

IBM Goes Full Steam Toward Eliminating AI Problems



by
Torsten Volk
Enterprise Management Associates

Organizations that quickly transform the exponentially growing internal and external data into actionable knowledge individualized for each employee will outcompete the slower ones. Employees can make data-driven decisions based on their roles and individual situations. The more an enterprise replaces implicit assumptions with data-driven ones, the better each decision will be.

To get to the point where AI can be successful, AI problems must be addressed. "The ultimate competitive advantage is to create an organization that can outlearn everyone," said Ginni Rometty, IBM's CEO, at the IBM Think 2018 event in Las Vegas, where she explained how specific IBM tools tackle AI bottlenecks.

AI problem No. 1: Obtain the training data

The availability of sufficient training data is critical for the success of any AI project. Often, AI projects fail at the data collection stage because training data is only available in an unstructured format that must be cleaned before the data can be used.

IBM offers a set of new tools that address both challenges:

- Data Refinery, which is included in Watson Studio, offers a simple visual interface to connect to and explore most unstructured and structured data sources. For example, it can check the reputation and expertise of all parties

involved in a negotiation and include as much background information about each individual person as is available openly online. It includes connection wizards for Twitter and LinkedIn. A user can download freely accessible information about the other parties and then run a sentiment analysis on this data to find out if they are likely to have a positive attitude toward the company's brand, if they have a good reputation in the marketplace or if they are affiliated with a competitor.

- Watson dashboards enable business analysts to quickly weed through vast amounts of data from different sources in search of interesting correlations. Like Data Refinery, Watson Dashboards lower the threshold for business staff and full-stack developers to work with data and accomplish many tasks without requiring help from a data scientist.

- Watson Data Kits enable enterprises to fill in data gaps to help add enhanced, AI-driven capabilities to other apps. For example, IBM now offers a data set that includes 300,000 restaurant menus in 21,000 cities. Based on this data and on demographic data obtained through a service like NeighborhoodScout, a user could create a map that shows viable locations for restaurants based on demand derived from the demographic data and competitive menus shown in the Watson Data Kits. This data could be further enriched by adding ratings from Yelp to increase the weight of higher-rated competitors.

AI problem No. 2: Trust in security

Data owners often feel uncomfortable consenting to the

use of their data assets due to concerns about data privacy. The recent Facebook breach has shown that to surrender control over data sets can lead to severe legal and PR consequences. Therefore, many promising AI projects never get off the ground due to data owners' inability to gauge their own security risk of signing over access to the project team.

IBM Cloud Private for Data, delivered in container format for rapid deployment to any Kubernetes platform, provides a centralized layer for data governance, integration and analysis to abstract the management of data from applications and APIs. This architectural separation addresses a big part of the trust issues to make data available for AI analysis, as the new platform can transparently enforce data privacy and deliver compliance reports to all stakeholders.

AI problem No. 3: Shortage and cost of data scientists

For everyone who believes that every business and IT process can benefit from AI, the limiting factor for most AI projects in terms of cost and availability is data scientists. Two goals must be clear to reduce the dependency on data scientists:

1. Help data scientists be much more effective with collaboration and data science tools to provide their critical services faster for self-service consumption by the business.
2. Help developers and business analysts solve certain problems without the help of data scientists.

Watson Studio provides visual tools and templates that support developers and business analysts to define, train and test their own machine learning and AI (ML/AI) models. If they get stuck, they can pull in data scientists for help through the built-in collaboration capabilities. The tools can also make data scientists more effective by providing pointers about what is important within various data sets.

Watson Studio enables almost everyone to get their hands dirty with data analysis and AI, without the need to stand up their own complex environments. Some definitions and training of deep learning models may be best left to data scientists, but business staff can create image categorizers, sentiment analyzers or text analysis tools without help. And the more they get their hands dirty, the more new ways of leveraging the 80% of currently untapped data they will come up with.

AI problem No. 4: Stand up and operate the environment

Requesting new ML/AI server environments is a hassle and leads business units to either engage in shadow IT by moving data to public cloud services, such as AWS SageMaker, or to simply abandon a project entirely.

IBM Cloud Private for Data and Watson Studio are integrated to leave data on premises as required by compliance policies but still use IBM's hosted GPUs to train ML/AI models quickly and cost-effectively, without the need to buy hardware. This lowers the threshold for experimentation, which is critical to spur interest around the potential of ML/AI.

AI problem No. 5: Integration between data center and cloud

Based on recent Enterprise Management Associates (EMA) research, the lack of integration of management tools for bare metal, virtual machines (VMs) and containers that are hosted in different data center and cloud locations is a crucial pain point for enterprises. The lack of a unified security model for these different deployment options and the inability of current IT staff to support all of these options contribute to this problem. To effectively utilize ML/AI with its tremendous hunger for data and CPU capacity, it is critical to eliminate any kind of management silos within corporate IT.

IBM Cloud is based on a centralized architecture for the modular addition of management, data and application services through Kubernetes containers. This architectural paradigm helps customers to centrally deploy and manage applications in the data center, IBM Cloud and further Kubernetes-based public clouds. Ultimately, every enterprise needs to get to a stage where only policy requirements determine whether an application or data is hosted on bare metal, containers, VMs, platform as a service or framework as a service and if that application should remain on premises or move to the public cloud.

IBM's Power9 server chip is optimized for ML/AI algorithms, and the company claims that this new CPU architecture completes the standard clickstream benchmark 46 times faster than what Google can offer today. While EMA has not confirmed this claim, customers can simply try out these new chips through IBM Cloud without significant risk to determine business viability.

IBM's new "Let's put smart to work" slogan translates into a product strategy for solving various AI problems. IBM needs to compete with the also very aggressive ML/AI strategies of AWS, Microsoft and Google.