

REPORT REPRINT

IBM Watson OpenScale builds trust in AI via monitoring, fairness and explainability

MAY 01 2019

By Paige Bartley

At the IBM Think conference in early 2019, 'trust in AI' was a central theme. The IBM Watson OpenScale platform, available in the cloud and on-premises, offers tools to monitor model performance, facilitate explainability and detect and mitigate fairness issues for models in production, allowing the enterprise to manage AI implementations at scale.

THIS REPORT, LICENSED TO IBM, DEVELOPED AND AS PROVIDED BY 451 RESEARCH, LLC, WAS PUBLISHED AS PART OF OUR SYNDICATED MARKET INSIGHT SUBSCRIPTION SERVICE. IT SHALL BE OWNED IN ITS ENTIRETY BY 451 RESEARCH, LLC. THIS REPORT IS SOLELY INTENDED FOR USE BY THE RECIPIENT AND MAY NOT BE REPRODUCED OR RE-POSTED, IN WHOLE OR IN PART, BY THE RECIPIENT WITHOUT EXPRESS PERMISSION FROM 451 RESEARCH.



Introduction

The IBM Think Conference in February had three broad themes: digital and AI strategy, hybrid and multi-cloud architecture, and trust in AI and corporate stewardship. The third of these themes is where IBM is leveraging its existing ecosystem of AI and machine learning products to offer capabilities that help organizations ensure model performance, explainability, and detection and mitigation of bias in models in production. The IBM Watson OpenScale platform, which became generally available in late 2018, was specifically highlighted at the conference for its ability to achieve these objectives.

Worldwide, transparency and explainability are becoming table stakes for AI and machine learning, driven by increasing regulatory pressure and consumer rights, as well as the business need to simply understand how and why certain outcomes were reached. Visibility into models and their performance is needed not only for defensibility but for scalability as well.

451 TAKE

‘Trust in AI’ was perhaps the least emphasized of the three main themes at the IBM Think conference in 2019, but is the one in which IBM likely has the greatest potential of showing product and brand leadership. IBM Watson OpenScale’s capabilities and ability to integrate into IBM’s existing offerings such as IBM Watson Studio and Watson Machine Learning allow the enterprise to address the explainability and transparency requirements that are difficult to achieve given the increasing complexity of models and interconnected ecosystem of tools.

Transparency will be critical to both regulatory compliance and efforts to scale and operationalize. Watson OpenScale’s focus on explainability for business users creates visibility and control for an important end user audience, helping bridge the AI talent gap and assisting in efforts to confidently deploy fair models in production scenarios.

Context

The genesis of IBM Watson OpenScale began in Q4 2017 with IBM’s recurring Global Technology Outlook program, led by IBM Research. Out of that work, which has provided insight into long-term business technology trends for decades, explainability and interpretability emerged as key themes for the year. Businesses were looking to implement AI and machine learning, but the ability to have transparency regarding the result of models was increasingly a concern for deployment at scale and in production environments. Internal factors, such as the need for repeatability for business performance as well as external factors such as regulations like the EU’s General Data Protection Regulation (GDPR), put pressure on organizations to mitigate bias and explain results.

IBM’s development of Watson OpenScale began in early 2018, and by mid-year, the company began to focus on fairness in models. Because models are inherently trained on data sets generated by humans, the most accurate models will accurately reflect human bias, causing legal and regulatory repercussions. Therefore, product capabilities around identification and mitigation of bias, and adjustment of models to make them fair – rather than just perfectly accurate – were necessary. Watson OpenScale was announced in mid-October 2018 and became generally available in December 2018.

REPORT REPRINT

The IBM Think conference in February, then, was the product's first major 'hands-on' event with the public. Development of the product is ongoing and is aligned with IBM's focus on open architecture and new 'Watson Anywhere' deployment methodology announced at the show. At the conference, SPSS collaboration and deployment services for Watson OpenScale was a major announcement for the product, indicating focus on developing further integration with the existing IBM portfolio.

Product

IBM Watson OpenScale is an open platform to operate and automate AI across its lifecycle, with a focus on providing tools for explainability and bias mitigation, as well as monitoring, for models in production. The core value is to provide confidence in AI and ML to business stakeholders, helping connect model outputs to business KPIs. While it is highly complementary and integrates tightly with the IBM Watson Studio model building and training IDE as well as with Watson Machine Learning to deploy models, models can also be built in other third-party frameworks such as TensorFlow, Scikitlearn, Keras, SparkML and PMML. With IBM's increased emphasis on open architecture, Watson OpenScale supports AI deployed in various run-time environments, such as Amazon ML, Azure ML and custom run-time environments behind the enterprise firewall. Watson OpenScale itself is currently available on IBM Cloud and IBM Cloud Private for Data.

While Watson Studio is geared more toward data science practitioners and subject matter experts, Watson OpenScale provides a highly visual, drill-down interface so that data-savvy business users can explore the effects of variables on models and adjust as necessary to meet certain desired or regulatory-driven objectives for fairness and bias mitigation. In addition, there is a flexible, open data mart that allows metrics to be exported out or explore in place using any SQL-based BI tool. Because business roles frequently interface with the enterprise's compliance workflows and interact directly with customers, having explainability tools and metrics geared toward this audience is important for successfully scaling AI processes within the organization and bridging the AI skills gap.

Watson OpenScale provides a model monitoring framework and helps give various business stakeholders visibility into AI in three primary ways:

- Track and measure business outcomes in production
- Meet regulatory constraints and govern AI in production
- Adapt AI to changing business situations in production

Functionality is geared toward these three objectives. The product logs the input and the output of the model and can also send dummy inputs through the model to see how it would behave under certain circumstances. Users can configure bias thresholds for regulatory or business purposes and bias percentages can be calculated. Tracking of ongoing performance of the model is also performed, with data collected and analytics being provided on model performance over time. Users can additionally toggle their visualization between training data and runtime data in order to ascertain where bias originated: to see if the original training data was biased or if the model evolved during runtime.

Roadmap

Immediate plans for the IBM Watson OpenScale roadmap fall broadly into two categories: integration with both IBM products and other environments and leveraging AI directly within the product to improve usability and outcomes.

The first, and most obvious, area of development will be deeper integration into IBM's existing suite of products, particularly for governance. Because knowing the lineage of models is important for scaling AI efforts while providing appropriate governance, integration with the existing Watson Knowledge Catalog is planned. The goal with this integration would be to provide users with one-click access and visibility into model lineage from the Watson OpenScale interface, with lineage information automatically being sourced from the catalog. Integration with other non-IBM statistical and advanced analytics platforms is also on the immediate roadmap.

The use of machine learning and AI within the product itself is also an area of focus, reflecting the trend in the broader enterprise software market. This concept of using AI to improve AI can take several forms. Existing capabilities include a model within the product that is actively learning how the business's models are making decisions. But on the roadmap are plans for additional uses of embedded AI used to drive 'smart' recommendations and to identify next-best actions for users, improving usability.

Competition

In the data science space at large, IBM Watson OpenScale competes with a wide range of open source offerings, although traditionally, explainability and bias mitigation have often been the 'missing link' in data science pipelines built from open source products. The free ML.NET machine learning library, in its 0.8 build, included model explainability APIs (originally used internally at Microsoft) that allow for the understanding of feature importance of models. TensorFlow offers some explainability capabilities for certain model types, such as boosted trees, but does not have a dedicated interface for isolating variables and drilling down into sources of bias. A key difference with OpenScale vs. these other options in the market, however, is that it is optimized for business users to evaluate explainability and model performance at runtime.

Prominent commercial players in the cloud and AI space offer natural competition to IBM. Microsoft, with its Azure Machine Learning, offers a studio environment to build, deploy and share predictive analytics solutions. While explainability is currently not a prominent product focus, Microsoft's collaboration with ZestFinance in late 2018 was formed to create explainability solutions for highly regulated finance use cases. Amazon SageMaker on AWS for building, training and deploying models, might potentially offer more direct competition to IBM if it were to add in additional functionality focused on explainability and bias mitigation, particularly if it were to focus on a business user (rather than technical) audience.

Smaller niche players focused on management and governance of the machine learning process also provide competition. Immuta, which specializes in data management for data science, provides a granular policy engine that allows conditions to be set for how data is used in certain models, and mechanisms for access controls. This upstream approach focuses on mitigating bias and protecting data privacy before models are deployed, although explainability in runtime, offered by OpenScale, may offer stronger regulatory and compliance use cases. Interpretable AI takes a solution-centric approach, using explainable models and algorithms for use cases in specific, largely regulated, industries.

SWOT Analysis

STRENGTHS

Few products share IBM Watson OpenScale's purpose-built focus on model explainability, bias mitigation and monitoring, which are all increasingly important in production use cases. The business user audience bridges an important skill gap and helps integrate into enterprise-wide AI workflows.

WEAKNESSES

Integration efforts are still not currently complete for IBM Watson OpenScale within the full IBM ecosystem of offerings. This is especially notable for broader data governance and lineage functions, though more extensive integration, such as with Watson Knowledge Catalog, is on the immediate roadmap.

OPPORTUNITIES

The regulatory environment and court of public opinion are increasingly demanding AI explainability. By being early to the market with a dedicated product built on a rich legacy of research, IBM can position itself as a leader for 'trust in AI' and ethical frameworks for fair use.

THREATS

Major cloud providers, with their own AI and ML offerings, are also publicly discussing fairness and explainability. With the need evident, it is just a matter of R&D to build in competing functionality. Despite IBM's open approach, customers may prefer a 'one-stop-shop.'