TATA CONSULTANCY SERVICES

Building on belief

Streamlining machine learning-based model management framework: A BFS industry perspective

Abstract

Mass digitization and data-driven transformation have led to greater uptake of complex deep learning or neural network algorithms by financial institutions, enabling reimagination of the digital core and innovation in business models. However, due to the black-box or opaque nature of these models, maintaining auditability and transparency have become equally important areas of consideration. Subsequently, ensuring the interpretability of these complex models is an emerging challenge due to the regulatory implications.

A comprehensive and resilient machine learning model management framework should address fair training or learning, transparency, data governance, and streamlined operationalization in order to ensure conceptual soundness and accelerated validation cycles of the model deployment process. In this white paper, we review the critical aspects of the model management framework that help in reducing time-to-market for new-age machine learning (ML) models and ensuring that the model's operational performance is not compromised.

Managing AI and ML models

The adoption of AI and ML has gone up manifold across the value chain of financial models with the aim of driving efficiency, enhanced customer experience, and diverse data-driven analytics. ML models are being adopted across various touchpoints—customer onboarding, fraudulent transaction detection, regulatory reporting, and alternative (alt-data) financial data ingestion. This, in turn, necessitates a robust build and deployment framework to transform the AI/ML setup and integration pipeline for speeding time-to-market and enabling continuous quality assessment.

Reimagining the model management components

Before institutions resort to this build framework, they need to reimagine the traditional model risk management framework.[1] The essence of this framework can be captured within three core components. First, quick rebuilding, training, and deployment of the models using new technologies and alternative data, allowing them to fail and pivot fast. A streamlined model pipeline structure (ModelOps) is needed to achieve this in an agile mode. Second, clean and trusted data is needed to achieve a fair, explainable, and reliable model outcome over the lifetime of the model. This requires an integrated data and model validation framework, which should be closely coupled with the model governance process. Last, setting risk tolerance levels and developing model materiality assessment scorecards are essential for model oversight and control. As the components evolve and get more integrated with the overall model risk framework, model governance and oversight framework become more intuitive with real-time insights.

---

[1] Federal Reserve, Supervisory Guidance on Model Risk Management, April 4, 2011, Accessed January 14, 2021, https://www.federalreserve.gov/supervisionreg/srletters/sr1107a1.pdf

Let us take a look at the three critical components of ML model management.

Data validation

Data validation comprises discovery, preprocessing, feature engineering, and drift monitoring for the model input data and is handled by the first line of defense. Let us exemplify this with a retail borrower risk scoring model that classifies the creditworthiness of a borrower. By analyzing a customer's demographic details, financial status, behavioral records, and social information, potential risk factors can be identified. Here, data discovery or feature engineering intends to pick up the most critical set of input parameters correlated with an output class. For example, a customer with 'low' creditworthiness may be strongly linked with his/her credit history, annual income, and criminal record input parameters. It is also possible to hand-pick the appropriate algorithm by exploring the data and analyzing the task in hand – be it a classification, regression, or forecasting task. Since creditworthiness is a classification task (high, medium, or low creditworthiness), the best-performing classification algorithm such as decision tree, random forest, or K-means clustering could be selected through a ranking exercise.

When it comes to the model fairness aspect, there is also a great deal of data validation procedures involved. 'Sensitive features' or 'sensitive attributes' must be detected to identify features such as age, gender, and political affiliation that are sensitive towards determination of the outcome of a data segment. Sensitivity analysis or SHAP scores can help identify the correlation between input features and the outcome.

While detailed data helps build a robust and diverse model, protecting personally identifiable data (PII), like age, address, gender, and ethnicity, among others, without significantly compromising model performance is yet another critical decision. In such scenarios, consolidated PII parameters can be encapsulated in composite parameters (like age and location) to depict model behavior in the model explainability phase.

Segregated data validation components such as data mining, visualization, quality check, and ETL can be orchestrated using a streamlined DataOps workflow. DataOps helps to synchronously manage these tasks along with model training and model inference pipelines to achieve a seamless ML model training, validation, and testing workflow.

Outcome validation

Model outcome validation is handled by both the first and second lines of defense. It plays an important role in running performance-based recalibration of an existing model in a periodic fashion. The performance of an ML model can significantly sway due to variance in the input data distribution (data drift), irregular tuning of models, or inappropriate selection of algorithms (concept drift). This causes periodic decay in the model performance and may eventually leave the model unreliable for business decisions. In our earlier example of a retail borrower risk scoring model, if the majority of borrower data becomes defaulted customer data over time, the learning process gradually becomes biased with more probability of inaccurate prediction due to an imbalanced learning dataset. Further, cross-validation with various sets of data can help validate the outcome of a model and frequently assess the performance deviation from the benchmark. With more and more automation in the model deployment pipeline, model outcome validation can be integrated to trigger alerts in case of a threshold breach.

Model oversight and control

An organization needs to adopt a resilient model oversight framework by defining stringent model design development policies, periodic reviews, and materiality scoring systems. This is primarily handled by the second line of defense.

For ongoing oversight of 'black-box' models, we also need to leverage open-source community-driven frameworks which can help assess and compare performance, disparity, sensitive features, and performance-fairness tradeoff metrics across a host of models. The insights gained out of this ensures the models provide non-discriminatory outcome without compromising accuracy levels. Many of these frameworks are built by developer communities and can be leveraged using popular open-source languages and APIs.

From our example of the retail borrower risk scoring model, oversight and control can be exercised by documenting all the various input features of the model and their correlation with the outcome (conceptual soundness), the periodic output of the model, and the materiality risk of the model. Setting appropriate controls will ensure that the model inaccuracy level or other inherent risks do not outweigh the reward or surpass the risk threshold.

Model oversight and control components can be integrated with the centralized model repository systems to generate management reports containing model risk scores and commentaries.

Streamlining ML model workflow and scaling: ModelOps

'Black-box' model risk management must be more agile and less susceptible to new-age model management nuances such as alt-data bombardment, high-velocity data, data drift, etc. Often, there is a significant increase in time-to-market for ML models due to time-consuming data cleansing, preprocessing, model training, and hyperparameter tuning procedures.

ModelOps is a target state for creating a seamless release pipeline of ML models by harmonizing their build, training, validation, release, integration, and periodic review phases in an agile world. In the aforementioned components of model risk management, ModelOps prevails across aspects such as monitoring of model performance, drift, and health; understanding data distribution, features, their interrelation, and causality with the predicted output; periodic audit, and reporting. Thus, ModelOps is also the starting point to make ML model frameworks more transparent, agile, and auditable.

While the complex ML algorithms are spreading deeper and wider across the model development and validation value chain, it is becoming imperative that financial institutions establish operational measures (such as prediction accuracy, prediction fairness, prediction reliability, prediction explain-ability, and so on) for people to be able to justify and rely on their AI-driven decisions. Adoption of an overarching technology governance framework across data models and strengthening ModelOps capability can address and remediate potential vulnerabilities of AI-based financial black-box models and position institutions as a trusted partner in their digitalization journey.

Conclusion

There is no denying the fact that ML systems are gradually becoming the lifeline of financial institutions. Unless controlled through strategic guardrails and thresholds, unintended autonomous decisions taken by these systems could lead to dangerous, unethical, and widespread impacts on the financial industry and society.

However, we cannot slow down the pipelines for operationalizing these models due to compliance obligations. Therefore, we need the ModelOps space to continuously evolve so that innovation in ML labs can be released into the market at an optimal yet sustainable pace, without compromising on the realizable business benefit.

About the authors

Ushasi Sengupta

Ushasi Sengupta is a senior research analyst in TCS' Corporate Marketing Research group. She specializes in the banking, financial services, and insurance sector. She has a bachelor's degree in Electronics and Telecommunication Engineering from West Bengal University of Technology, Kolkata, India.

Sanjukta Dhar

Sanjukta Dhar leads the BFSI CRO Strategic Initiative for TCS Canada. She brings to the table more than 18 years of risk transformation program experience across major banks and financial services organizations. Sanjukta holds a bachelor's degree in Civil Engineering from Jadavpur University, India.

TATA CONSULTANCY SERVICES

Awards and accolades

TOP 3 IT SERVICES BRAND

FASTEST GROWING IT SERVICES BRAND FOR THE DECADE 2010 - 2020

Contact

For more information on TCS' Banking, Financial Services, and Insurance (BFSI) unit,

Visit https://www.tcs.com/banking-financial-services and https://www.tcs.com insurance

Email: bfsi.marketing@tcs.com

About Tata Consultancy Services Ltd (TCS)

Tata Consultancy Services is a purpose-led transformation partner to many of the world's largest businesses. For more than 50 years, it has been collaborating with clients and communities to build a greater future through innovation and collective knowledge. TCS offers an integrated portfolio of cognitive powered business, technology, and engineering services and solutions. The company's 500,000 consultants in 46 countries help empower individuals, enterprises, and societies to build on belief.

Visit **www.tcs.com** and follow TCS news **@TCS**.

Corporate Marketing | Design Services | M | 11 | 21