



Market Insight Report Reprint

Assessing the field in the MLOps race

October 15 2021

by **Krishna Roy**

Machine learning operationalization blossomed this year as organizations have made machine learning projects a critical part of their digital transformation initiatives. Software vendors are seeking to get in on the MLOps action, which has made it into a vibrant but somewhat confusing product sector, which we aim to demystify.

451 Research

S&P Global

Market Intelligence

This report, licensed to Dominion Data Labs, developed and as provided by S&P Global Market Intelligence (S&P), was published as part of S&P's syndicated market insight subscription service. It shall be owned in its entirety by S&P. 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 S&P.

Introduction

Products to support the operationalization of machine learning models – aka MLOps tools – are increasingly being adopted across the enterprise. Almost three-quarters (74%) of enterprises have either invested in MLOps tools or planned to do so within the next 12 months, according to 451 Research's Voice of the Enterprise: AI & Machine Learning, Use Cases 2021 survey. For companies that have not yet made that investment, the choices are bountiful; MLOps tools are available from a wide variety of vendors. However, with that choice, comes ever more difficult purchasing decisions as the MLOps tools landscape becomes harder to navigate. Furthermore, this issue is made more difficult by the fact MLOps is currently suffering the fate of every hot tech buzzword – confusion over its meaning and value.

THE TAKE

The MLOps tools market is still relatively immature – albeit evolving quickly – and every approach has its benefits and pitfalls, adding another layer of complexity. Nonetheless, MLOps needs to be understood because it is a linchpin for the ethical and responsible use of machine learning. While the adoption of MLOps products won't guarantee this outcome, they will help. Continually ensuring that machine learning models remain accurate, up-to-date and fit for purpose, as well as making certain they are explained, are prerequisites to ensure models are used fairly. That means the predictions they make aren't biased or misinformed, which has ramifications not only for the company using them to make decisions but for society at large.

Context: ML what?

As we explored in a report earlier this year, MLOps originally concerned itself with the deployment of machine learning models into production and their ongoing monitoring and retraining, which is known as operationalization – hence MLOps. It didn't refer to the earlier stages of the ML development process such as data acquisition, integration and annotation, or the initial training of the model – as is the case today. The expanded definition of MLOps means more vendors can address it, bringing more choice to organizations that haven't already invested in it. But it also makes for a confusing and complex vendor landscape. We are not intending to provide an exhaustive list of offerings available but rather a look at various approaches – as well as a window into the drivers behind its adoption.

Market drivers

A key driver behind the increasing need for MLOps is the shift from static manual development to adaptable machine learning systems, which are more complex in nature, but were also designed to address many of the problems associated with a manual and static approach.

For most data scientists and other machine learning practitioners, training, validating and deploying one model into production, as well as performing other tasks associated with its workflow, is straightforward and can be carried out manually. The problems start to mount up when the volume of models they craft significantly increases – and the machine learning pipelines become more complex – because it creates scalability issues from a technical as well as a cultural perspective. Data scientists understandably don't want to spend their valuable time and expertise 'babysitting' models in production. That means IT is often left with the task – if the organization hasn't already recruited ML engineering personnel – or data scientists haven't agreed to take it on. IT is obviously reluctant to become over-burdened but sometimes has little choice, which creates more issues.

Furthermore, a manual machine learning approach is flawed in other ways. It lacks the robustness that MLOps tools promise to deliver and the automation that makes certain steps easier and quicker. Equally important, MLOps paves the way for adaptable machine learning systems that dynamically adjust to changes in data and other factors, making models stay relevant. Manual approaches, in contrast, risk causing stale models because they need updating manually.

MLOps specialists

Many MLOps specialists focus on the ‘last mile’ of machine learning involving operationalization processes once models have been created, which is essentially the classic purview of this discipline. Others focus on model optimization, which is the stage after models have been trained and before they are deployed. Either way, using them may require an organization to purchase a separate offering for a specific part of the machine learning workflow, such as model creation, because most MLOps specialists of this ilk don’t see this process as part of their purview.

While specific capabilities clearly vary from vendor to vendor, MLOps specialists generally address a few key arenas. The ability to monitor data as well as model drift is one critical capability. Moreover, they are interconnected. Maintaining a handle on unexpected and undocumented changes to the data used to feed the model is as vital as keeping an eye on the model to ensure it doesn’t decay, because both will erode the overall success of the machine learning initiative.

Monitoring other aspects of the model, such as its features and accuracy, are other model-monitoring capabilities provided by specialists. Additionally, these vendors also focus on monitoring the hardware infrastructure necessary for models to be placed into production and kept running successfully in public clouds, on-premises and often a mixture of both. Equally important, model bias and ‘explainability’ to create confidence in the model are other vital capabilities addressed by specialists. Model optimization specialists – as the name implies – concentrate primarily on optimizing certain types of models, such as deep learning models, which have their own specific set of operationalization challenges.

Seldon, Datatron, Verta and Modzy are MLOps specialists. Iguazio also goes to market as an MLOps specialist. ModelOp is another purveyor of MLOps, although strictly speaking, the vendor is true to its name because it addresses the ModelOps discipline, which essentially involves embracing other model types too, such as rule-based and statistical models. Deeplite, Neural Magic and DarwinAI are model optimization specialists.

MLOps from the hyperscalers

Hyperscalers ‘get’ the overall software development process – be it machine learning-focused or not. This understanding has led them to build out portfolios that span the entire machine learning development process from data acquisition to model creation, deployment and ongoing retraining and monitoring – and eventually, most things in between.

By embracing the entire machine learning lifecycle, AWS, Google and Microsoft focus on providing organizations with a ‘one stop shop’ approach. However, AWS et al. also assume that an organization will want to use their respective cloud infrastructure only for machine learning initiatives to get these benefits, necessitating organizations to standardize on cloud infrastructure from one of these suppliers.

MLOps from purveyors of enterprise data science

Companies that provide enterprise data science platforms also increasingly address more aspects of machine learning than in the past. These vendors seek to make data science enterprise-ready by providing a set of open source and proprietary capabilities that are designed to embrace the tools, languages and frameworks data scientists already know and use in an integrated software stack for easier deployment, scalability and automation purposes.

Most enterprise data science platform vendors only addressed the data ingestion, preparation, training, and testing phases of machine learning projects until a couple of years ago - and didn’t focus on subsequent aspects of it at all. However, that’s no longer the case because an increasing number of enterprise data science platform providers have delivered capabilities to cater to the ‘last mile’ of machine learning, involving the deployment of models into production, as well as ongoing retraining and monitoring tasks.

Furthermore, some have done so using strategic tuck-in acquisitions, while others have taken an internal development route. For example, DataRobot has acquired MLOps specialists ParallelM and Algorithmia, while the likes of Domino Data, Dataiku and CognitiveScale have pursued homegrown development strategies to achieve this endgame. IBM, SAS Institute, H2O.ai, Alteryx, RapidMiner, KNIME and RStudio are other vendors in the enterprise data science platform sector.

Careful scrutiny of specific functionality in enterprise data science platforms on a vendor-by-vendor basis is required. While all share the same high-level objective of making data science corporate-ready by touting an end-to-end stack, they will have gaps, which is an inevitable consequence of any evolving tech market.

MLOps from data platform providers

Data platform providers have embraced machine learning as an analytical workload that is suited to the data platforms they already provide, thus providing 'value add' to existing customers already using them. For example, Cloudera provides machine learning offerings as a user experience for the company's Cloudera Data Platform (CDP). Cloudera delivers a development environment for data scientists as well as visual apps for non-experts. The company's approach to address data science for data scientists essentially involves offering a proprietary workbench or integration with popular open-source tools as part of CDP. Cloudera also provides capabilities to support explainable machine learning, as well as machine learning models in production.

Databricks is another example of a vendor that provides a data platform-centric approach to machine learning model development and operationalization as the vendor delivers machine learning as a use case and workload specifically for its Unified Analytics Platform. Databricks acquired German startup 8080 Labs earlier in October to add faster and easier exploration and transformation of data for non-expert data scientists, by reaching for 8080 Labs' bamboolib, which provides a graphical user interface to the Pandas data and analysis manipulation tool for the Python programming language. Databrick's pickup of 8080 Labs continues the company's strategy to address non-expert data scientist requirements through strategic tuck-ins. Databricks also nabbed open source dashboarding and visualization specialist Redash in 2020 with this objective in mind.

CONTACTS

The Americas

+1 877 863 1306

market.intelligence@spglobal.com

Europe, Middle East & Africa

+44 20 7176 1234

market.intelligence@spglobal.com

Asia-Pacific

+852 2533 3565

market.intelligence@spglobal.com

www.spglobal.com/marketintelligence

Copyright © 2021 by S&P Global Market Intelligence, a division of S&P Global Inc. All rights reserved.

These materials have been prepared solely for information purposes based upon information generally available to the public and from sources believed to be reliable. No content (including index data, ratings, credit-related analyses and data, research, model, software or other application or output therefrom) or any part thereof (Content) may be modified, reverse engineered, reproduced or distributed in any form by any means, or stored in a database or retrieval system, without the prior written permission of S&P Global Market Intelligence or its affiliates (collectively, S&P Global). The Content shall not be used for any unlawful or unauthorized purposes. S&P Global and any third-party providers, (collectively S&P Global Parties) do not guarantee the accuracy, completeness, timeliness or availability of the Content. S&P Global Parties are not responsible for any errors or omissions, regardless of the cause, for the results obtained from the use of the Content. THE CONTENT IS PROVIDED ON "AS IS" BASIS. S&P GLOBAL PARTIES DISCLAIM ANY AND ALL EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, ANY WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE OR USE, FREEDOM FROM BUGS, SOFTWARE ERRORS OR DEFECTS, THAT THE CONTENT'S FUNCTIONING WILL BE UNINTERRUPTED OR THAT THE CONTENT WILL OPERATE WITH ANY SOFTWARE OR HARDWARE CONFIGURATION. In no event shall S&P Global Parties be liable to any party for any direct, indirect, incidental, exemplary, compensatory, punitive, special or consequential damages, costs, expenses, legal fees, or losses (including, without limitation, lost income or lost profits and opportunity costs or losses caused by negligence) in connection with any use of the Content even if advised of the possibility of such damages.

S&P Global Market Intelligence's opinions, quotes and credit-related and other analyses are statements of opinion as of the date they are expressed and not statements of fact or recommendations to purchase, hold, or sell any securities or to make any investment decisions, and do not address the suitability of any security. S&P Global Market Intelligence may provide index data. Direct investment in an index is not possible. Exposure to an asset class represented by an index is available through investable instruments based on that index. S&P Global Market Intelligence assumes no obligation to update the Content following publication in any form or format. The Content should not be relied on and is not a substitute for the skill, judgment and experience of the user, its management, employees, advisors and/or clients when making investment and other business decisions. S&P Global Market Intelligence does not endorse companies, technologies, products, services, or solutions.

S&P Global keeps certain activities of its divisions separate from each other in order to preserve the independence and objectivity of their respective activities. As a result, certain divisions of S&P Global may have information that is not available to other S&P Global divisions. S&P Global has established policies and procedures to maintain the confidentiality of certain non-public information received in connection with each analytical process.

S&P Global may receive compensation for its ratings and certain analyses, normally from issuers or underwriters of securities or from obligors. S&P Global reserves the right to disseminate its opinions and analyses. S&P Global's public ratings and analyses are made available on its websites, www.standardandpoors.com (free of charge) and www.ratingsdirect.com (subscription), and may be distributed through other means, including via S&P Global publications and third-party redistributors. Additional information about our ratings fees is available at www.standardandpoors.com/usratingsfees.