

Machine Learning in 2019: Getting it Right from the Cloud to the Edge



Machine learning and artificial intelligence will be front and center in all business processes

As machine learning (ML) and artificial intelligence become more pervasive, data logistics will be critical to your success.

In *Machine Learning Logistics: Model Management in the Real World* (O'Reilly, 2017), the authors note that 90 percent of the effort required for success in machine learning is not the algorithm, model, framework, or the learning itself. It's the data logistics. Perhaps less exciting than these other aspects of ML, it's the data logistics that drive efficiency, continuous learning, and success. Without data logistics, your ability to continue to refine and scale are severely limited.

Good data logistics does more than drive efficiency. It is fundamental to lower costs now and improved agility in the future. As ML and AI continue to evolve and expand into more business processes, enterprises must not allow early successes to turn into limitations or problems long-term. In a paper by Google researchers (*Machine Learning: The High Interest Credit Card of Technical Debt*), the authors point out that although it is easy to spin up ML-based applications, the effort can result in costly data dependencies. Good data logistics can mitigate the difficulty in dealing with these complex data dependencies to avoid hampering agility in the future. Using a proper foundation like this can also ease deployment and management as well as allow the evolution of these applications in ways that are impossible to foresee precisely today.

In 2019, we'll see a shift from complex, data-science-heavy deployments to a proliferation of initiatives that can be best described as KISS (Keep It Simple to Start). Domain experience and data will be the drivers of AI processes that will evolve and improve as experience grows. This approach will offer another advantage: it also improves the productivity of existing personnel as well as expensive, hard-to-find, -hire, and -retain data scientists.

This approach also eliminates the concern over picking "just the right tools." It is a fact of life that we need multiple tools for AI. Building around AI the right way allows continuous adjustment to take advantage of new AI tools and algorithms as they appear. Don't worry about performance, either (including that of applications that must stream data in real time) because there are continual advances on that front. For example, NVIDIA recently announced RAPIDS, an open source data science initiative that leverages GPU-based processes to make the development and training of models both easier and faster.

Multi-Cloud Deployments Will Become More Commonly as a Way to Prevent Lock-In

As organizations expand their use of ML and AI across multiple lines of business, they will need to access the full range of data sources, types, and structures on any cloud

while avoiding the creation of data silos. Achieving this outcome will result in deployments that go beyond a data lake, and 2019 will mark the increased proliferation of global data platforms that can span data types and locations.

Organizations will move to deploy a commonly data platform to synchronize and drive converge of (and optionally preserve) all data across all deployments, and through a global namespace provide a view into all data, wherever it is. A commonly data platform across multiple clouds will also make it easier to experiment with different services for a variety of ML and AI needs.

To be sufficiently agile for whatever the future might hold, the data platform will need to support the full array of disparate data types, including files, objects, tables, and events. The platform must make input and output data available to any application anywhere. Such agility will make it possible to fully leverage the global resources available in a multi-cloud environment, thereby empowering organizations to achieve the cloud's full potential to optimize performance, cost, and compliance requirements.

Analytics at the Edge Will Become Strategically Important

As the Internet of Things (IoT) continues to expand and evolve, the ability to unite edge, on-premises, and cloud processing atop a commonly, global data platform will become a strategic imperative.

A distributed ML/AI architecture capable of coordinating data collection and processing at the IoT edge eliminates the need to send massive volumes of data over the WAN. This ability to filter, aggregate, and analyze data at the edge also facilitates faster, more efficient processing and can result in better local decision making.

Organizations will strive to have a commonly data platform -- from the cloud core to the enterprise edge -- with consistent data management to ensure the integrity and security of all data. The data platform chosen for the cloud core will, therefore, be sufficiently extensible and scalable to address the complexities associated with distributed processing at a diffuse and dynamic edge. Enterprises will place a premium on a "lightweight" yet capable and compatible version appropriate for the compute power available at the edge, especially for applications that must deliver results in real-time.

A Final Word

In the next year we will see an increased focus for AI and ML. Enterprises will keep it simple to start, avoid dependencies with a multicloud global data platform, and empower the IoT edge so ML/AI initiatives deliver more value to the business in 2019 and well into the future

The State of AI in South Africa



As part of the ongoing process of rebuilding after decades of political unrest and economic uncertainty, companies across South Africa are engaged in the development of a nascent artificial intelligence (AI) sector. So, what have been some of the most recent developments in the AI sector in the country? What is the long-term potential for South Africa to emerge as a global center of expertise in AI development? What strengths and weaknesses does the country currently possess in this area? And how best might IT professionals contribute to the ongoing development of the country as a global hub?

Growing sector

In recent years, a growing number of AI companies have emerged in South Africa - and there are signs that the country is gradually building up a solid base of expertise in the technology. One interesting example is Cape Town-based AI Software and Solutions company Cortex Logic, which works across a range of industry sectors to augment and automate legacy processes, ultimately transforming them into intelligent systems via its AI Engine. As Jacques Ludik, Founder and Managing Director at the company, and also Former Former President of the Machine Intelligence Institute of Africa (MIIA), explains, the Engine works by solving 'strategic and operationally relevant problems by unlocking the value offered by the fourth industrial revolution and mobilizing data science, Internet of Things (IoT), Big Data and analytics'. Another example is AI start-up DataProphet, a machine learning-based solution that analyses and interprets a company's data using a prescriptive AI system - and enables production teams to eliminate defects and scrap, and minimize downtime by prescribing the optimal parameter control limits and set points.