

WHITE PAPER

Opening the “Black Box”

THE PATH TO DEPLOYMENT OF AI MODELS IN BANKING

Opening the “Black Box”

THE PATH TO DEPLOYMENT OF AI MODELS IN BANKING

CONTRIBUTING AUTHORS:



Ram Ananth

Avantage Reply, Head of Quantitative Practice



Seph Mard

DataRobot, Head of Model Risk Management



Peter Simon

DataRobot, Data Scientist, Financial Markets Practice

Emerging Trends for Model Validation in the Age of AI

Recent technological advancements have accelerated the integration of AI and machine learning models into more and more banking processes. In today’s banking industry, institutions not using AI and machine learning risk losing their competitive edge, as competitors are increasingly enhancing their strategic decisions with the powerful analytical capabilities of AI and machine learning.

However, due to an increased reliance on models for everyday business processes and decisions, model risk must be effectively managed. If left unchecked, the consequences of model risk can be severe; where model risk is defined as the risk of financial or reputation loss due to errors in the development, implementation or use of models. Therefore, AI and machine learning models require constant monitoring and effective validation. This is not only a regulatory requirement, but it is also sound business practice.

This white paper presents the cornerstones of effective modern model risk management in the age of AI and machine learning by first providing an overview of AI and machine learning in banking, summarizing the regulatory background and the machine learning model lifecycle, and then finally presenting the challenges and emerging best practice for the validation of models, in an ever-changing world of AI and machine learning.

The Rise of AI and Machine Learning

OVERVIEW OF MACHINE LEARNING

Most banks today are quite familiar with machine learning, a subset of AI in which algorithms learn by example from historical data and are used to predict outcomes and uncover patterns that are not easily spotted by humans. However, not all classes of machine learning algorithms are recent developments. In fact, many have been used in the banking and financial services industry for decades. These include basic statistical methods and techniques like linear and logistic regression (collectively, ‘generalized linear models’ or GLMs) and rule-based modeling. These traditional methods are relatively simple, well-understood by practitioners and easily explainable and interpretable. However, this comes at the price of reduced performance compared to that achievable with modern machine learning techniques. Table 1 presents a summary of fundamental machine learning concepts and methods.

METHOD	DESCRIPTION
Supervised Machine Learning	<p>Supervised machine learning algorithms uncover insights, patterns, and relationships from a labeled training dataset – that is, a dataset that already contains a known value for the target variable for each record. Because you provide the machine learning algorithm with the correct answers for a problem during training, it is able to “learn” how the rest of the features relate to the target, enabling you to uncover insights and make predictions about future outcomes based on historical data.</p> <p>Two common examples of supervised machine learning are regression and classification:</p> <ol style="list-style-type: none"> 1. Regression, in which the algorithm returns a numerical target for each example, such as how much revenue will be generated from a new marketing campaign. Linear Regression is a common method used to model the relationship between a continuous target variable (i.e., dependent variable) and other features from the data (i.e., explanatory variables) by fitting a linear equation to observed data. Such models are called linear models. 2. Classification, in which the algorithm attempts to label each example by choosing between two or more different classes. Choosing between two classes is called binary classification, such as determining whether or not someone will default on a loan. Choosing between more than two classes is referred to as multiclass classification. Logistic regression is a common algorithm for binary classification.
Unsupervised Machine Learning	<p>Unsupervised machine learning algorithms infer patterns from a dataset without reference to known, or labeled, outcomes. Unlike supervised machine learning, unsupervised machine learning methods cannot be directly applied to a regression or a classification problem because the values for the output data are unknown, making it impossible to implement a supervised learning approach. Unsupervised learning can instead be used for discovering the underlying structure of the data.</p> <p>Some applications of unsupervised machine learning techniques include:</p> <ol style="list-style-type: none"> 1. Clustering allows you to automatically split the dataset into groups according to similarity. Often, however, cluster analysis overestimates the similarity between groups and doesn’t treat data points as individuals. For this reason, cluster analysis is a poor choice for applications like customer segmentation and targeting. 2. Anomaly detection can automatically discover unusual data points in your dataset. This is useful in pinpointing fraudulent transactions, discovering faulty pieces of hardware, or identifying an outlier caused by a human error during data entry. 3. Association mining identifies sets of items that frequently occur together in your dataset. Retailers often use it for basket analysis, because it allows analysts to discover goods often purchased at the same time and develop more effective marketing and merchandising strategies. 4. Latent variable models are commonly used for data preprocessing, such as reducing the number of features in a dataset (dimensionality reduction) or decomposing the dataset into multiple components.

Table 1: Overview of Traditional Machine Learning

As computer processing power has increased exponentially in recent years, computationally intensive modern machine learning approaches such as Artificial Neural Networks, Deep Learning, Decision Tree Algorithms, Regularization Algorithms, Ensembling, and so on have gained increasing traction. ('Modern' in this context is a relative term; for instance, neural networks date back to the 1950s and Support Vector Machines to the 1960s, but only in the last two or three decades has sufficient computing power become available at a low enough cost to deliver on the promise of these techniques.) These approaches can deliver much-improved performance compared to traditional statistical techniques but may be viewed as less explainable or interpretable than more traditional statistical modeling. This is due in part to the complexity of these methods that require more sophisticated non-parametric estimation techniques and hyper-parameter optimization, as presented by Table 2 below. However, as we will review in detail shortly, recent advances in modern interpretability techniques have reduced this problem substantially by providing an outcomes-focused approach to interpretability, which make it possible to gain insights into what is driving models, individual predictions, and the sensitivities therein.

METHOD	DESCRIPTION
Tree Based Algorithms	<p>Tree based models in their simplest form are decision trees, which fit a sequence of decisions to data. As these often lead to naive models which perform poorly once deployed, advanced tree-based algorithms combine multiple such decision trees to produce models which are robust and highly performant.</p> <p>Some common example of advanced tree-based algorithms include: Decision Trees (DT), Random Forests (RF), Classification And Regression Trees (CART), and Gradient-Boosting Machines (GBM).</p>
Artificial Neural Networks	<p>Artificial Neural Network algorithms are a collection of models which are adept at capturing non-linear patterns, or patterns that are allowed to reuse variables. Modern neural networks are building blocks that allow model builders to design models that exactly represent the problem they wish to solve. Neural network libraries provide tools (such as auto-differentiation) to speed up the process of fitting that model to data.</p> <p>Some common examples of Artificial Neural Networks algorithms include: Perceptron, Feed-forward Backpropagation Networks, Hopfield Networks, and Radial Basis Function Networks (RBFN).</p>
Deep Learning	<p>Deep Learning algorithms run data through several "layers" of neural network algorithms, each of which passes a simplified representation of the data to the next layer. The ability to process large numbers of features makes Deep Learning algorithms very powerful when dealing with unstructured data (such as images, audio and video). However, Deep Learning algorithms can be overkill for less complex problems because they require access to a vast amount of data to be effective and may be outperformed by Tree-Based Algorithms in these cases.</p> <p>Instead of organizing data to run through predefined equations, Deep Learning algorithms characterize basic parameters about the data and trains the computer to learn on its own by recognizing patterns using many layers of processing.</p> <p>Some common examples of Deep Learning algorithms include: Recurrent Neural Networks (RNN), Deep Boltzmann Machines (DBM), Deep Belief Networks (DBN), Convolutional Neural Networks (CNN), and Stacked Auto-Encoders.</p>
Regularization Algorithms	<p>Regularization is a common modification of other supervised machine learning methods (typically regression) that penalizes models based on their complexity, favoring simpler models that are also better at generalizing unseen data that was not used for training.</p> <p>Some common examples of Regularization algorithms include: L1, L2, Ridge Regression, Least Absolute Shrinkage and Selection Operator (LASSO), Elastic Net, Least-Angle Regression (LARS).</p>
Ensembling	<p>In contrast to less sophisticated machine learning, ensemble methods combine multiple models (weak learners) and construct a set of hypotheses to solve the same problem (strong learners).</p> <p>Some common examples of ensembling techniques include: Boosting, Bootstrapped Aggregation (Bagging), and Stacked Generalization (blending).</p>

Table 2: Advanced Machine Learning Methods

Model Risk Management: Regulatory Requirements

Model risk has seen heightened scrutiny from regulators worldwide as model complexity has increased and models are being used in more and more areas of banks. This scrutiny was kick-started by the aftermath of the global financial crisis, in which it became apparent that at least some incumbent models at the time were unfit for purpose, which was exacerbated by a lack of appropriate governance around models and their use. The inevitable supervisory reaction has been increased regulatory scrutiny for how models are built, approved and maintained. In 2011, the Federal Reserve and the Office of the Comptroller of the Currency (OCC) published the SR 11-7 supervisory guidance on Model Risk Management. As part of the Comprehensive Capital Analysis and Review (CCAR), banks are required to submit documentation on model risk management policies and practices. SR-11-7 has emerged as the de facto regulatory standard for model risk management.

Regulators elsewhere have also increased scrutiny. In 2017, the European Central Bank (ECB) published guidelines known as the Targeted Review of Internal Models (TRIM). TRIM expects financial institutions to have a well-developed model risk management framework in place, and to apply this to all models used for business decision-making; especially those pertaining to capital requirements. Also in 2017, the Office of the Superintendent of Financial Institutions (OSFI) in Canada published the E-23 guidelines on model risk management for deposit-taking institutions covering expectations on practices for an enterprise-wide model risk management framework. In the United Kingdom (UK), the Prudential Regulation Authority (PRA) is also heavily focused on model risk; with stress testing models getting a lot of attention. In 2018, the PRA published a supervisory statement on model risk management for stress testing (SS3/18), covering firms' development, implementation and use of stress testing models.

Algorithmic trading has also seen increased regulatory focus in the UK. The PRA published a supervisory statement on Algorithmic Trading (SS5/18) outlining expectations on the risk management and governance of Algorithmic trading, where machine learning models are increasingly being used to facilitate automated trading decisions. These developments are mirrored in many other jurisdictions around the world.

The global regulatory requirements for model risk share common pillars:

- Financial institutions must implement an effective model governance framework underpinned by robust policies, procedures and controls to manage model risk
- Financial institutions must evaluate model risk through the full lifecycle covering the development, implementation and use of models
- Financial institutions must establish a clear definition of what constitute models and ensure that the model inventories are effectively maintained
- Financial institutions must provide effective challenge of models through appropriate independent model validation activities that evaluate the conceptual soundness, performance and weaknesses of models

Model Lifecycle - The Path to Deployment in Banking

The path to model deployment in banking is traditionally cumbersome. Deploying a model is hard and takes time, resources and coordination across many teams throughout the bank. Gains in efficiency of this process, whether technological or process-driven will certainly accelerate the speed-to-deployment of new models into production. However, the challenge remains in balancing speed to implementation with a set of appropriate controls, to ensure that model errors are detected quickly and can be corrected in a timely manner, without having a detrimental impact on the bank’s business.

The model lifecycle is a robust workflow that describes the different stages of production of a model, such as development, implementation, validation, approval and use (see Figure 1 below and Table 3 at right).

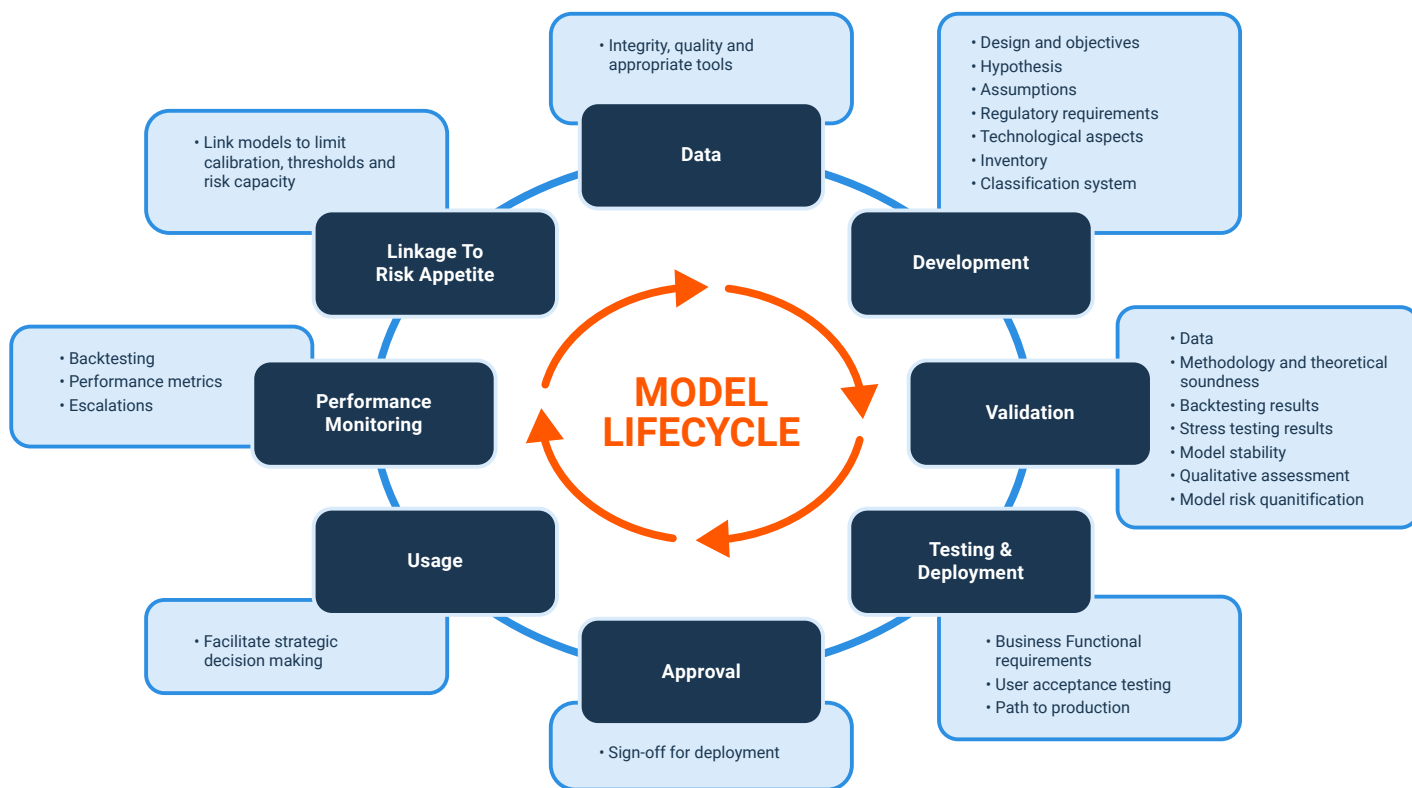


Figure 1: The Model Lifecycle

METHOD	DESCRIPTION
Origination, Design and Development	The model owner needs a clear understanding of the purpose of the model – why the business needs it. Clearly-documented business requirements enable the model builder to construct a model that is aligned with business needs. Furthermore, one needs to ensure that the data used for the model is complete, timely and accurate.
Implementation	<p>Implementation requires the infrastructure to develop, test and run the model. The model owner, together with IT, ensures that platforms and tools are in place to support the model throughout its entire life. Technical and functional specification documents specify the implementation of the model.</p> <p>Testing should expose model strength, limitations and constraints beyond which model performance deteriorates significantly. If models are provided by a third party, their documentation needs to be reviewed and results examined to ensure the model works as intended and that the model owner has full understanding of its limitations.</p>
Validation	<p>Independent validation provides assurance that the model performs as expected, meets business needs and satisfies regulatory requirements. The scope of review during validation includes:</p> <ul style="list-style-type: none"> • Assessment of Model Purpose; • Review of model design, assumptions, and development; • Review of model performance and stability during training, and review of the ongoing model performance plan; and, • End-to-end review of the model documentation, i.e., can the model be reconstructed independently using the provided model documentation? <p>For validation purposes, it is important to use challenger models to test the key assumptions and compare different approaches in terms of predictive power and speed. Having the ability to construct benchmark models quickly is very important to be able to provide the effective challenge. It is crucial that there is no conflict of interest when carrying out the independent validation.</p>
Approval	<p>All models are reviewed by relevant committees before being approved, and the board of directors informed. The committees may also advise and refer the final decision to the board of directors. The materiality of the models is a significant factor when determining the approval process.</p> <p>The decision uses the reports from the model owner, which include adequate information to support the review. Reports provide concise remediation actions and plans in response to validation reports and issues identified. Similarly, the committees receive and review all validation reports.</p>
Use	Only when the model has been approved can it be deployed in a production environment. The model owner documents the scope of and concise restrictions on model use. Controls to ensure intended model use can also be put in place.
Ongoing Monitoring	Ongoing monitoring of all models ensures that the model performance is within agreed parameters and enables identifications of model limitations and weaknesses. It is an important source of insight into the operational implementation of a model, e.g., data issues, methodology complexity and technological considerations.

Table 3: Key Components of the Model Lifecycle

Organizational Considerations and Effective Challenge

The “effective challenge” of models is a guiding principle of model risk management: critical analysis by objective, informed parties who identify model limitations and assumptions and advocate appropriate revisions. Effective challenge depends on a combination of incentives, competence and influence. Incentives to provide effective challenge to models are stronger when there is greater separation of that challenge from the model development process.

To be effective, challenge must have the influence to ensure that actions are taken to address model issues, through a combination of explicit authority, stature within the organization and commitment and support from higher levels of management. Thus, competence is a key to effectiveness, since technical knowledge and modeling skills are necessary to conduct appropriate analysis and critique. Having clear demarcation in terms of these three lines of defense, with clearly defined roles and responsibilities, is very important in this regard.

As models have broad applications across many lines-of-business, it is important that the committee structure is reflective of the diverse expertise required to appropriately review and challenge the framework, as well as to advise the board. Factors that can influence the committee structure include financial products involved in the model, regional regulatory disparities, type of models which can span different time scales (algorithmic trading models vs. credit scoring models), operational complexity and risk classes. Best practice roles and responsibilities are shown in Table 4 at right.

FUNCTIONAL ROLES

Responsible for executing policies and procedures prescribed in the model risk management framework.

ROLE	RESPONSIBILITIES
1 FIRST LINE OF DEFENSE: Manage risk in model specification and building	
Model Owners/ Users	<ul style="list-style-type: none"> • Set out model requirements and business context • End user of the final model output
Model Builders	<ul style="list-style-type: none"> • Ensure the model complies with business requirements and framework requirements • Responsible for the end-to-end model lifecycle: design, build, use and performance monitoring • Produce and maintain documentation • Implement controls and respond to breaches identified by model validation
2 SECOND LINE OF DEFENSE: Control and model validation	
Independent validators	<ul style="list-style-type: none"> • Perform independent validation of documentation provided by model owner, model assumptions and implementation, model use and limitations on use, model inter-dependencies; can the model be reproduced by a third party using existing documentation? • Monitor remediation of breaches performed by the model owner and corresponding documentation • Monitor production results against expected results • Ensure that any external independent validation required has the expertise and skill to perform the validation
Control functions	<ul style="list-style-type: none"> • Own the overall Model Risk Management Framework and Institutional Policy • Ensure adequate first line of defense controls are in place • Manage conflict of interest, e.g., between model owner and model validator • Approve use of independent external validation
3 THIRD LINE OF DEFENSE: Manage the risk managers	
Internal Audit	<ul style="list-style-type: none"> • Assesses the performance of the model risk management framework: governance, risk management and internal controls • Evaluates existing policies and procedures to determine whether they are adequate and comply with regulatory requirements and with industry best practices

Table 4: Three Lines of Defense Against Model Risk

Reducing Model Risk with Automated Machine Learning

Individual, hand-coded models are not only difficult and time-consuming to build, they are also error-prone and introduce implementation risk to the modeling process. Traditional machine learning model development methods are time-consuming, tedious, and subject to user error and bias, which introduces additional model risk. Automation also means systemization. By systematizing technical aspects to follow a standardized process — for example, out-of-sample validation, model tuning, variable selection, model selection, and so on — implementation risk is greatly reduced. Automated machine learning software handles these tasks behind the scenes, with a single administrative interface, allowing data scientists and modelers to handle more projects and perform more high-value work — like engaging with their internal stakeholders to understand the business problem, to explain business results, and collaborating with model validation to ensure streamlined productionalization of models.

As a result, using automated machine learning, organizations can more consistently adhere to best practices, streamline model validation, and, ultimately, dramatically accelerate the speed-to-deployment. This systemization provided by automation allows for a high confidence of replicability of the model development process, which provides model validators with a more efficient solution to provide effective challenge and independently assess and review the model and its potential limitations, ultimately enabling the validation process to effectively manage model risk.

Figure 2 at right presents how a traditional machine learning process is implemented manually and how automated machine learning creates massive efficiency by automating these steps and ultimately ensuring best practices are implemented, while also providing the necessary transparency needed for validation.



Figure 2: Machine Learning Workflow

It's clear that automation reduces operational and model risk through programmatic consistency for feature engineering and extraction, model building and hyperparameter optimization, model diversification through challenger models, model selection, automatic documentation, model performance testing, and ongoing monitoring and model management. We will dig deeper into exactly how automated machine learning can help drive successful model risk management in the age of AI and machine learning in the next section, as well as the best practices for validating AI and complex machine learning models.

Successful Independent Model Validation in the Age of Artificial Intelligence & Machine Learning

The successful use of AI and machine learning models requires that validation teams are able to provide effective challenge prior to their use. However, these more advanced methods are commonly misconceived to be "black-box" algorithms; these models are sometimes thought to give up transparency, interpretability, and explainability which are needed to validate these models, with the benefit of added modeling accuracy. However, a solution to the "black box" problem is an outcomes-based approach to interpretability.

DataRobot is a leader in AI explainability and provides numerous model agnostic interpretability features that provide optimal transparency into sophisticated machine learning models. The technology supporting explainable and interpretable AI and machine learning has made substantial progress in recent years, and DataRobot is allowing users access to highly accurate machine learning models without a tradeoff of transparency or interpretability, which is a requisite for these models to be effectively validated.

With linear regression and other generalized linear models, we can gain considerable insight into the structure and interpretation of a model by examining its coefficients. For example, coefficients from a linear model explain both magnitude and directional impact of that feature, holding all else constant. For more complex modern models like regularization algorithms, tree based approaches, artificial neural networks, deep learning, and ensembling methods considered earlier, no comparably simple parametric description is available, making the interpretation of these models substantially more difficult. Modern machine learning interpretability approaches deal with this problem by focusing on models' outcomes; by analyzing and understanding how models react to different inputs, we can build a good picture of not only model performance and stability, but also the drivers of a model's behavior and sensitivities to various scenarios, including extremes.

MODEL DIVERSIFICATION

While some machine learning techniques tend to consistently outperform others, it is rarely possible to say in advance which will perform best for a given business problem. Therefore, during the modeling and validation processes, it is best practice to build and test independent challenger models. DataRobot provides an out-of-the-box "model leaderboard," ranking models based on their out-of-sample predictive performance in order to allow model selection across a high number of modeling techniques.

DataRobot automatically builds models that are highly interpretable (e.g, simple linear models) and challenges them with more sophisticated modelling approaches (e.g, ensembling methods, neural networks, decision-tree boosting, etc.) or pre-processing. Therefore, the leaderboard includes a diverse pool of modeling methods that and an infrastructure that allows for easy benchmark comparison between models. DataRobot supports this activity by ranking a wide variety of modeling approaches on unseen data using common evaluation metrics which are appropriate to the class of modeling problem in hand, as well as making a number of other evaluation metrics available during model selection, as presented below in Figure 3.

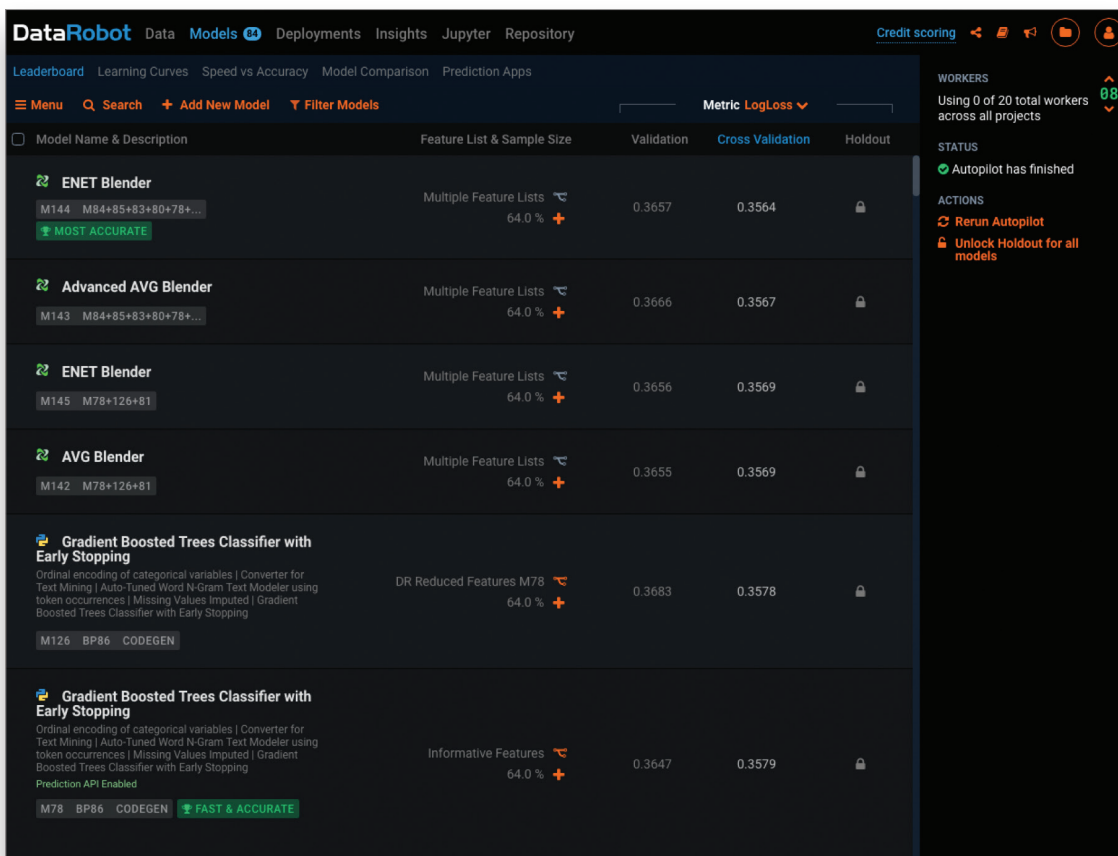


Figure 3: Model Diversification – DataRobot Leaderboard

A solution to the “black box” problem is an outcomes-based approach to interpretability.

MODEL STABILITY AND VALIDATION

To find patterns in a dataset from which it can make predictions, an algorithm must first learn from historical examples – typically from a historical dataset that contains the output variable to be predicted. However, if a model is trained too closely on its training data then it may be overfitted; where overfitting is a modeling error that occurs when a model is too closely fit to training data and therefore performs poorly on out-of-sample data (previously seen data not used to train the model). Overfitting generally results in an overly complex model that explains idiosyncrasies and random noise in the training data, rather than the underlying trends that the model was intended to capture. There is substantial risk to a bank’s bottom line if overfitted models are allowed to be deployed, as such models will not be stable in the light of production data and may produce outcomes which result in significant financial losses; this is a particular risk in financial markets, where models are particularly prone to overfitting due to markets’ ever-changing nature.

To avoid overfitting, the best practice is to evaluate model performance on out-of-sample data (or out-of-time for date/time partitioned data). If the model performs very well on in-sample data (the training data) but poorly on out-of-sample data, then that can indicate a model that is overfit. Model overfitting can be avoided by training (and tuning) a model on one dataset and then validating the model’s predictions on unseen “training” data, with the final model being further evaluated for stability on another set of unseen data (the “holdout” set). Depending on the modeling algorithm, business problem and the size and stability of the underlying data, this may take the form of cross-validation, training-validation-holdout, out-of-time validation and/or backtesting. A discrepancy in model performance across model training, validation and holdout sets is indicative of overfitting, and therefore should be closely scrutinized during validation, possibly considering the deployment of other more stable models, even if their headline performance is slightly worse. DataRobot fully supports the different validation techniques, ranking models on their out-of-sample performance.

MODEL PERFORMANCE

A lift chart sorts and groups model’s predictions into equal sized bins and compares them to the actual outcomes in the model validation data. This depicts how well a model segments the target population and how capable it is of predicting the target by visualizing modeling accuracy for each bin. The chart is sorted by predicted values—lowest to highest predictions, for example—which highlights whether there are systematic differences in model performance for certain groups of predictions. Looking at the lift chart, the left part of the curve indicates where the model predicted a low score on that section of the population while the right side of the curve indicates where the model predicted a high score.

The points on the Lift Chart indicate the average percentage in each bin. The “Predicted” blue line displays the average prediction score for the rows in that bin. The “Actual” orange line displays the actual percentage for the rows in that bin. In general, the steeper the Actual line is, and the more closely the Predicted line matches the actual line, the better the model. A close relationship between these two lines is indicative of the predictive accuracy of the model.

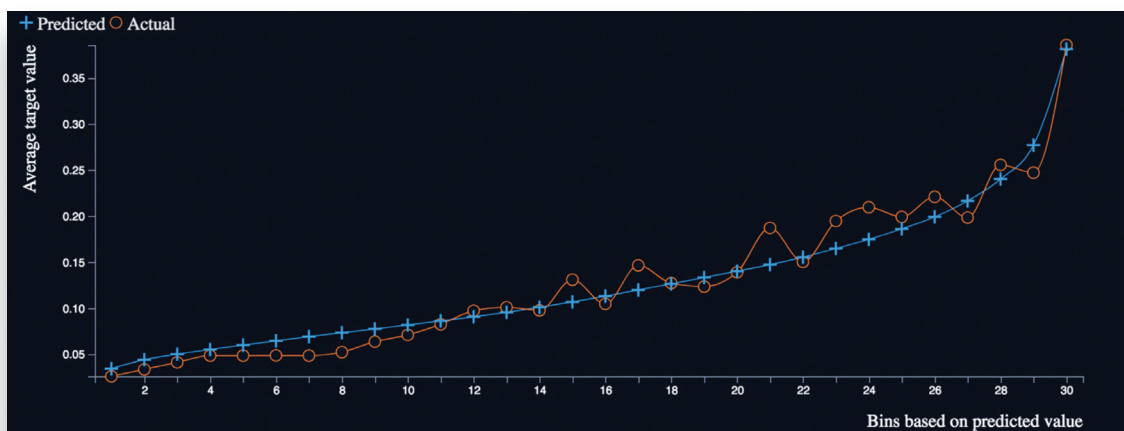


Figure 4: Lift Chart

Similar analyses can also be carried out by sorting the chart on other modelling inputs. This allows the detection of model biases in certain areas or groups of cases, which permits the identification of areas which might benefit from the building of separate models. (For instance, in Figure 5 below, the left-hand pane shows that the model under-predicts default rates for higher credit grades and overpredicts in the middle grades. The right-hand pane shows that the model accurately models credit risk for 36-month loans but over-predicts for 60-month loans; as a result of this insight, model validators could suggest that a separate model is built for the longer loan term.)



Figure 5: Identifying Model Segmentation

FEATURE IMPACT – Identifying Important Features in the Model

The feature sensitivity testing technique (sometimes called Permutation Importance) works by altering model input data and observing the effect on the resultant predictions. The Feature Impact for a given column measures how much worse a model’s accuracy would be if predictions are made after randomly shuffling that column (while leaving other columns unchanged), thus effectively neutralizing that column; therefore, a larger drop in modeling accuracy denotes the importance of a particular feature. It follows that columns with higher feature impact scores are more important in driving the model’s predictions, and ranking data features by their feature impact therefore allows modelers and validators to understand which features have the greatest impact on model performance. It is common practice to normalize the scores so that all features’ impacts are expressed as percentages of that of the most impactful column.

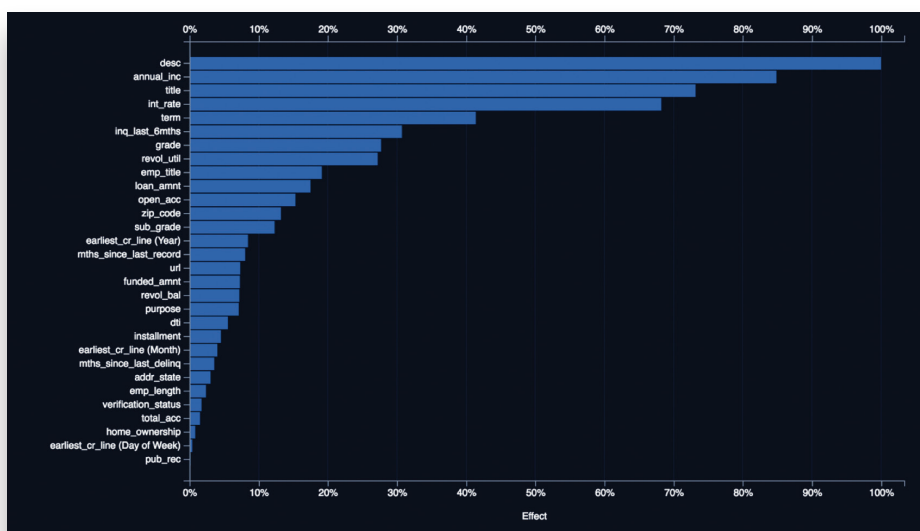


Figure 6: Feature Impact – Permutation Importance

MODEL SENSITIVITY – Partial Dependence Plots

For more complex models, a comparably simple parametric description is not available, making the interpretation of these models more difficult. To address the difficulty of sensitivity analysis for his gradient boosting machine, Friedman (2001) proposed the use of partial dependence plots. Partial dependence plots show the average partial relationship between a set of predictors (the input variables) and the predicted response (the target variable).

A partial dependence plot is depicted below in Figure 7. The yellow partial dependence line depicts the marginal effect of a feature on the target variable after accounting for the average effects of all other predictive features, thus providing a view of “all else equal” sensitivity of the target in respect of a given feature. It indicates how, holding all other variables constant, the value of this feature affects the prediction; this is calculated by setting all values in a given feature to the same value in turn (for a range of values) and then plotting the average response of the target. These values help determine how the value of each feature affects the target and the shape of the yellow data points describes the model’s view of the marginal relationship between the selected feature and the target.

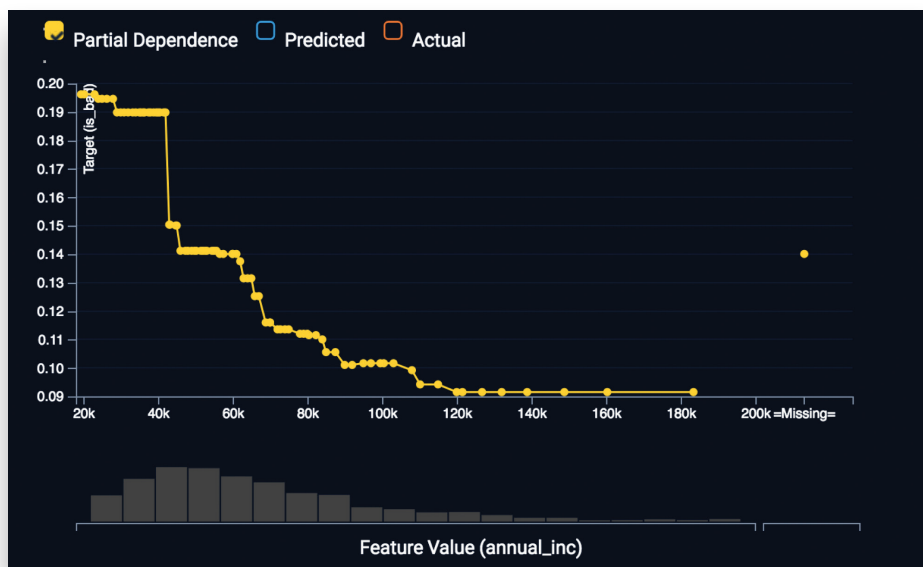


Figure 7: Partial Dependence Plot

ONGOING MODEL MANAGEMENT AND PERFORMANCE MONITORING

The ongoing monitoring and periodic review of a production model’s performance is a critical component of a sound model risk management framework and it occurs after a model has been deployed. This periodic review process should incorporate information that has become available since the model was developed, which allows validators to affirm the model’s ability to continue to meet the model objectives for which it was approved. Effective model management is complex at scale since there can be hundreds or even thousands of models that have been deployed into production at banks. If these models are ineffectively managed, then that can introduce substantial model risk across the enterprise.

Model Management is a powerful feature of DataRobot that allows users to proactively monitor and manage all deployed machine learning models (models created both inside and outside of DataRobot) to maintain peak performance and ensure compliance with the appropriate regulatory and institutional policies. By doing so, model stakeholders can have confidence that the machine learning models driving their business are accurate and reliable throughout changing market conditions.

DataRobot’s Model Management allows stakeholders to monitor and manage all their inventory of deployed models from a central portal and apply embedded best-in-class data science expertise to review model performance and detect model decay sooner, as shown below in Figure 8 below. This consolidated overview encourages collaboration across groups and model stakeholders – from model developers, IT, and model validators – and provides an aggregate summary of the health, usage, and auditable historical lineage (i.e., change-log) of each deployed model across the model inventory. Therefore, model risk can be assessed and effectively managed by allowing rapid, real-time drill down into each model to analyze latency, throughput, error rates, and customizable performance alerts across the entire model inventory.

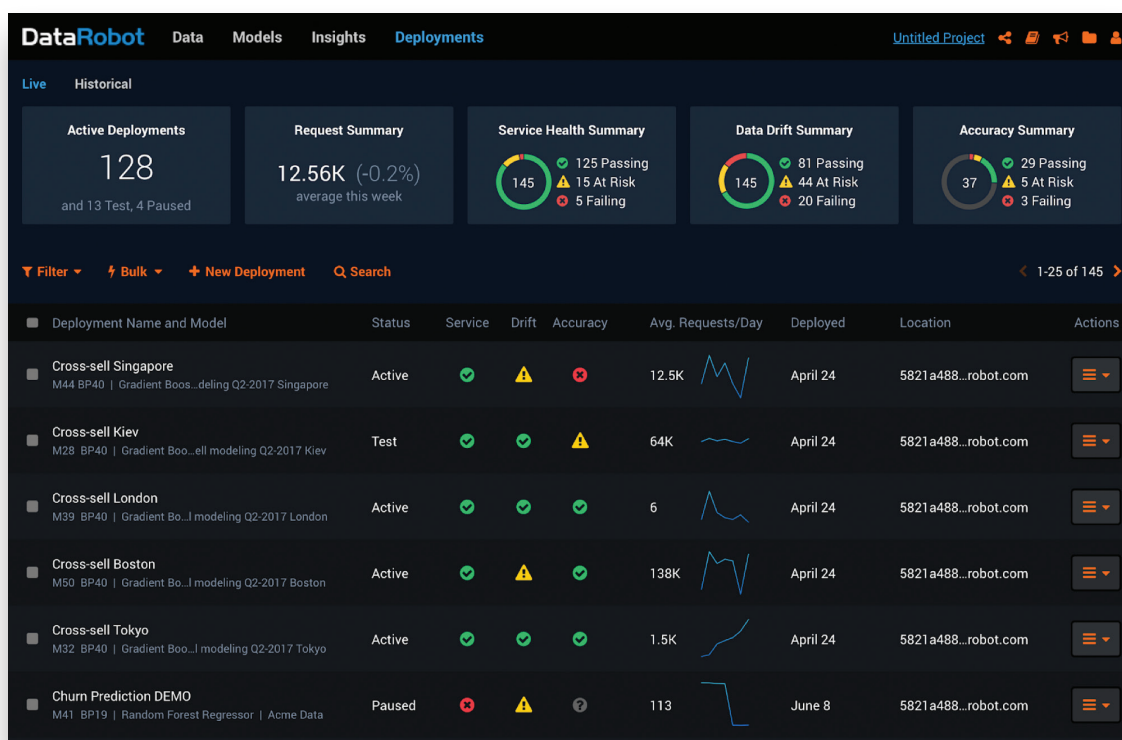


Figure 8: Model Diversification – DataRobot Leaderboard

Using DataRobot’s competitive advantage in machine learning automation, the ongoing performance of deployed models is automated and allows model stakeholders the rapid ability to detect model decay or potential performance issues much sooner than traditional manual processes. Stakeholders can interactively explore key performance metrics to effectively assess the root cause and severity of model decay (e.g., model or data drift) by examining how model performance has changed over time, as shown in Figures 9 and 10 below.

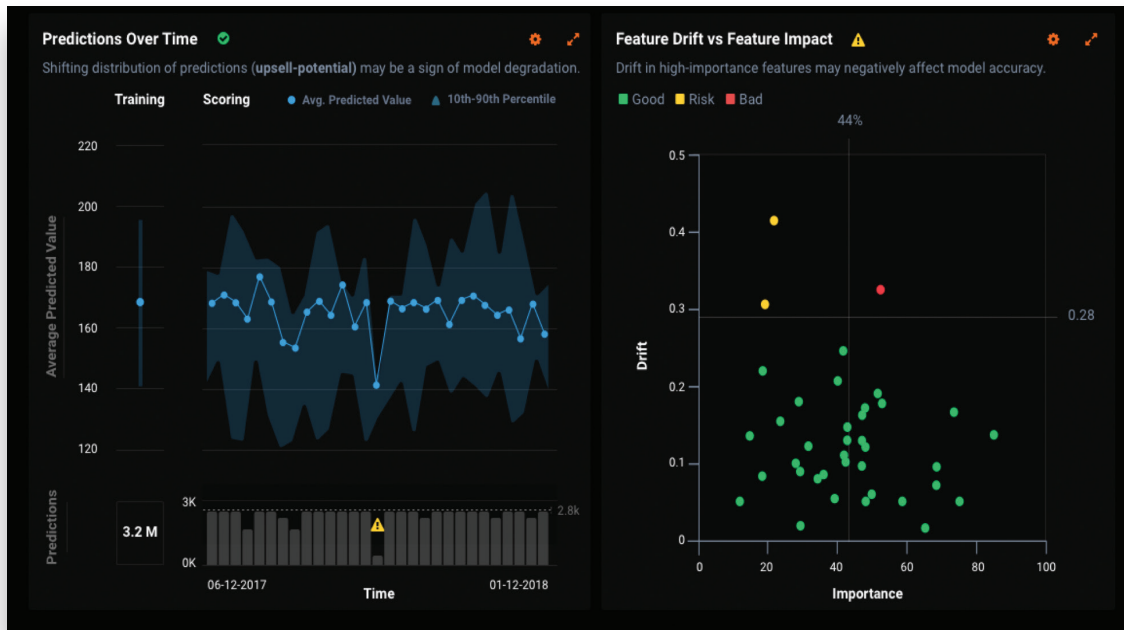


Figure 9: Model Drift

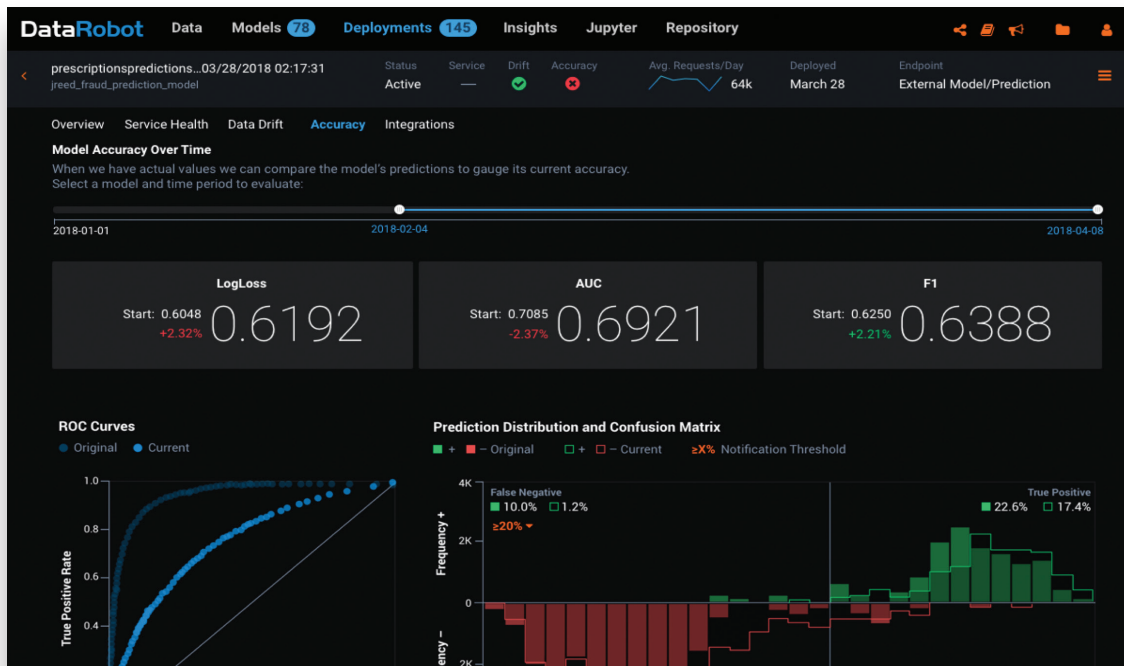


Figure 10: Model Accuracy Over Time

Conclusion

AI and machine learning models provide an excellent opportunity for financial institutions to improve their predictive modeling capabilities in a number of areas; providing competitive advantages in revenue generation, increased efficiency and risk management. Senior management needs to have a stake in the increased uptake of machine learning models, with key drivers being accuracy, transparency, speed, cost, flexibility and auditability.

However, it is important to hold the right balance between risk and reward; ensuring that there are robust controls in place on the design, development, implementation, validation and use of models. Utilizing the right technologies will help firms protect their franchise and stand the test of regulatory scrutiny, while also enabling effective strategic decision making. In this regard, a firm's model framework should be a triangulation of theoretical modeling, empirical evidence and expert judgement. This needs to go beyond just regulatory box-ticking with firms linking model outcomes to strategic levers that can be pulled by senior management in order to influence key performance indicators such as revenue, cost, regulatory capital, liquidity and leverage.

The stakes in managing model risk are at an all-time high. Fortunately, when complemented with an effective model risk management framework, automated machine learning provides an effective solution to ensure a robust model risk management framework for your enterprise and drive strategic outcomes for financial institutions in the age of AI and machine learning.



Established in 2004, Avantage Reply (a member firm of Reply) is a pan-European specialised management consultancy delivering change initiatives in the areas of Compliance, Finance, Risk and Treasury.

Visit Avantage Reply on LinkedIn 

www.avantagereply.com

DataRobot

DataRobot helps enterprises embrace artificial intelligence (AI). Invented by DataRobot, automated machine learning enables organizations to build predictive models that unlock value in data, making machine learning accessible to business analysts and allowing data scientists to accomplish more faster. With DataRobot, organizations become AI-driven and are enabled to automate processes, optimize outcomes, and extract deeper insights.

Learn more at datarobot.com