



www.ibm.com

590 Madison Avenue, New York
NY 10022, USA
Tel: +1 800 426 4968
Email: askibm@vnet.ibm.com

IBM Watson OpenScale

The company

International Business Machines Corporation (NYSE: IBM), is an American multinational technology and consulting corporation, with headquarters in New York. IBM manufactures and markets computer hardware and software, and offers infrastructure, hosting and consulting services in areas ranging from mainframe computers to nanotechnology. The company was founded in 1911 and over the years has grown through acquisition into its present state.

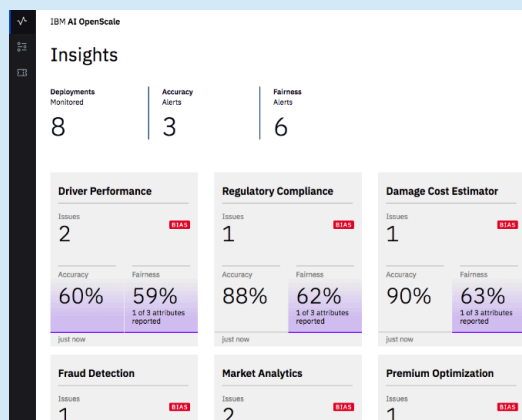


Figure 1 – OpenScale dashboard

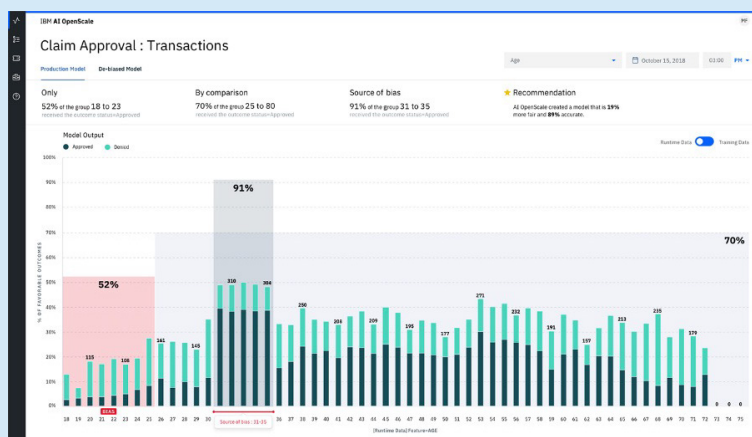


Figure 2 – OpenScale detect and mitigate function

IBM has a global presence, has initiated various emerging open standards and open environment initiatives (such as Eclipse, Jazz, OSLC, RELM and so on; some more open than others) and owns many significant brands such as Db2, WebSphere, Watson, System z and so on.

“ IBM also makes available, as open source products, AI Fairness 360, which provides bias detection and mitigation, and the Adversarial Robustness Toolkit, which helps to prevent model hacking. ”

What is it?

IBM Watson OpenScale is a platform that is specifically targeted at operationalising AI (augmented intelligence) models. It runs on both IBM Cloud and IBM Cloud Private, and it integrates with any tools for building and running AI, from competitive offerings like Microsoft Azure ML and Amazon ML, to open source environments such as TensorFlow or Spark, to IBM’s own tools, Watson Studio and Watson Machine Learning.

What does it do?

Watson OpenScale has various capabilities that enable business to operate and automate AI across its lifecycle, with the aim of bringing enterprise-scale AI more easily within reach. First, the product allows you to monitor how the deployment of AI is affecting business outcomes, whether positively or negatively, no matter where models were developed or deployed.

Figure 1 shows a dashboard that is used to automatically detect unfair or inaccurate results based on user-defined conditions. As can be seen in this example, “driver performance” has two issues: it is both inaccurate (that is, performing badly) and is producing unfair outcomes. In total there are three accuracy alerts and six bias alerts.

A second function of Watson OpenScale is to detect, and automatically mitigate, bias: see [Figure 2](#). Various companies have announced such capabilities for training purposes but, as far as we know, Watson OpenScale is the only such product that is also available at run-time, when it compares the performance of the deployed models with recommended de-biased models. No retraining is required. However, it will provide a way for a user to download and inspect the de-biased model before deploying it into production. Thirdly, in addition to audit capabilities that allow you to trace the recommendation any decision, Watson OpenScale provides an explainability function. This is illustrated in [Figure 3](#). It allows you to explore the factors involved in any decision from AI, how much those factors contributed to the decision and how different factors might have changed the outcome. As can be seen, this also provides a confidence level for the transaction in question, and the ability for the business to explain a decision to a customer, regulatory body or other party.

“ Watson Machine Learning (and IBM SPSS) provides complementary facilities for monitoring model performance, model management capabilities, and the Experiment Assistant, which is a suite of tools for managing model training. Workflow to support functions such as approval processes are built into the various IBM platforms upon which Watson OpenScale runs. ”

Under the covers the product also provides what is known as payload logging. This means that all the data being scored is logged and persisted to a database and it is against this data that the various dashboards and other capabilities of Watson OpenScale are based. You can query the database using IBM Cognos or using third party tools such as Looker or Tableau, in order to get performance metrics. In addition, these metrics can also be exported and combined with application metrics to measure business outcomes.

Why should you care?

AnalyticOps concerns the operationalisation, monitoring, improvement, management and governance of augmented intelligent models. As AI becomes more widely deployed, AnalyticOps capabilities will become more vital. Where most companies today have a handful of deployed models and applications, if that, we expect large organisations to have thousands or tens of thousands of models to deploy and manage in the not-too-distant future. This is going to be impractical without tools of the type described here.

The Bottom Line

IBM is one of only a handful of vendors to offer end-to-end model management, despite the fact that this is a technology that has been around for a long time. It is one of even fewer companies that is actively addressing operationalisation in order to connect the work done by a data scientist to business owners, in business terms. Indeed, IBM is taking a lead in so far as its ethical positioning is concerned. We know of no other company that offers bias detection and mitigation along with explainability tools. We are very impressed.

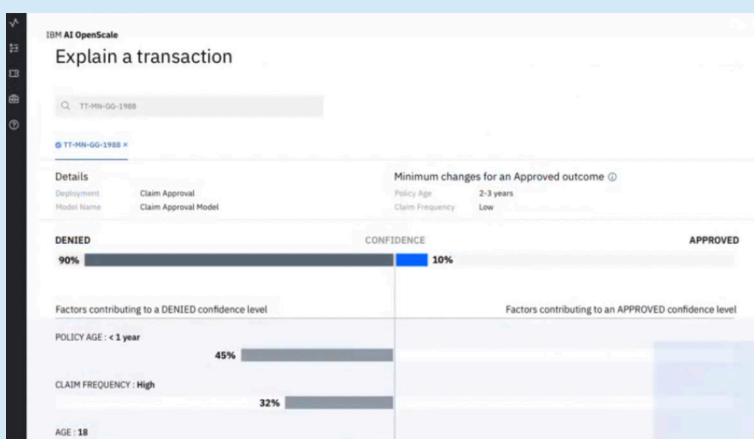


Figure 3 – OpenScale explainability function

FOR FURTHER INFORMATION AND RESEARCH [CLICK HERE](#)