Thorogood
Whitepaper

The Business Case for MLOps

A compound of "machine learning" and "operations", many may erroneously assume that MLOps is best left to technical crowds. However, for organizations to bring machine learning into the mainstream of their decision making and derive tangible value from their investments in the space, a fundamental understanding of MLOps, and its acute benefits to the business, is crucial.

Several developments over the past decade have created conditions allowing the application of ML models to take on new heights. Advances in digital technologies and changing consumer behavior have resulted in a proliferation of data, and the advent of cloud computing has made storing and processing that data cheaper and easier than ever before. The availability of open-source libraries and the evolution of analytic tools have made it possible to apply statistical techniques in novel ways and to do so at a scale and speed that was once impossible or at least impractical for many organizations.

With these conditions acting as enablers, many organizations have invested heavily in experimenting with machine learning models and AI in their organizations, but to what avail?

Challenges to Recognizing Value

A recent study between MIT Sloan and Boston Consulting Group, which surveyed organizations across 29 industries globally, found that 71% of respondents felt they understood how AI would change the way their business generates value. But only 11% reported seeing significant financial benefit last year.

Recognizing the value that ML can provide, companies have invested in the development of machine learning models, but many have struggled to scale those models in a way that is enterprise-ready and robust enough to affect change.

Until recently, many companies had built, tested, and deployed ML models in tactical ways, with data scientists using tools and environments they were most familiar with and tackling problems which were often divisional or particular to their domain expertise. While this approach may work for finding pockets of value, to give a model enough credence for enterprise businesses to employ it widely and allow it to guide decision making, it is crucial that companies consider how to evolve their thinking toward supporting ML in a more operationalized ecosystem.

What is MLOps?

To make ML models robust enough for decision making, there is a strategic need for frameworks and architectures that support their ongoing monitoring, maintenance, adoption, and governance. Over the past few years, MLOps principles have begun to emerge to help define the parameters for successfully deploying ML models in production environments. These principles enable greater collaboration between data scientists, data engineers, and business stakeholders on the development, delivery, use and support of ML models within the business.

MLOps practices focus on ensuring models are implemented robustly, efficiently, and in an automated and repeatable fashion. MLOps also focuses on putting in place disciplines for monitoring models over time so that any degradation in performance or accuracy can be addressed before it runs the risk of misguiding the business. Operationalizing ML models in this manner helps make them more scalable, more easily understood, more relevant, and more trusted by the business.

An effective MLOps function can provide feedback to the community of data scientists across an organization – often a dispersed and disparate community – to help them adopt practices and approaches that will help them individually and the organization overall. To do this successfully, we find it invaluable to have overlapping skills sets in the MLOps team, with members proficient in data science, data engineering, and business understanding. Multiskilled teams can be smaller, communicate better, and produce results faster. Because it is particularly important that the productionized outputs match the data scientists' outputs exactly, having sufficient data science capability within the MLOps team ensures not only efficient dialogue but also a second line of review for the model before it's used.

As you can imagine, being able to establish a successful MLOps practice is critical for achieving true business benefit from your machine learning practice.

Business Benefits of MLOps

Uncover value more quickly

In a world of rapid change and volatility, a business's ability to recognize previously imperceivable changes as they are happening provides great advantage. One of the most powerful benefits of MLOps lies in its ability to enable businesses to deploy and adapt ML models faster, which means being able to leverage the insights the model provides with less lead time.

MLOps supports this speed because it is founded on existing DevOps principles, but with a lot more added. DevOps is a set of practices and philosophies created to streamline and accelerate the build, test, and release of software systems to the business faster and more reliably. Beyond DevOps, MLOps must enable the tracking of data sets used to train and re-train models, for example, and to be on the lookout for drift in model accuracy. Integrating Machine Learning concepts with established DevOps principles allows companies to manage their ML lifecycle with greater responsiveness, reliability, and speed.

This means as business needs arise or change, data scientists can work swiftly and collaboratively with data engineers and business leaders to adapt and deploy models that can drive returns for the business.

Increased accuracy, trust, and awareness

A model that is not trusted is not likely to be employed by the business, and a model that is not used by the business will not drive value. A good MLOps framework will ensure models are equipped with enhanced auditability, are properly monitored, and are responsive to change. It will also provide a methodology for combining rapid feedback with automated monitoring to ensure that the quality and performance of the models do not degrade over time.

These components of MLOps help maintain the accuracy of models over time and increase the trust that data scientists and business users place in model outputs. By placing your ML models under enterprise control and giving them the proper governance and care, organizations can foster greater trust in the models and a greater openness to embracing a culture that makes the most of data science and machine learning.

Amplify the value of your data scientists

In a world where strong data science skills are a scarce resource, it is in organizations' best interest to consider how to amplify the value of their data scientists. In many organizations, it is common for the data scientist who developed a model to oversee refreshes, maintenance, and monitoring. As they develop more models, the amount of time spent on support and monitoring can become a significant overhead, limiting the new analyses their experience can be applied to. By automating these processes, data scientists are empowered to spend their time instead on adding new features to existing models and innovating to solve other business challenges.

By laying out frameworks and architectures that support enterprise scale machine learning, organizations can empower their data scientists to perform even more valuable analyses by using tools that support greater scale, faster processing, and wider functionality. For example, when working with a data engineer, data scientists' models with long run times could be reengineered to take advantage of an enterprise architecture that enables greater parallelization options, allowing strands of the model to be run simultaneously and processing times to be dramatically reduced.

Similarly, tools that provide better tracking of experiments can provide data scientists with more robust environments for model building, training, and testing. In addition to tools, alignment on MLOps standards allows for simpler training and spreading of best practices, both among data scientists and MLOps teams. Empowering data scientists to perform their experiments in enterprise environments, and providing them guidance on best practices for scaling, can equip them with more powerful tools and a faster path for deploying models to the business. By bringing the powers of data science and data engineering together, organizations can reap the full benefit of the technological advances that are revolutionizing the data science space.

Manage machine learning consumption costs

Though cloud computing is often cited as a cheaper alternative to traditional on-premise options, analytics solutions in particular can become very resource-heavy if not designed intentionally and with cloud consumption considerations in mind.

For example, a data scientist may create a highly accurate forecasting model for a specific business function and market. In the future, the model may need to be scaled for more data and more markets. MLOps teams can be invaluable in bringing knowledge around options for parallelization given the requirements and refactor the code so that it can run as efficiently as possible as more data and markets are added. The parallelization will lead to increased efficiency in the form of lower end-to-end runtimes, lower compute cluster costs, and in general, more responsible and sustainable consumption of compute resources.

Conclusion

Today machine learning and artificial intelligence are no longer competitive tools used only by those in the vanguard of technological innovation. As organizations of all types bring advanced analytics into the mainstream of their decision making, it's imperative that investments focus not just on model development, but the proper ecosystem that will allow those models to guide decisions responsibly and to best effect. Establishing a strong MLOps practice will require devising architectures, frameworks, governance and learning paths that build and support that ecosystem.

Find out more:
Contact: Amanda.Teschko@thorogood.com
Amanda Teschko – US Chairman, Business Intelligence & Analytics Consultant at Thorogood Associates