Machine Learning Operations (ML Ops) For Healthcare

Deploying and managing machine learning models for use in operational or clinical workflows is no easy task.

> *Approximately 90% of models never make it into production, and successfully maintaining the 10% that do requires comprehensive ML Ops capabilities.[1]*

ML Ops enhances the DevOps procedures used to deploy traditional software with the unique requirements of ML systems. For models to be integrated into workflows and drive improved health and financial outcomes, data scientists must integrate with standard IT DevOps procedures while also handling enormous volumes and varieties of data from diverse sources, monitoring model performance, reacting to degradation in model accuracy, and maintaining robust governance and change management processes.

Seamless ML Ops with ClosedLoop

ClosedLoop's ML Ops enables data scientists to seamlessly deploy models and drive success at scale. To support continuous improvement and learning in an ever-changing healthcare environment, ML Ops provides a suite of capabilities to not only configure and manage deployments, but automate model performance monitoring and effortlessly retrain models, with support for audit and governance.

🚀 Comprehensive Deployment Management

With ML Ops no-code deployment management, data scientists can **efficiently configure and automate every step in the pipeline – from data uploads, data prep, and normalization to model runs and report generation** – with built-in error handling for schema changes and data anomalies. Data scientists can configure deployments for Batch or Realtime models, which can also be managed through REST APIs.

Configurations are customizable end-to-end and can be reused across an organization's deployments, ensuring standardized ETL and preventing retrained models from disrupting operations when released. Moreover, data scientists can effortlessly govern deployments with functionality for staging, versioning, and archiving.
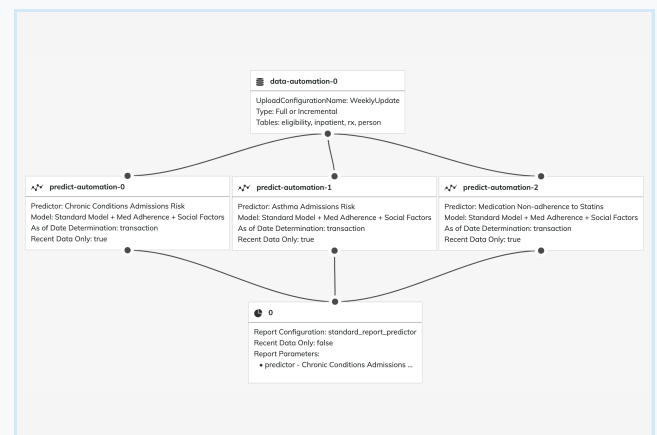

Sample Deployment

### ⏱ Automated Model Performance Monitoring

Automated monitoring of a model's performance streamlines identification of degradation in model accuracy and feature drift, preventing inaccurate predictions from influencing business operations. To this end, data scientists can compare test and live model accuracy, with built-in model retraining benchmarks.

Data scientists can dig into the specific causes of model degradation using Feature Drift. This capability uses the
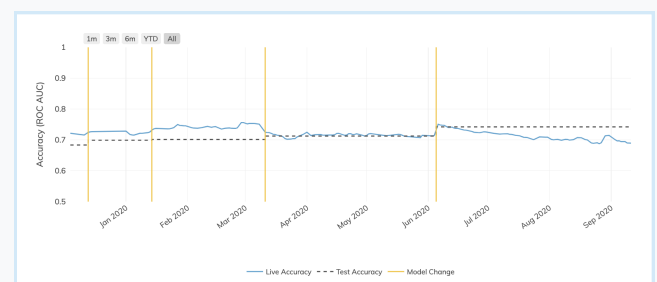
Accuracy Over Time

Population Stability Index to monitor the distribution of feature values among the current population versus the population used to train the model. This shift is assessed in tandem with the feature's relative importance, helping teams evaluate the need for model retraining and avoid alert fatigue.
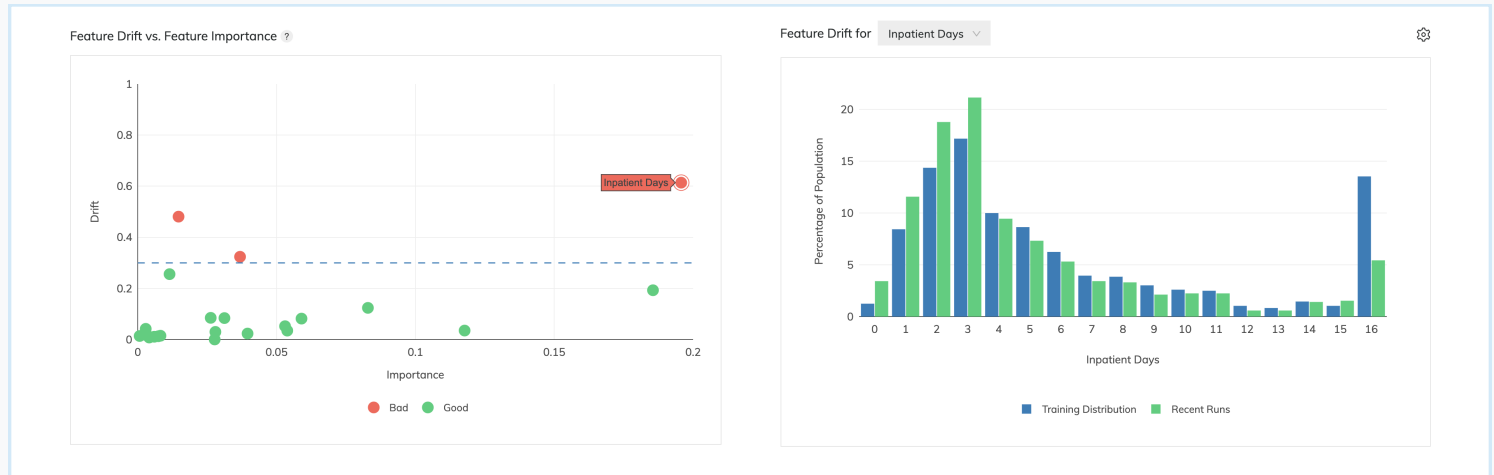
Drift Versus Importance

⚙ Centralized Audit and Governance

ML Ops supports service quality monitoring, **acting as a central repository for tracking data uploads, model runs, and reports**. Data scientists can audit and govern deployments through access to a full history of event logs. They can institute and maintain guardrails by configuring warnings, adjusting individual permissions, and managing version history. Comprehensive versioning across data processing pipelines supports rollback to previous versions at any time.

ML Ops also supports audit and measurement of algorithmic biases which may exacerbate health disparities. Bias is measured with Group Benefit Equality (GBE), a metric that explicitly addresses the inadequacy of standard fairness metrics in healthcare. GBE is transparent, explainable, and uses clearly defined thresholds to assess model bias.

Why it Matters for AI-Enabled Healthcare Organizations (HCOs)

ClosedLoop combines best-in-class ML Ops capabilities with healthcare-specific content to help HCOs address common deployment pitfalls and facilitate successful integration of machine learning models.

As an industry, healthcare is evolving. This transformative process will continue to accelerate, and it is vital that HCOs are equipped with robust ML Ops capabilities enabling them to respond to emerging challenges which include:

- **Accuracy and Bias Monitoring** – Models must constantly be monitored and audited for bias and accuracy. To this end, ML Ops capabilities are necessary and cannot be substituted.
- **Dynamic Conditions and Shifting Populations** – ML Ops is essential to keep deployed models constantly updated and running reliably in production amidst frequent changes to healthcare systems, interventions, code systems, program eligibility requirements, chronic condition prevalence, and more.
- **The Explosion of Healthcare Data** – Data scientists must incorporate more data than ever before from a burgeoning variety of new sources (e.g., public, interoperability feeds, new clinical sources, proprietary intervention data).
- **New Strategic Opportunities** – HCOs are increasingly engaging in direct contracting, value based partnerships, and other risk-taking endeavors. This demands rapid scaling of existing and new machine learning use cases, with robust deployment management to ensure continued model performance monitoring.

To evolve along with healthcare demands that HCOs have the capability to seamlessly deploy models and comprehensively manage, monitor, and govern them. ClosedLoop's robust ML Ops gives HCOs this capability.

1. "Why do 87% of data science projects never make it into production?" Venturebeat, 19 Jul. 2019, https://venturebeat.com/2019/07/19/why-do-87-of-data-science-projects-never-make-it-into-production/. Accessed 14 May 2021.