IBM AI Openscale

Learn how to operationalize AI across its lifecycle with the right open platform for your business

94% of companies believe AI is key to their competitive advantage (IDC), but only 1 in 20 has extensively incorporated AI into their business processes and workflows (MIT Sloan Management Review). Why?

AI requires new management considerations
AI is a learning system that continuously evolves and learns based on experience and the data it encounters. This introduces a new set of operational and cultural paradigms that must be considered. AI systems change behaviors over time, which needs to be understood, controlled, and accepted. They also need data from production systems to "feed" the learning loop, which introduces new considerations around data security, operational policies, and DevOps.

Different teams build and run AI systems
Those who build AI models have skills in statistics and math, but those who run and support applications have critical IT and DevOps support skills. Bridging this gap to run a continuously evolving system is difficult and requires new skills and tools.

AI is often a black box
Businesses do not have the confidence to trust and explain how certain decisions or recommendations were reached, and they can't always ensure fair outcomes for customers.

Building and retraining AI takes time
The process of building custom models can stretch over the better part of a year—six to nine months is common—which is a significant time and resource burden for data science teams. That's to say nothing about the extensive process of running, maintaining, and retraining models on new data, or about the fact that by the time the models are built, the business opportunity may have passed.

Overview

The IBM AI OpenScale platform allows enterprises to automate and operate AI at scale—wherever it resides, across its entire lifecycle—with transparent, explainable outcomes, automatically freed from bias. AI OpenScale provides confidence in AI outcomes, and it spans the gap between the teams that operate AI and the business units that use these applications.

Features

Tooling and deployment flexibility
— Monitor models developed in Watson Studio or any IDE or open source framework, and hosted in WatsonMachine Learning or any third-party or private model serve engine.
— Deploy behind an enterprise firewall, wherever your data and applications reside.

Model fairness
— Continuously monitor how AI makes predictions about individual transactions.
— Detect and automatically mitigate bias at both build and runtime to ensure fair outcomes.
— Synthesize a de-biased model, executing in the background and running alongside currently deployed models, to generate trust prior to deployment.
— Configure fairness metrics for continuous evolution, providing quality checks so that no biased data or models can drive unfair decisions in business applications.
— Visualize the fairness score of deployed models, so model builders can act swiftly to mitigate bias.

Operations dashboard
— Measure model accuracy for its ability to deliver outcomes aligned with accuracy established by knowledge workers.
— Proactively troubleshoot underperforming deployments.
— Continuously evolve AI models embedded in business applications.
— Discover data to be labeled to improve AI outcomes.

Explainability
— Explore the factors that influenced an AI outcome, in easy-to-understand business terms.
— Understand how changing factors would produce improved outcomes through contrastive explanations.

Traceability and auditability
— Trace AI recommendations to business outcomes.
— Gain insights into AI and application behavior by capturing the request sent to the AI model or Python function, along with statistics about its health.

Open data mart
— Provide operations staff with access to third-party reporting tools for customized dashboarding.
— Integrate model insights with business KPIs in existing enterprise data marts, data warehouses, and data lakes.

Model ops
— Script or configure models into application deployment pipelines to enable actions on models, including start, stop, and redeployment.

Neural Network Synthesis (NeuNetS)
— Automatically design a new, customized neural network, optimized for a particular use case.

Top use cases of AI OpenScale

Insurance

Problem
Insurance underwriting is becoming increasingly competitive and faces significant challenges. Insurance professionals struggle with evaluating risk exposure due to lack of uniformity and control, and insurers are unable to efficiently process application submissions due to the reliance on unstructured data.

Solution
Machine learning models trained on historical customer and claims data can help underwriters make more consistent and accurate risk assessments. These models can provide price suggestions for individual customers based on different features in their profile. To adhere to regulatory compliance standard, insurers need to be able to explain how these models work. AI OpenScale's explainability features allow underwriters to see the exact considerations prioritized by the model in making a recommendation. Additionally, its bias detection and mitigation features help underwriters ensure these models continue to produce fair recommendations after they've been deployed.

Telecommunications

Problem
Effectively maintaining physical assets and infrastructure is essential for telcos. Asset failure can lead to service outages, a key cause of customer churn. Maintenance prioritization is an expensive and difficult process, and it's often challenging to catch asset failure in the field before it leads to a problem on the network.

Solution
Machine learning models trained on historical asset failure data—including sensor data, images captured from drones, and old maintenance reports—can help telcos predict asset failure before it causes an outage. Network engineers and IT operations need to ensure their AI models are accurately predicting failure, but the data are incredibly complex. AI OpenScale's traceability features allow teams to monitor the health of their models at runtime and connect business outcomes with predictions, so they can be confident in the accuracy of their models.

Supply Chain

Problem
Effective demand forecasting is essential to keeping operational costs down while still meeting consumer demand, but it's very difficult. Companies aren't equipped to deal with the volume and diversity of data needed to account for real-time changes in demand. Forecasts that can't adapt to constantly changing variables in today's market can lead to multimillion-dollar miscalculations, severely damaging a company's bottom line.

Solution
Machine learning models can be trained to help a company improve its Forecast Value Added (FVA) metrics, learning from historical, successful, and unsuccessful forecast override data. These models help demand planners make better adjustments to their demand forecast. AI OpenScale's model dashboard allows the supply chain team to monitor their models' accuracy over time, so that they can check that their AI-powered applications are consistently delivering outcomes as accurate as those produced by knowledge workers.

Financial Services

Problem
Banks spend up to one-fifth of their operating costs on anti-money laundering efforts, and failure to do so effectively can result in massive fines from government regulators. Most institutions employ rules-based systems to detect potentially suspicious activity, but the data are so complex that banks must employ a large workforce of analysts to manually comb through rule-generated alerts— 95-98% of which are false positives.

Solution
Banks can use machine learning and deep learning to prioritize and analyze AML alerts. Models trained on historical transaction data can identify suspicious patterns that a rules-based system would miss, because they learn at the same rate as that at which the bad actors learn to beat the system. AI OpenScale's explainability and traceability features help banks keep up with ever-changing regulations, and allow financial crime analysts to understand the reasoning behind their models' alert analysis, so they can make better decisions about which alerts to dismiss and which to escalate.