



**How to future-proof  
machine learning  
operations during  
major market shifts**

# How to future-proof machine learning operations during major market shifts

A practical guide for data science practitioners and business leaders on responding to and mitigating the impact of disruptions to machine learning models

During periods of severely abnormal data, analytics teams and their dependent business processes need to retune their models to begin generating relevant insights now, while also mitigating the impact of disruptions on future performance.

So how does that work in practice? This perspective is designed to help data science practitioners work with business stakeholders to create process and standards for identifying, reacting to and future proofing against regime shifts that impact their machine learning models.

We are all too familiar with the business disruption caused by COVID-related shifts — changes in consumer behavior, disruptions in production and supply chains, and additional regulatory pressures, to name a few. The impact on machine learning (ML) models is less obvious, but just as significant. After all, ML creates value by leveraging large amounts of data to create faster and more accurate insights, helping data scientists and analytics leaders — and the business processes that rely on their outputs — to model risk, forecast market forces and guide strategic decision making.

Although these ML models are designed to respond to changes, they are also fragile and perform poorly when the input data differs considerably from the data they were initially trained on.

Data science practitioners understand this drastic shift of input data as a regime shift, caused in this case by a significant change in market forces and consumer behavior that are the inputs and targets of ML models. Classic examples of regime shifts in the field of macroeconomics, for example, are breaks in economic cycles that result in profound and unpredictable political and economic consequences, as well as financial ruptures that culminate in recessions. While the mathematics of identifying and responding to regime shift is important, data-driven enterprises first need to understand the strengths and weaknesses of their data operations. Then they can put their mathematical prowess into play to continue to drive business outcomes related to inventory management, asset economics, sales forecasting, fraud detection, risk management, customer behavior and marketing.

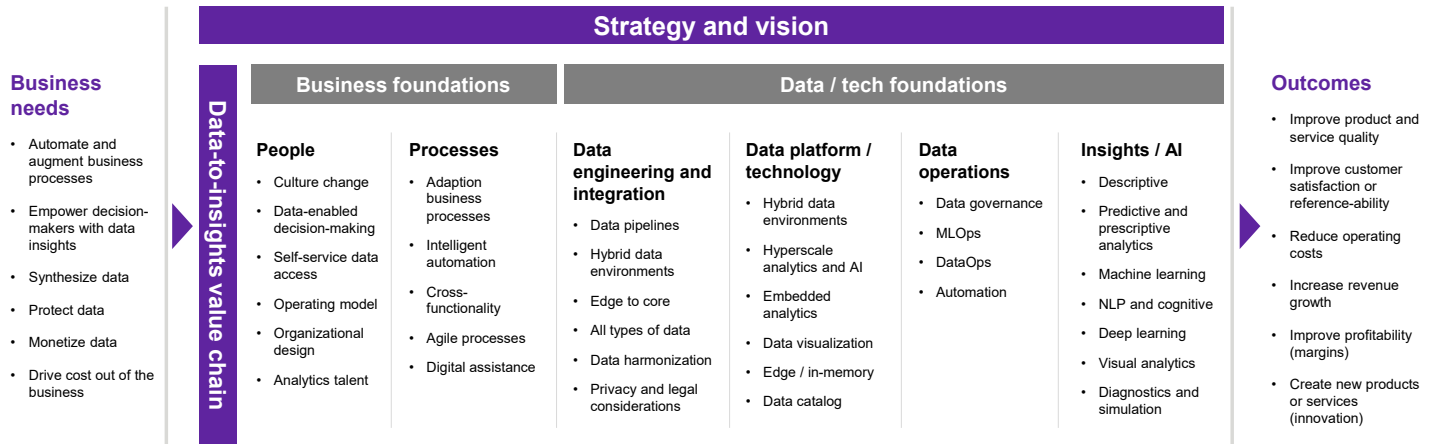
To enable that resiliency, this paper constructs a set of adaptive guidelines to monitor, react and respond to changes that impact the quality of ML models. This is built upon alignment between data science teams and business stakeholders who both design the operational models and data-driven culture needed for value-add analytics operations. This is not just about 2020, but about the inevitable changes that came from it. Wise data and analytics teams understand the certain and ongoing degradation of models and the likelihood and large impact of future massive disruptions.

## Regime shifts:

“...substantial reorganizations of complex systems with prolonged consequences.”

Source: [Complexity Theory for a Sustainable Future, Jon Norberg and Graeme Cumming \(editors\), Colombia University Press, 2008](#)

Successful data and analytics teams are building the data operations to track, react to and counter the impact of perpetual change and monumental regime shift (**Figure 1**).



**Figure 1.** A data value chain map can ensure that data operations are strategically aligned to key business outcomes.

## How to create robust standards for reacting to disruption and degradation

Few, if any, data scientists could have predicted the drastic increase in demand for toilet paper during 2020 and the ripple effect on consumer behavior and retail supply chains. Regardless, as data science practitioners situated within a value chain critical to the success of global companies, we have a responsibility to understand the inevitability of unforeseen disruption to source data, and how to maintain the continuity of insights that guide strategy and inform action.

Two initial steps are critical:

- Understand exactly what will cause ML models to falter
- Establish detailed, documented and agreed-upon plans for identifying and reacting to different levels of failure

In crisis periods, ML models do not outperform conventional statistical models, as they are, by design, incapable of identifying biases in the data. Model risk will usually increase with incorrect assumptions and inputs, making it extremely difficult to gain business intelligence even from vast amounts of data. Standard knowledge among DS practitioners is that both traditional and ML models' forecasting power depends on the quality of the data being fed and the set of assumptions made regarding the underlying data-generating process.

It is not surprising to see model pass/fail thresholds breached more often during a crisis than during normal times. In finance, for example, it is widely documented that the process governing a credit or loan portfolio performance differs across various regimes. A failure to adapt to structural breaks leads to larger errors and poor model performance statistics.

A robust modeling approach requires detailed standards connected to pass/fail thresholds. These thresholds should not be considered a simple KPI or alarm bell

for the model, but should instead trigger documented, detailed and agreed-upon processes for responding to the nuanced indicators of poorly performing models.

Here are three prompts to help you establish standards for identifying and communicating model failure or performance degradation:

**Document the factors that influence the model's performance in formats that facilitate strategic discussion.** During the construction and operation of a model, assess the different factors' varying degrees of influence on the model's predictive power. Create a framework for describing the impact of each factor on the model's performance. Use the framework to proactively assess the mathematical implications, but also the business implications of regime shift. Generate documentation and referenceable assets that data science practitioners and business stakeholders can use to discuss changes to input data and the resulting impact on model performance.

**Create a communication plan for calling out the factors leading to larger errors and deterioration of the model's predictive power.** What is the difference between minor disruptions and a complete breakdown? How significant is a 2-week disruption compared to a 2-month disruption? Extend your framework for describing changes to the model, to establish triggers, escalation standards and contingency plans for consistency in your reactions over time. Stakeholders should not be left to interpret the severity of disruption in an ad hoc manner, but instead should be able to start from an agreed-upon reference and designation.

**Decide what managerial judgment should be factored in to determine when and how models are recalibrated.** Identify who is responsible for and dependent on the models, and who has the final say for the various strategic and tactical decisions that must be made. Document their roles in spotting, assessing and reporting disruption as well as in researching, designing and committing to remediation. Something as simple as a responsibility assignment matrix for reacting to different thresholds of disruption can be a worthwhile investment.

## If it's broken, fix it

Standards for identifying a disruption are not the same as appropriately reacting to it. So how do you know when to make a change to the model?

Making changes to machine learning models is not simple. The pandemic has revealed how entangled our lives are with AI. There is a subtle co-dependence between behavioral changes and how AI models work, so a tremendous set of variables about the data and related business drivers need to be factored in.

Here are three key tips for determining when and how to make changes to your models:

**Don't be afraid to take action just because it is supposed to be "automated."** It is important to keep a cautious eye on automated systems that are just barely holding up. When you are faced with extraordinary circumstances, you cannot set models on auto pilot and be fully hands off.

Stepping in with a manual adjustment in automated systems when needed is paramount.

**Make sure everyone knows what they are taking on when adapting to regime shift.** The spectrum of approaches available for data science practitioners ranges from simple to sophisticated. New mathematical and computer models will surface as potential solutions. They should be assessed not only for their feasibility to improve the model performance, but also for the extent of change management and implementation costs that will be required to adopt them. Take the time to ensure you are up to date on industry best practices and the details of real use cases for new solutions.

## Hurdles faced by AI at scale:

“The focus of data science teams has traditionally been on developing analytical assets, while dealing with the operationalization of these assets has been an afterthought.” — Gartner

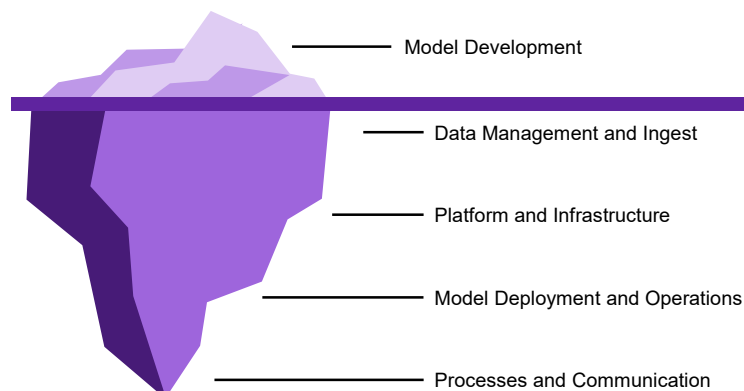
Source: [Use Gartner’s 3-Stage MLOps Framework to Successfully Operationalize Machine Learning Projects](#), 2 July 2020

**Use design best practices and be transparent about uncertainty.** At the center of your focus is data science and machine learning, but don’t forget about the people. Understand your changes as more than just mathematics, and consider the needs from a variety of stakeholders. Utilize design thinking best practices to ensure you have defined the problem space appropriately. When meeting, clearly distinguish between generative discussions about possibilities and directional alignment, compared to focused efforts to scope designs, refine solutions and move forward down a chosen path. Leverage input from engineers, P&L owners, users of analytics reports and a breadth of stakeholders. This will allow you to design the best solution to the disruption, while also maintaining a high level of transparency around the problem and the path to resolution.

Realistically, how are design thinking and data analytics capabilities integrated? Learn more about [data and analytics solutions](#).

## Future proof with resilient models and data operations

The more time data scientists spend triaging and reacting to model degradation, the less time they spend building new and innovative ways to drive business outcomes with data. It is possible to create models that hold up even when the environment changes — provided the ranges for the features remain broadly identical to what have been tested previously (**Figure 2**).



**Figure 2.** Sustainable value from ML comes from more than the models themselves. Ensure you have visibility into connections beneath the surface of your data to insights value chain.



Consider the following ways to make current and future models more resilient to disruption:

**Unwind the reliance on historical data if the model identifies regime shifts.**

Since machine learning models are trained on data, if you don't believe that the environment from which the data was obtained is comparable with the current setting, you should try simpler models that rely on fewer features or that deploy variables that explain the final output or result of the models. Reducing the number of features will prevent you from being misled by a questionable dataset and help you understand what will still work and what might not in the new environment.

**Establish requirements and best practices for how to handle missing or severely disrupted data.** Data science practitioners should take proper care in the way they handle missing data. When manufacturing, supply chain, industrial and other business processes halt during forced business shutdowns, missing data can adversely impact the future training of ML models.

To alleviate future issues related to missing data, keep a log of events and how they are represented in the data during this time. In order to minimize time spent on future data cleansing, document these data issues, whenever possible, being sure to include the reasons for the exception and the potential consequences expected.

It is possible to overcome issues in data irregularities and limit the alteration to the original time series either by omitting the data or inputting it manually. For instance, consider applying shorter time frames to your model retraining frequency. This is particularly effective for industries — financial markets, for example — where issues arising from data irregularities are short-lived. Don't be too concerned about model performance in this scenario, as the remedy is only temporary and won't affect overall performance.

Data from a previous crisis (usually economic) may appear to be the most likely candidate to inform management or the model developers about the overlays. However, be warned that the data-generating mechanism can significantly vary across crisis periods, hence relying on previous crisis data could be a poor benchmark for the future.

**Retrain more frequently and with smaller data sets if statistically accommodable.** When faced with a temporal component in the data, models in multiple business applications should be re-trained frequently with smaller data sets if it is statistically accommodable and economically sensible. The purpose is to become more agile and better capture the quickly changing nature of the real-world socioeconomic dimension they model.

## Caution!

Inefficient ML and data operations can be costly, from cloud computing costs to engineering productivity. Apply DataOps best practices to increase visibility and strategic control.

In addition, the following best practices help create ML models that are less susceptible to catastrophic failure:

**Integrate more data sources where possible, increase frequency and granularity of data collection.** Increasing the frequency and granularity of data collection from weekly to daily, or daily to hourly, for example, can help shorten learning cycles and improve understanding of changes in data segments.

In the case of financial markets forecasting, for instance, models should be trained not just on the ups and downs of the last few years, but also on black swan events such as the Black Monday stock market crash in 1987, the dot-com bubble burst in 2000 and the 2007-2008 global financial crisis. And of course, a health crisis such as the coronavirus is a perfect trigger to build better ML models.

**Create a log of events and how they are represented in the data, and document reasons and expected consequences.** It is important for data scientists to resist the temptation to throw 2020 data out of their forecasting models entirely due to anomalies and instead rely on 2019 and 2018 data.

It is true that governments and regulators are implementing blunt instruments to kickstart the economy, creating anomalies in consumer behavior. However, these anomalies should not be ignored, and you should train your models to guide you through them, incorporating dummy variables to capture and account for the anomaly period if necessary.

Although it is unlikely that the economy will face such a severe global impact anytime soon, we more frequently do encounter regime shifts in regional context. For instance, shutdowns due to severe weather or natural disasters can have a significant localized impact and are far more common. It is important, therefore, to train models through anomalies (or abnormal circumstances) to understand the impact, duration and recovery rate of the business, so that you can make smarter decisions at scale and in less time.

**Acknowledge the existence of bias.** Biases can arise from ML models when they are tasked with implementing policy decisions as a result of unrepresentative training data sets. Although it is not the subject matter of this paper to elaborate further, it is important to explore the intended and unintended consequences of bias on the ML algorithms.

For example, when the input data used to train the models is more representative of particular groups of people based on gender, race and other factors, the forecasts from the model may also be systematically biased towards these groups. Hence, these biases can lead to decisions that can adversely affect specific groups of people without the developer explicitly modeling them.

**Create a more agile practice — MLOps.** To quote Sebastian Klöser, solution principal, AI and Data Science at DXC Luxoft, “Most companies have focused on building machine learning muscle — hiring data scientists to create and apply algorithms capable of extracting insights from data. This makes sense, but it’s

## How MLOps helps

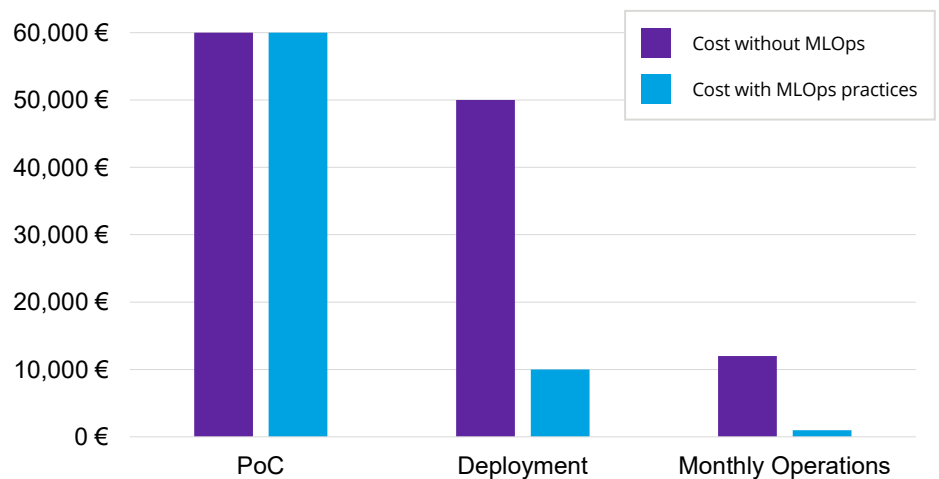
MLOps reduces costs from POC to operation at scale compared to native and ad hoc machine learning initiatives (**Figure 3**), regardless of your...

- state of maturity
- legacy systems
- preferred tools of choice
- current organizational model

a rather limited approach. Think of it this way: They've built up the spectacular biceps but haven't paid as much attention to the underlying connective tissues that support the muscle."

MLOps is that connective tissue. It integrates machine learning with software development best practices to ensure businesses get real and continuous value from machine learning and artificial intelligence applications. This by no means obscures the ability of data scientists to flex their muscles. As [Klöser observes](#), "Data scientists can continue to use the tools and methods they prefer, their output accommodated by loosely coupled DevOps and DataOps interfaces. Their ML algorithm development work becomes the centerpiece of a highly professional factory system, so to speak."

### Money spent across the ML lifecycle



**Save 75% per AI Use Case**

**Figure 3.** Savings realized from DXC MLOps, assuming average runtime of one year.



“Alternative modeling approaches can act as a dependable mechanism or interim solution, or even offer an opportunity to harness new trends and patterns emerging from the disruption.”

## How are others doing it?

Clearly, a complexity of factors influences machine learning models and their integration into an enterprise context. As a best practice, it is always worth staying up to date on how changes are being handled outside of your organization.

When examining partners, platforms, solutions and techniques, it is important to consider not just the mathematical validity of their approach, but also their relationship to problem solving. Were they facing the same problem you are? Did they have the same resources and intended outcome? At the core, we are trying to solve problems through machine learning, not simply produce beautiful math.

One traditional econometric approach used for regime shift is the Markov-switching model. In this regime-switching model, shifts between regimes evolve according to an unobserved Markov chain, a mathematical system that transitions from one state to another according to certain probabilistic rules.

A regime-switching model is a parametric model of a time series in which parameters are allowed to take on different values in each of some fixed number of regimes. This method is widely used by econometricians and quantitative analysts in financial market modeling, and is widely applied to time series data.

Despite attempts to recalibrate models with traditional approaches, it is not abnormal for previously well-established models to struggle as new data regimes emerge. Alternative modeling approaches can act as a dependable mechanism or interim solution, or even offer an opportunity to harness new trends and patterns emerging from the disruption.

Augmenting supervised machine learning techniques with alternative techniques such as reinforcement learning can help businesses rapidly adapt to new economic behavior. This machine learning technique enables an agent to learn in an interactive environment by trial and error, using feedback from its own actions and experiences. Reinforcement learning uses rewards and punishment as signals for positive and negative behavior.

Another modeling approach well suited to major regime shifts is Bayesian modeling. This technique has intrinsic abilities to measure uncertainty, include prior knowledge and embed structural relationships between variables. Bayesian networks also support causally grounded counterfactual analysis to infer what happens to a target (dependent) variable when an intervention is performed on a feature (predictor) variable.

Other strong candidates to successfully model the relatively short data sets obtained are hierarchical linear models (HLM), Gaussian processes and other probabilistic graphical models (PGM).

“DXC’s Data and Analytics team can help you successfully establish model governance at scale, enhance collaboration between data and business teams, and deploy the foundation for resilient MLOps.”

## Conclusion: We are in it together

If it is any consolation, you can be sure that your peers across private and public organizations are dealing with the same disruptions you are. As processes are “digitalized” and data is “democratized,” we are having less trouble accessing data, and more trouble ensuring it is clean and then figuring out exactly what to do with all of it.

If you have decided to invest in the foundations of a robust yet adaptive data science practice, you are already ahead of the curve. But beware of one pitfall: Most of the challenges businesses encounter with their ML models are because an increasing number of them are buying off-the-shelf ML systems and lack the in-house know-how needed to maintain and retrain them. Only a handful of businesses have the skillset and data infrastructure to put sophisticated alternative approaches into practice.

For businesses that assumed all automated ML models could run by themselves, the current situation is a wake-up call.

Now is a time to take stock at those newly exposed systems and ask how they might be designed better and made more resilient. If machines are to be trusted, we need to closely monitor them. A team of experienced data scientists can help you connect what’s going on in the world to what’s going on with the algorithms, and introduce operational best practices and a well-considered DataOps plan. This will enable ML to best serve the business now and into the predictably unpredictable future.

## About DXC in data analytics

DXC Technology has a cross-functional team of consultants, data scientists, designers, architects and engineers who assist customers in implementing advanced enterprise data solutions. Our Data and Analytics team can help you successfully navigate disruption by establishing model governance at scale, enhancing collaboration between data and business teams, and deploying the foundation for resilient MLOps. Learn more at [dxc.com/data-analytics](https://dxc.com/data-analytics).

## About the author

**Rajiv Thillainathan** is a senior data scientist for Analytics and Engineering at DXC Technology in the Asia-Pacific region. He has extensive experience in quantamental research and investing, having worked for global investment banks and asset management, hedge fund and technology firms. He has designed and delivered leading analytics solutions by discovering exploitable patterns in the data to realize profitable outcomes, and helped enhance performance by implementing systematic strategies based on advanced statistical analysis, ML, data mining and data visualisation techniques across structured and unstructured data.

Get the insights  
that matter.

[dxc.com/optin](https://dxc.com/optin)



Learn more at  
**DXC Data and Analytics services**