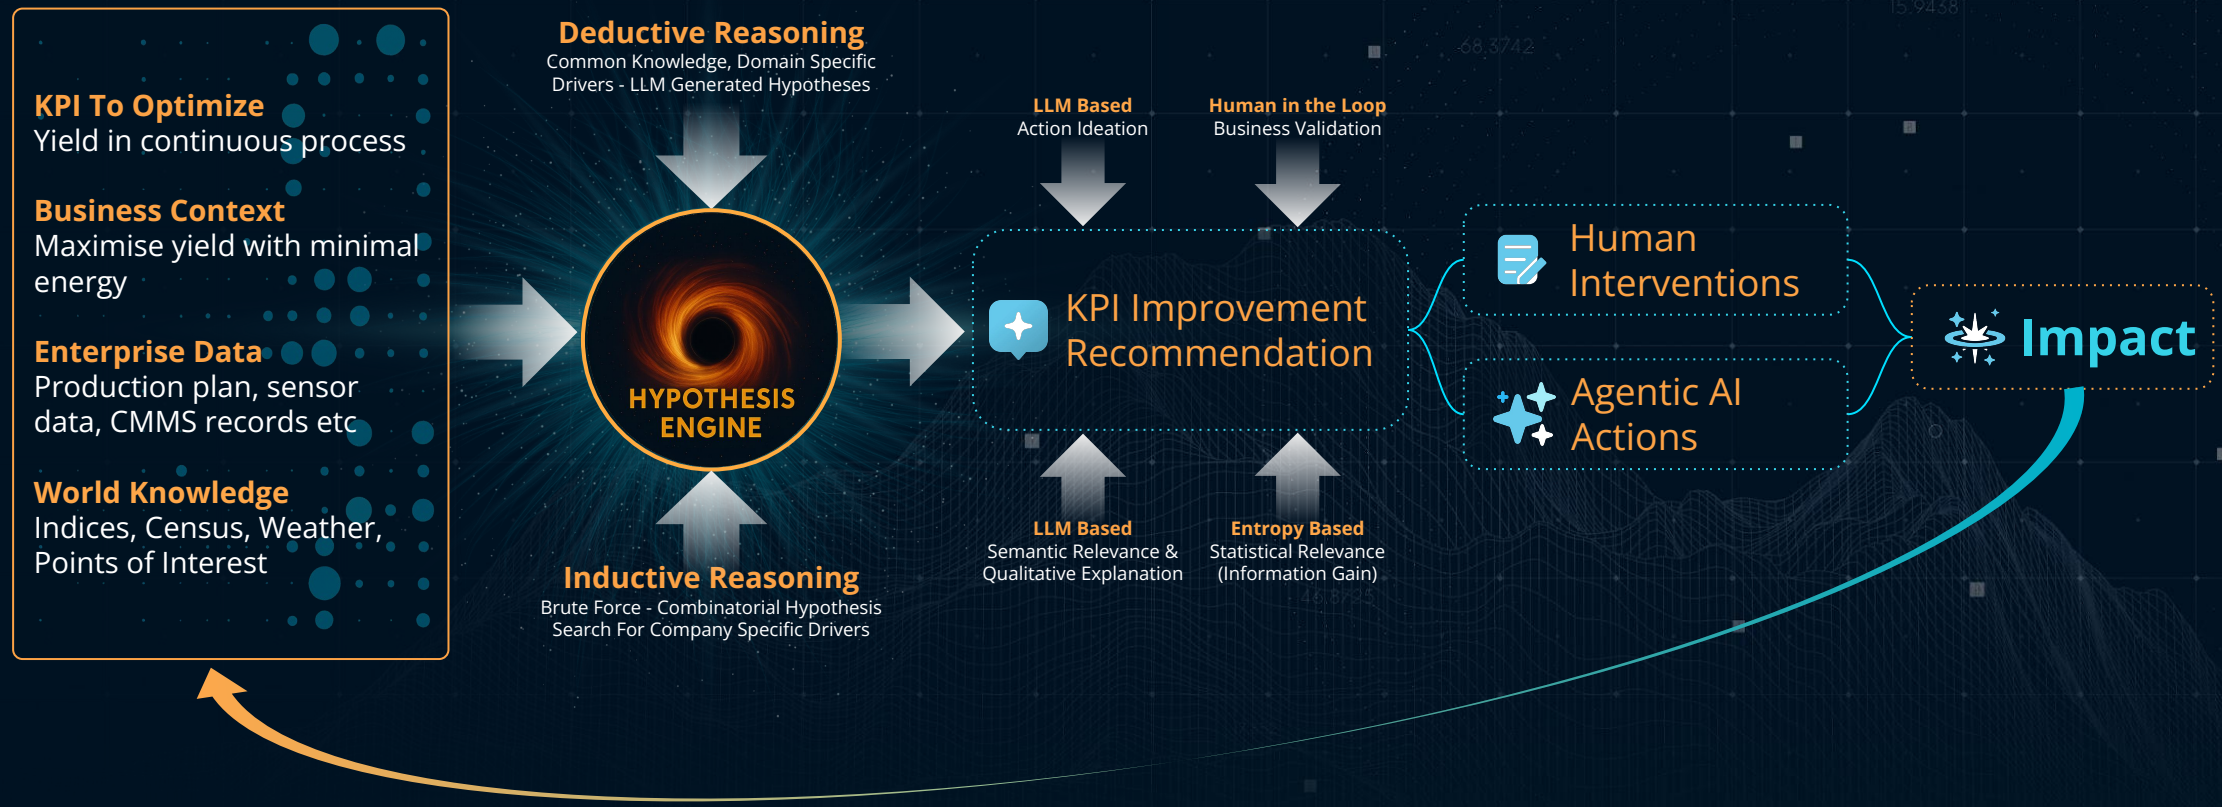

SPARK BEYOND

Making the paradigm shift to ‘Always Optimized’ KPI Optimization

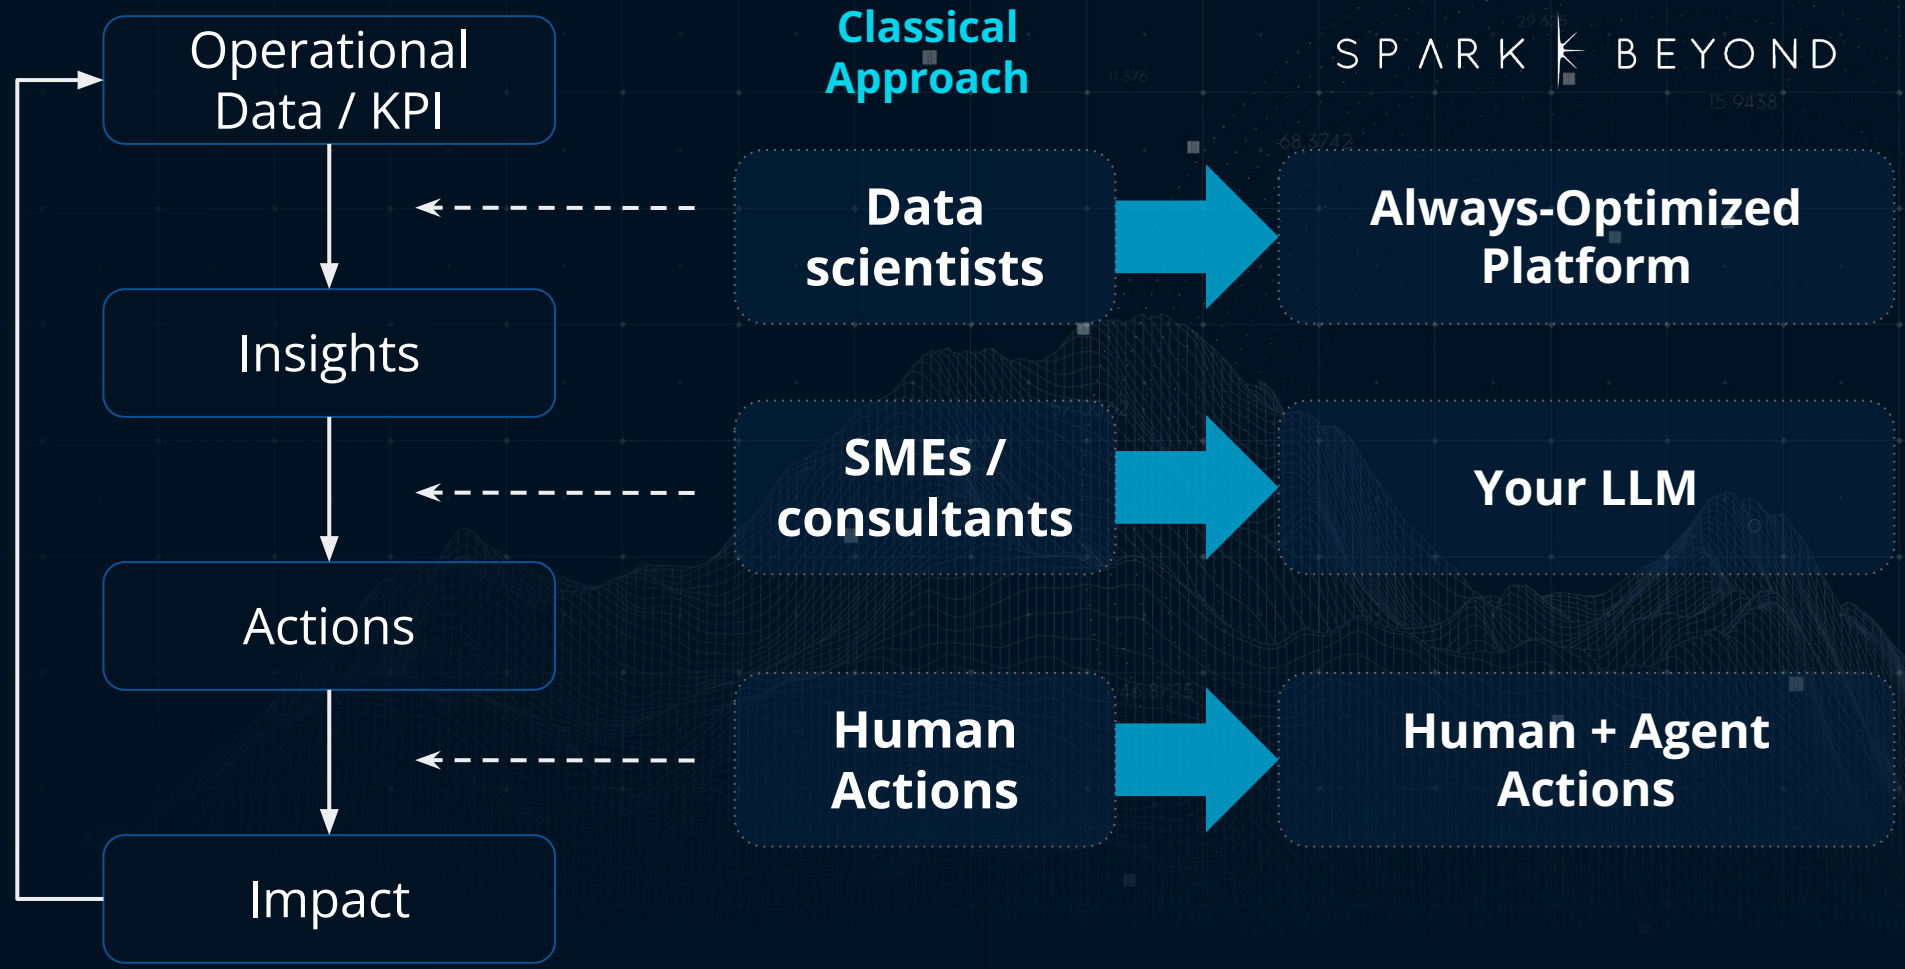

SPARK BEYOND

AI CoE Platform

Accelerate Results - No Large Data Teams. No Consulting Roadshows.

Discover & Prototype in Weeks

Bypass long discovery cycles. We'll help you rapidly prototype AI solutions that uncover the hidden drivers impacting your KPIs.

Prove ROI Before You Scale

Validate the P&L impact with a targeted pilot, delivering a concrete business case and tangible value from day one.

Scale to an "Always-Optimized" System

Empower a lean team to deploy and expand use cases, creating a continuous optimization loop for your business.

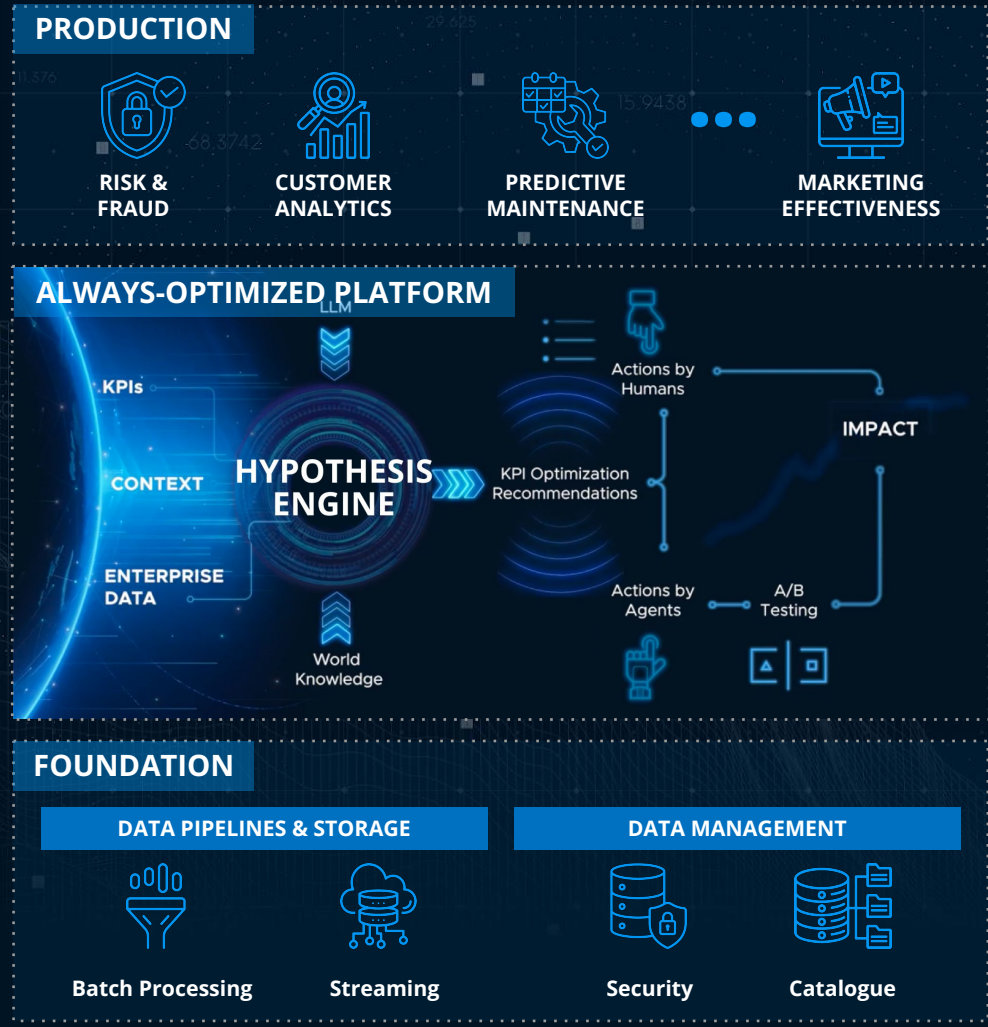

SPARK BEYOND

Why SparkBeyond for Your AI CoE Platform

Feature	Impact
Hypothesis Engine	Automatically discover and ***engineer new features*** from your operational and business data, surfacing ***hidden drivers*** of performance. Generate and test millions of hypotheses to identify true root causes and actionable levers for ***KPI improvement.***
Explainable AI	All insights and recommendations are delivered in clear, natural language—enabling business and operations leaders to ***understand, trust, and act*** on AI-driven findings. This transparency is critical for adoption and compliance.
Open Integration	SparkBeyond is cloud-agnostic and integrates seamlessly with leading hyper-scaler environments, CRMs, ERPs, and LLMs, ensuring ***flexibility, scalability, and alignment*** with your evolving tech stack.
Battle Tested @ Scale	The platform is **proven across 100+ Fortune 500 deployments**, with a track record of rapid time-to-value and measurable ROI.
Unified ML & Gen-AI	SparkBeyond is a ***unique combination*** of advanced machine learning and Gen-AI agent workflows providing a single foundation for a broad spectrum of AI use cases.
Rapid Deployment	Designed for ***fast implementation*** and operation in complex, distributed data environments—delivering actionable insights without requiring lengthy data warehouse projects.

SPARK BEYOND

Existing approaches to link LLMs to enterprise data are insufficient to address structured data needs

Overview of current approaches (not-exhaustive)

Pre-Training & Fine Tuning

What is it?

Pre-training a model on a selected corpus applicable to your enterprise domain
Fine-tuning LLMs to answer domain specific questions

Limitations

- Expensive to re-train
- Does not address structured data sources
- Fine-tuning is better suited to teaching specialized tasks or styles and less reliable for factual recall.

Retrieval Augmented Generation

What is it?

Retrieve data from outside a foundation model and augment your prompts by adding the relevant retrieved data in context

Limitations

- Structured data requires a query for RAG based solution to retrieve
- Retrieved query needs to be LLM compatible
- RAG is largely limited to searchable documents

Code Interpretation & Generation

What is it?

LLM task to translate a query spoken in natural language into SQL/code automatically

Limitations

- User needs to define the intent and insights
- Path to using the insight in an LLM use case is several steps away for a user

In-Context Learning

What is it?

One/few-shot learning example to gain new knowledge (e.g. feeding an existing ppt report about a quantitative analysis)

Limitations

- Context needs to be textual
- Context document can get easily outdated

Credentials

SPARK BEYOND

Boosting beverage production line efficiency with root cause analysis

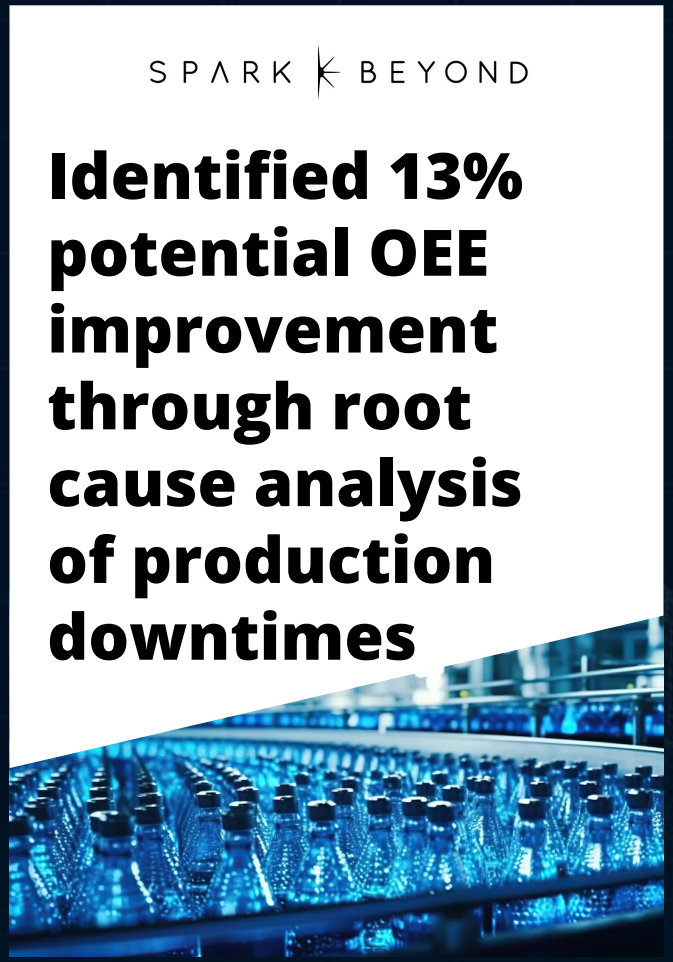

CHALLENGE

- A leading global beverage manufacturer wanted to improve **Overall Equipment Effectiveness (OEE)** for its filling and packaging lines
- Despite having **Lean, Six Sigma, and excellence programs** in place, they sought to use data to better understand **root causes of downtime** and to **foresee breakdowns** proactively

RESULTS

- Sparkbeyond autonomously pinpointed key root causes: volume output spikes, older product batches and failures at upstream machines
- Estimated **13% potential OEE improvement** from addressing these causes

APPROACH

- Collected **performance data** on OEE shifts, outages, process orders, and machine status
- Identified **"built-back" events** as the biggest negative driver of OEE
- Applied predictive analytics to **identify root causes** and recommended targeted actions

SPARK BEYOND

Achieving fuel efficiency in mining with up to 10% savings

SPARK BEYOND

Reduced fuel consumption by 10% in 4 months across mining fleet

CHALLENGE

- High fuel costs (30% of OPEX) pressured a mining fleet to cut consumption
- Needed to identify controllable drivers across a fleet of 100+ vehicles to reduce operational expenses
- Manual reviews are costly and time consuming prevent fuel used to be optimized

RESULTS

- Identified **~300 predictors** of fuel efficiency, including payload, tire pressure, and driver behavior
- Rolled out **dashboard monitoring** for key fuel drivers
- Achieved a **10% fuel consumption reduction** in 4 months, with ongoing optimization

APPROACH

SparkBeyond created a digital twin for each vehicle using:

- IoT Vehicle Management System (30+ onboard sensor data points)
- Fleet Management System (trip data)
- SAP maintenance logs
- External sources (weather, contextual data)

Each digital twin integrated 130+ data points per truck

SPARK BEYOND

Optimizing the Net Plant Heat Rate by 3.3% based on load conditions

SPARK BEYOND

Cutting fuel costs and emissions by optimizing plant heat rate with advanced analytics.

CHALLENGE

- Low efficiency due to boiler and combustion losses
- Heat rate varies with operational load
- High volume of sensor data from over 3000 sensors from PI system but low actionable insights
- Need to **reduce coal input costs and emissions**

RESULTS

- Identified 23 high-impact controllable drivers and predicted NPHR impact per driver
- Estimated **3.3% heat rate improvement** corresponding to **~3% reduction in CO_2 emissions**

APPROACH

- Model NPHR by load condition levels (Low, Medium, High)
- Ran learning experiments to rank thousands of features
- Prioritized initiatives and actions via SME workshops
- Built RCA and ideal parameter setting dashboard for engineers

SPARK BEYOND

$3m impact from optimizing alumina/caustic ratio in alumina refining

CHALLENGE

- Controlling the A/C ratio in the digester is critical to the alumina refining process and a major production driver
- However, the client was relying on an outdated predictive model, resulting in high variability throughout the process

RESULTS

- SparkBeyond identified **several ways to improve consistency** in the A/C ratio, with one improvement alone potentially driving over **$3 million** in added production value
- The client implemented changes valued at **$1.8 million annually**

APPROACH

The team used SparkBeyond to improve model accuracy through:

- Analyzed internal sensor data from the existing DBO model
- An **iterative process** to identify optimal time windows and interactions between sensor readings
- Built several models, from **linear regression to advanced boosted trees**, to maximize impact

SPARK BEYOND

Reducing non-conformance rates in shaft annealing process for a bicycle manufacturer

Direct cost improvement from reducing scrap

SPARK BEYOND

Identified adjustments for annealing process to reduce non conformance rates by 25%

CHALLENGE

- Automatically identify the root causes of **non-conformance** amongst a sample of **annealed metal parts,** augmenting engineering intuition of operators

APPROACH

- Combined **visual inspection data** with **furnace sensor** in historian**, power consumption, nitrogen, oxygen** levels and **environmental factors** (room temp, humidity)
- Generated hypothesis and **identified root causes**

RESULTS

- Identified **nitrogen pressure and oxygen content levels** during heating and cooling respectively, which when fixed, help reduce non-conformance rates from **29% to 4%**

SPARK BEYOND

Proactive identification and resolution suggestion for industrial chiller trips at a bicycle manufacturing plant

Avoid throughput losses

SPARK BEYOND

Prevent temporary loss of plant capacity by proactively preventing chiller trips

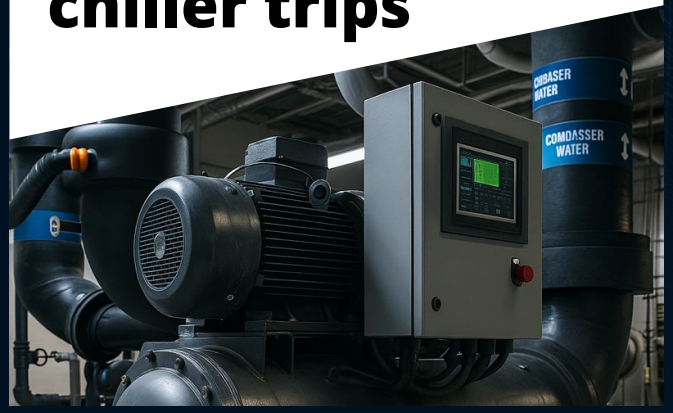

CHALLENGE

- Greenfield site with new **industrial chiller system for cooling** tripped frequently and abruptly with no consistent explanation - each trip while lasts a short duration leads to **plant downtime**

APPROACH

- Generate hypothesis from sensor and other data from **SCADA** and **operational systems**.
- Use LLMs to bring **engineering and physics knowledge** to boost insights.
- Use **LLMs + OEM manuals** to validate actions and recommendations.

RESULTS

- Developed a highly accurate predictive model with **85% recall and 80% precision** for identifying trips - $ value for preventing temporary plant shutdown not known to SparkBeyond

SPARK BEYOND

Predictive Maintenance of ESPs (Electrical Submersible Pumps)

SPARK BEYOND

$2M Impact per Early Warning for Key ESP Failures

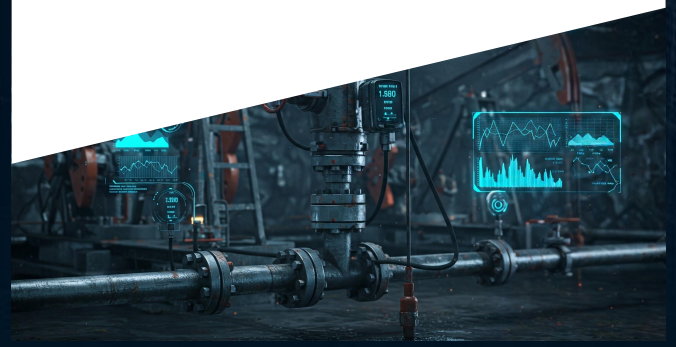

CHALLENGE

- Client needed to anticipate and manage ESP failures to minimize production loss
- Built a model to predict the probability of ESP failure within the next 100 days
- **Data Sets Used:** Past ESP failures, sensor readings, well trajectories, coordinates, completions

RESULTS

- **$2M impact per early failure alert**
- Enabled proactive maintenance, minimizing downtime and production loss

APPROACH

- Reframed task as remaining uptime prediction due to dataset imbalance
- Identified if an ESP is likely to fail within the next 100 days
- Used Discovery Platform with 8 datasets to uncover failure drivers
- Delivered insights as both code and natural language
- Provided daily SHAP-based predictions and explanations
- Outputs shared with maintenance teams to support preventive action

SPARK BEYOND

Proactive Predictive Maintenance for PCP Pump Failures

SPARK BEYOND

30+ Days Predictive Horizon | >50% Failure Detection Accuracy | 0–15% False Positives

CHALLENGE

- **Improve the life of wells** to increase production and reduce CapEx
- Build a predictive model despite incomplete data
- Integrate analytics with the client's **IT infrastructure**

RESULTS

- Identified **>50% of system failures** with a 0–15% false positive rate
- Defined 9 predictive parameters contributing to model power
- Achieved a 3**0+ day predictive horizon** with additional system data

APPROACH

- Focused on **Analytics & Insight and IT Integration**
- Built a live model to predict failures in PCP sub-systems from incomplete data
- Delivered an **integrated view of the IT infrastructure** and provided software improvement recommendations

SPARK BEYOND

Predictive Maintenance for Bearing Failure in CMP wafer process

SPARK BEYOND

Advanced analytics predicted failures weeks ahead, improving uptime.

CHALLENGE

- **Short lead time** limited predictive capacity
- **Manual process** unable to anticipate failures in advance
- **8-hour downtime** affects all 4 chambers

RESULTS

- **Earlier failure prediction** than manual process
- **Weeks of advance warning** enabled
- Improved **maintenance scheduling**
- Reduced **unplanned downtime**

APPROACH

- **Connected data** from vibration and maintenance sources
- Used **APC system** to track sensor signals up to 0.5Hz
- Leveraged **Maintenance records** with free-text.
- **Analyzed observations** before bearing failures

SPARK BEYOND

Optimizing wind farm production by continuous temperature adjustment

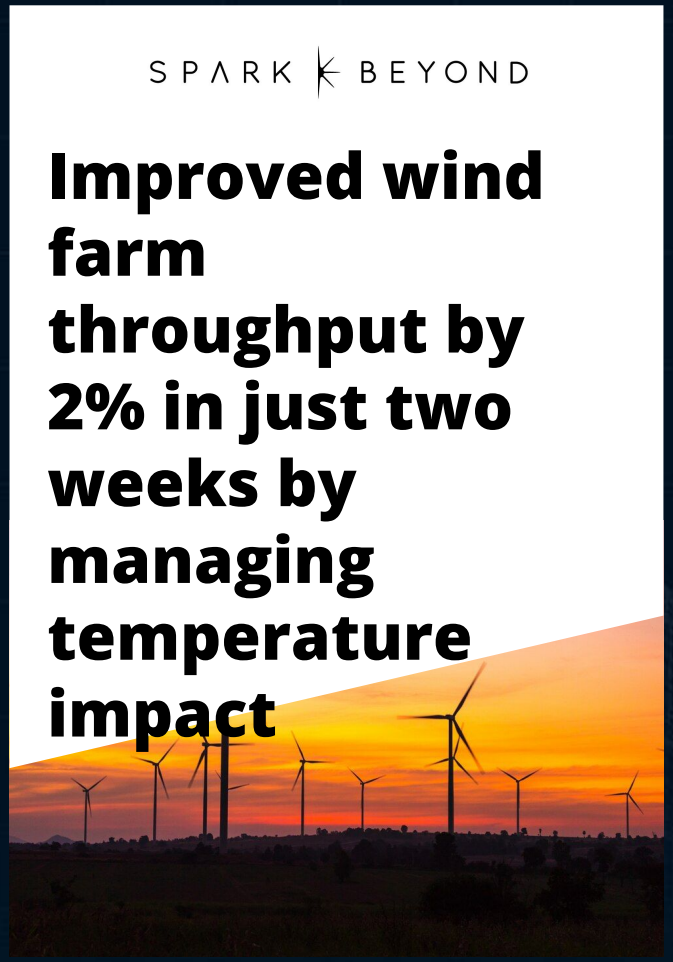

CHALLENGE

- A large European energy company faced **unexpected asset downtime** across multiple wind farms, impacting production and IRRs
- Traditional methods were **missing key drivers** behind the downtime

RESULTS

- SparkBeyond uncovered that **moving from cold to hot weather impacted turbine lubricant viscosity**, and reduced performance over time - so “temperature rate of change” was the most valuable metric
- Initiatives to control the rate of change led to a **2% improvement in throughput** in a two-weeks

APPROACH

- Automated **time series analysis of sensor data** across various periods (12h, 3 days, 1 week, 1 month)
- Identified **temperature sensors** as critical to performance
- **Augmented internal data with weather datasets** to uncover new drivers

SPARK BEYOND

Maximizing mean annual increment for a plantation

CHALLENGE

- A forestry company in Asia wanted to better identify **predictors of tree growth**
- They struggled with **inconsistent data** from two different MAI measurement methods and sought to **develop a consistent methodology** while **discovering new opportunities** to boost growth rates

RESULTS

- Identified human-controllable drivers of forest growth, validating expert knowledge
- Built an **elaborate model explaining ~40% of MAI variability**, even after excluding biological factors like tree diameter and height

APPROACH

- Reduced data preprocessing time and generated **100+ predictive models** across two datasets, including time series models
- Identified **critical features** (e.g., text list variables) improving MAI predictions
- Detected gaps in maintenance and planting methods
- Paved the way for **adaptive and autonomous AI systems** to ensure quality results

SPARK BEYOND

Advanced Analytics for Semiconductor Manufacturing

SPARK BEYOND

Boosted semiconductor line efficiency with >40% defect reduction and predictive insights.

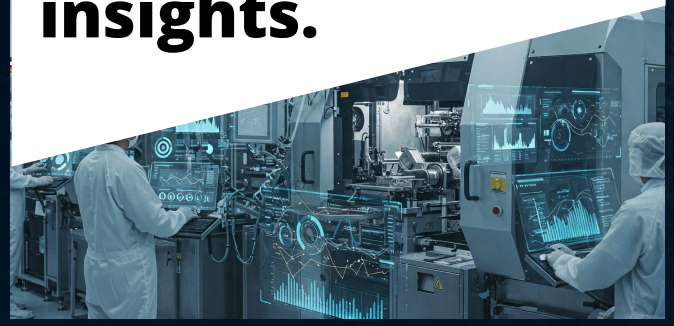

CHALLENGE

- Identify and reduce production inefficiencies.
- Improve productivity in high-precision semiconductor environment.
- Address recurring issues with long idle times and CPU overload.
- Predict equipment failures to prevent downtime.

RESULTS

- **40% reduction in yield detractors** in target workcenter.
- **5 root causes identified** with clear mitigation actions.
- **Predicted failures 4+ weeks** in advance with ~10% false positives.
- Boosted maintenance capacity and line stability.

APPROACH

- Analyzed 13 months of I/O data and **20+ parameters.**
- Used RCPS to identify **1,000+ statistical patterns** and top 15 root causes.
- Applied predictive maintenance using vibration sensor data.
- Built failure prediction models with 90% event coverage.

SPARK BEYOND

Reducing Particle Contamination in CVD Process

SPARK BEYOND

Eliminating root causes to reduce particle contamination in CVD.

CHALLENGE

- Particle defects drive downtime and wafer scrap.
- CVD issues cause **0.4% yield loss in production.**
- Root causes behind contamination remain unclear.

RESULTS

- **Uncovered key root causes** of particle defects.
- Defined actions to **reduce contamination levels.**
- Improved stability and production quality.

APPROACH

- Combined top-down hypothesis testing with bottom-up modeling.
- Automated root cause analysis using process data.
- Tailored diagnostics for high-priority failure points.

SPARK BEYOND

Predicting Failures in Semiconductor Fabrication

SPARK BEYOND

Advanced analytics reduced failures and improved fab performance.

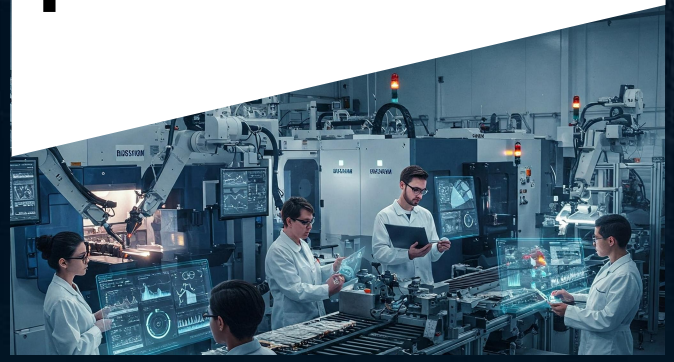

CHALLENGE

- **High data volume** from advanced fab instrumentation.
- **Limited insight** due to reliance on vendors.
- **Raw data underused** for performance improvement.
- Needed to **identify failure patterns** across systems.

RESULTS

- **87% accuracy** in predicting failure events.
- **14% false positives**, improving reliability.
- **86% correct** predictions of normal operation.
- Enabled **data-driven operations** and maintenance.

APPROACH

- **Integrated** and cleaned diverse data sources.
- **Modeled** key performance and failure drivers.
- **Extracted** explanatory variables from large datasets.
- **Validated** predictions and implemented improvements.


Thank You
SPARK BEYOND