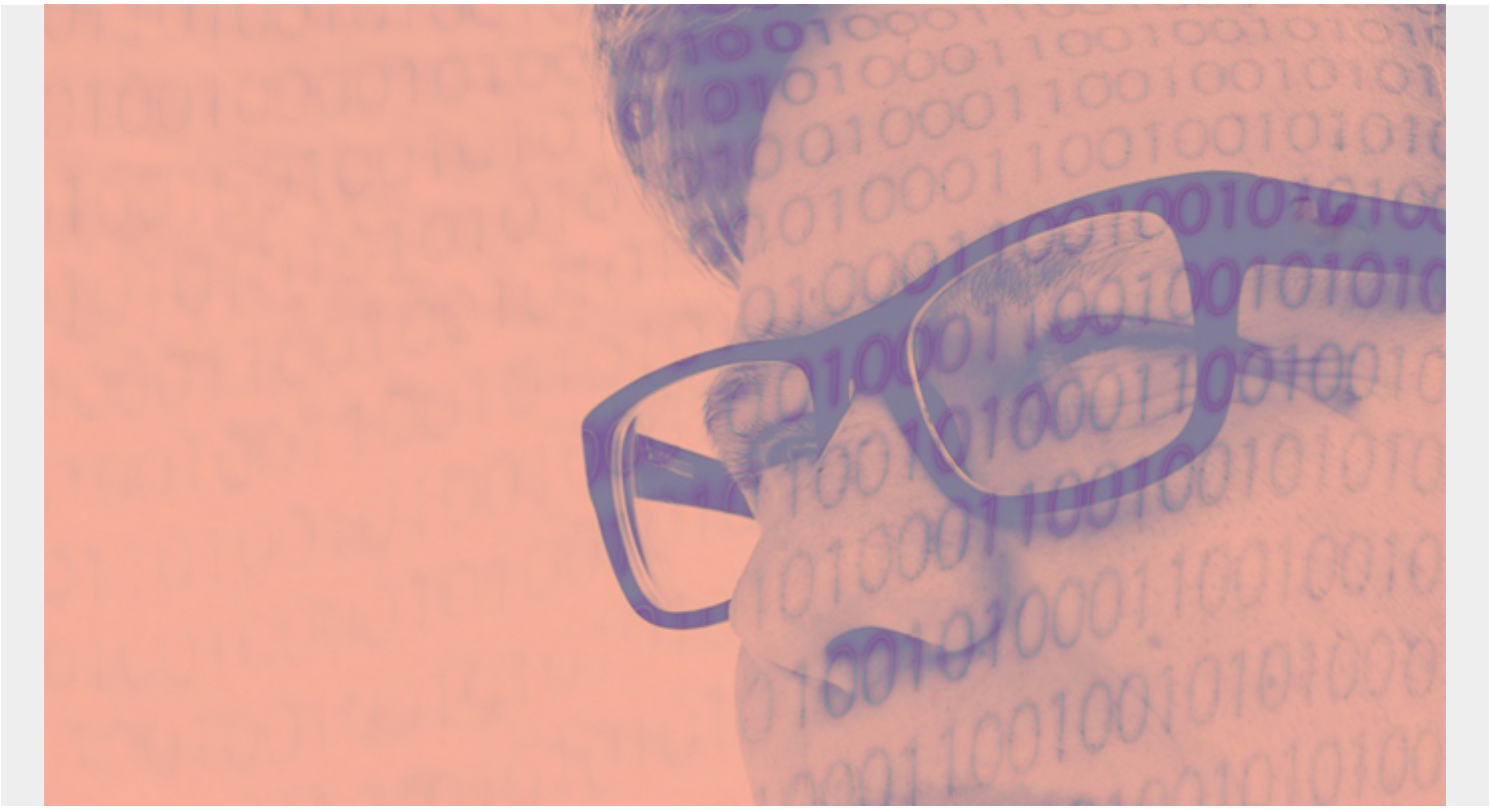
WHAT IS MLOPS? MACHINE LEARNING OPERATIONS EXPLAINED



If you are part of an IT or data team at any growing organization, you're familiar with the term machine learning.

Actually a method of computer function improvement that has been around since the 1950s, until recently—2015 to be exact—many people didn't understand the power of ML. But, with the influx of data science innovations and advancements in AI and compute power, the autonomous learning of systems has grown leaps and bounds to become an essential part of operations.

As Run.AI explains:

> *"Today, ML has a profound impact on a wide range of verticals such as financial services, telecommunications, healthcare, retail, education, and manufacturing. Within all of these sectors, ML is driving faster and better decisions in business-critical use cases, from marketing and sales to business intelligence, R&D, production, executive management, IT, and finance."*

The possibilities are endless and the result is that many organizations dedicate entire teams to ML operations. In this post we'll take a look at Machine Learning Operations (MLOps), including:

- Concepts
- Benefits
- Challenges
- MLOps best practices
- And more

Deciding if your organization is ready for an MLOps team starts here.

What is MLOps? 3 components of MLOps

MLOps is defined as "a practice for collaboration and communication between data scientists and operations professionals to help manage production ML (or deep learning) lifecycle. Similar to the DevOps or DataOps approaches, MLOps looks to increase automation and improve the quality of production ML while also focusing on business and regulatory requirements."

In short, MLOps is **all the engineering pieces that come together and often help to deploy, run, and train AI models**. With that, we can see that there are three tightly interwoven components of MLOps:

- Machine Learning
- DevOps (IT)
- Data Engineering

Each component contributes key elements that work to close the ML lifecycle loop within an organization.

With origins in the development of practices used to help data scientists and DevOps teams better communicate using machine learning, MLOps began as simple workflows and processes to deploy during implementations in order to manage the difficulties faced with ML.

Leaps and bounds ahead of where MLOps was just years ago, today MLOps accounts for 25% of GitHub's fastest growing projects. The benefits of dependable deployments and maintenance of ML systems in production are enormous. No longer just simple workflows and processes, now full-on benchmarks and systemization. IT and Data teams in all sorts of industries are trying to figure out how to better implement MLOps.

How MLOps Works

A deeper look into how MLOps works will reveal both the positive side and the problem side of this process. As discussed in an article from Medium:

"MLOps follows a similar pattern to DevOps. The practices that drive a seamless integration between your development cycle and your overall operations process can also transform how your organization handles big data. Just like DevOps shortens production life cycles by creating better products with each iteration, MLOps drives insights you can trust and put into play more quickly."

When considering data as a key business tool that directly relates to how an organization adapts future system operations, essentially MLOps is the process of taking both data and code in order to produce predictions that describe which deployment to put into production. This requires both operations (code) and data engineering (data) teams to work hand in hand.

Benefits of MLOps

Among many positive aspects of ML, a few topline benefits directly relate to any organization's ability to stay relevant and grow in this tech and information-driven world. Most experts agree, as outlined by Geniusee, that the MLOps positive impacts are:

- Rapid innovation through robust machine learning lifecycle management
- Create reproducible workflow and models
- Easy deployment of high precision models in any location
- Effective management of the entire machine learning lifecycle
- Machine learning resource management system and control

From data processing and analysis to resiliency, scalability, tracking, and auditing—when done correctly—MLOps is one of the most valuable practices an organization can have. Releases will end up with more valuable impact to users, the quality will be better, as well as performance over time.

The difficulties with MLOps

As exciting as ML may sound, the fact is, as this technology operations practice comes into play, there are many challenges an organization faces that stem from how to properly combine code and data to achieve predictions. As outlined in Wikipedia, such difficulties are:

- Deployment and automation
- Reproducibility of models and predictions
- Diagnostics
- Governance and regulatory compliance
- Scalability
- Collaboration
- Business uses
- Monitoring and management

With these difficulties in mind, as stated in run.ai, most organizations "never make it from the prototype stage to production. A commonly cited reason for this high failure rate is the difficulty in bridging the gap between the data scientists who build and train the inference models and the IT team that maintains the infrastructure as well as the engineers who develop and deploy production-ready ML applications."

However, with careful consideration and with knowledge of these difficulties, it is possible to reach a smooth MLOps goal with the implementation of standard practices.

Standard practices for MLOps success



Co-team operations

As we see from above, bridging the gap between DevOps and Data is one of the biggest issues to tackle the difficulties of MLOps practices. That's why the best thing an organization can do is create a "hybrid" team.

Towards Data Science explains, "The exact composition, organization, and titles of the team could vary, but the essential part is realizing that a Data Scientist alone cannot achieve the goals of ML Ops. Even if an organization includes all the necessary skills, it won't be successful if they don't work closely together. Another important change is that Data Scientists must be proficient in basic software engineering skills like code modularization, reuse, testing, and versioning; getting a model to work great in a messy notebook is not enough." This type of co-team operations will ensure communication and practice is smooth on all sides.

ML pipelines

The basic structure of data engineering involves pipelines that are essentially extractions, transformations, and loads. Normally formatted in graphs that display each node to represent dependencies and executions, these pipelines are a vital part of data management. With ML, data transformation will always be required. Therefore, pipelines are an essential standard.

Monitoring

For ML, keeping a close eye on operations is far more important than usual production operations. This is because, as Medium details,

> *"ML uses non-intuitive mathematical functions. The black box requires constant monitoring to ensure*

you're operating within regulation and that programs are returning quality information. You may have to retrain data periodically, and determining how and when to do so needs critical collaboration between the teams involved."

Within MLOps, managing and monitoring, both controllable and uncontrollable factors like latency, traffic, and errors, is a top priority.

Versioning

Expanding on a usual DevOps practice, as told in Geniussee, "In a traditional software world you need only versioning code because all behavior is determined by it. In ML things are a little different. In addition to the familiar versioning code, we also need to track model versions, the data used to train it, and some meta-information like training hyperparameters."

Validation

Again, expanding on a DevOps practice, testing, testing, and testing some more is vital to MLOps success. Both models and data require validation. For models, because they are not able to give full results, tests should be statistical and done in relevant segments to reflect data.

When approaching data, tests should be completed in a similar way to code domain testing with higher standards to account for feature changes. Statistical validation on all MLOps fronts is a good course of action.

The future of MLOps

Over the course of the few short years that MLOps has grown in popularity, a number of Open Source frameworks have emerged. A move that signifies that important of this practice, as data and technology continue to expand and reach new heights, developing ML strong strategies now, will assist organizations of all kinds to manage and succeed in the future.

Additional resources

For related reading, explore these resources:

- BMC Machine Learning & Big Data Blog
- Machine Learning: Hype vs Reality
- Machine Learning Architectures Explained
- Top Machine Learning Frameworks To Use
- ML with TensorFlow & Keras, a multi-part tutorial
- 3 Keys to Building Resilient Data Pipelines