



Model Deployment with DataRobot

DataRobot powers the AI-Driven Enterprise. Its automated machine learning platform harnesses hundreds of cutting-edge algorithms to discover the best predictive models for every situation, enabling users throughout the enterprise to make dramatically **smarter, faster business decisions.**

Solution Brief

Predictive models are not meant to sit on the shelf. These models are meant to drive real business value by serving up predictions continuously in production environments, or delivering predictions immediately on small projects. Even if an organization is using the DataRobot automated machine learning platform to quickly generate highly accurate predictive models, deploying those models is a separate step that requires the attention of data scientists, software developers, IT professionals, and business leaders alike.

DataRobot Model Deployment

As part of its focus on enabling the AI-driven enterprise, the DataRobot automated machine learning platform offers an end-to-end solution for predictive modeling, from development to deployment and operationalization. DataRobot offers five methods for addressing model deployment that enable organizations to quickly deploy the winning model.



Enterprise Automation

On-Demand Analysis

Portable Prediction

Separate Development and Production Environments

Working with Big Data

Enterprise Automation

Every model DataRobot builds can publish a REST API endpoint, making it easy to integrate into modern enterprise applications and effectively turning DataRobot into a prediction engine. Instead of waiting months to write scoring code and deal with the underlying infrastructure, organizations derive business value from machine learning in a fraction of that time. Whether new data comes at 10 times per second or 10 times a day, automated predictions improve mission-critical applications by incorporating the most up-to-date machine learning techniques and offering a path to continual model upgrades through the DataRobot platform.

With DataRobot, the model and the prediction server are as rock-solid and reliable as the database and application, offering security, low latency, and throughput SLAs. Unlike the difficult process of working with a developer trying to duplicate a model's feature engineering process, DataRobot makes it easy to quickly develop and deploy highly accurate predictive analytics.



On-Demand Analysis


In some cases, it is important to make predictions on an ad hoc basis. These cases include on-demand predictions for weekly or monthly activities, such as which customers are likely to churn, or one-time predictions, like an analysis of which marketing programs will result in business opportunities. DataRobot enables on-demand analysis for any model, providing a quick and easy way to deliver predictions that either do not require inclusion in production systems or are special requests.

DataRobot's Predict functionality, a drag-and-drop prediction interface, removes the dependency on external teams like software development and IT and allows users to get predictions when they need them. This capability for on-demand analysis is great for agile businesses that change frequently and need rapid prototyping flexibility to test models before integrating them into production systems.

Portable Predictions

For some industries, the combination of data volume and latency requirements means that models need to be installed where the predictions are needed, which is often some type of low-powered device. Network latency or connectivity problems may introduce issues that make it too risky to deliver predictions from centralized servers with the reliability the application requires.

For these cases, DataRobot's Scoring Code Export offers a simple, self-contained download of the chosen model. The code is available as an executable .jar file or as Java source code, and can be deployed anywhere Java runs. The generated scoring code contains all of the dependencies and preprocessing steps and model parameters needed to produce predictions at the required levels of speed and accuracy.



Separate Development and Production Environments

For highly regulated enterprises like banks, keeping development work from affecting models in production is often a key consideration in model deployment. This separation stops bugs and errors from progressing from development into deployed models, which in turn prevents critical errors by the model in question – or worse, the accidental deletion of critical models that are in use.

DataRobot's Standalone Scoring Engine separates these two environments so that models can be tested and implemented in a stable, isolated environment. The Standalone Engine has the capability to run imported models without ever touching the development server from which they were exported.

Working with Big Data

Transferring large amounts of data is resource-intensive, and sometimes not possible depending on the amount of data. Usually, data being scored by a machine learning algorithm is stored on and accessed from a central server. But with extraordinarily large datasets, this poses a problem for enterprises that are unable to transfer and host that much data at once.

Spark Scoring with DataRobot allows enterprises to score data for machine learning where it is located instead of having to transfer and host that data on a central server. This allows businesses to run models produced using DataRobot on potentially huge data sets without changing the storage location of the data from its current instantiation on a Hadoop network.



Deployment Services

Need help deploying your predictive models? The DataRobot AI Acceleration Service offers a number of services to jumpstart your AI success, including support for model deployment in your production environment. [Find out more here.](#)

Deployment Considerations

Not sure which DataRobot deployment option is right for you? Use this chart to help determine which option would work best in your environment, or contact DataRobot to discuss your needs.

Deployment Options	Use Case Requirements and System Architecture						
	Frequent Model Refreshes	Scheduled Batch Scoring	Quick Deploy	Disconnected from Network	Big Data	Low Latency Predictions	Separate Development and Production Environments
Rest API	✓	✓	✓		✓	✓	
Scoring Code Export		✓		✓	✓	✓	✓
Spark Script	✓	✓	✓	✓	✓		✓
Ad hoc via GUI	✓						
Standalone Scoring Engine	✓	✓	✓	✓	✓	✓	✓

Contact Us

DataRobot
One International Place, 5th Floor
Boston, MA 02110
www.datarobot.com
info@datarobot.com