

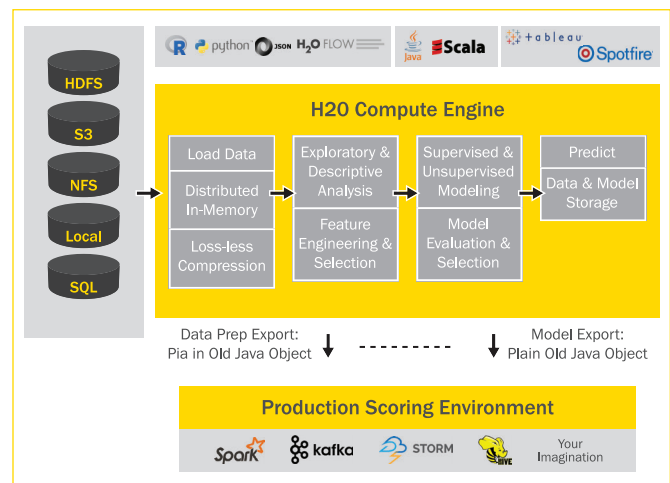


H2O Platform

100% open source, fully distributed in-memory machine learning platform with linear scalability

H2O makes it possible for anyone to easily apply machine learning and predictive analytics to solve today's most challenging business problems. It intelligently combines the following features:

- **Best of Breed Open Source Technology** – Enjoy the freedom that comes with big data science powered by open source technology. H2O was written from scratch in Java and seamlessly integrates with the most popular open source products like Apache Hadoop® and Spark™.
- **Easy-to-use WebUI and Familiar Interfaces** – Set up and get started quickly using either H2O's intuitive web-based Flow GUI or familiar programming environments like R, Python, Java, Scala, JSON, and through our powerful APIs. Models can be visually inspected during training.
- **Data Agnostic Support for all Common Database and File Types** – Easily explore and model big data from within Microsoft Excel, R Studio, Tableau and more. Connect to data from HDFS, S3, SQL and NoSQL data sources. Install and deploy anywhere, in the cloud, on premise, on workstations, servers or clusters.
- **Massively Scalable Big Data Munging and Analysis** – H2O Big Joins performs 7x faster than R data.table in a benchmark, and linearly scales to 10 billion x 10 billion row joins. Train a model on complete data sets, not just small samples, and iterate and develop models in real-time with H2O's rapid in-memory distributed parallel processing.
- **Real-time Data Scoring** – Rapidly deploy models to production via plain-old Java objects (POJO), model-optimized Java objects (MOJO) or REST API. Score



new data against models for accurate and fast predictions in any environment.

H2O supports the following distributed algorithms:

Supervised Learning

- Statistical Analysis
 - Generalized Linear Models
 - Naïve Bayes
- Ensembles
 - Distributed Random Forest
 - Gradient Boosting Machine
- Deep Neural Networks
 - Deep learning

Unsupervised Learning

- Clustering
 - K-means
- Dimensionality Reduction
 - Principal Component Analysis
 - Generalized Low Rank Models
- Anomaly Detection
 - Autoencoders

Advanced Features for Data Scientists

- Automatic standardization of the predictors
- Automatic initialization of model parameters (weights & biases)
- Automatic adaptive learning rates
- Ability to specify manual learning rate with annealing and momentum
- Automatic handling of categorical and missing data
- Regularization techniques to manage complexity
- Automatic parameter optimization with grid search
- Early stopping based on validation datasets
- Automatic cross-validation

Quick Start

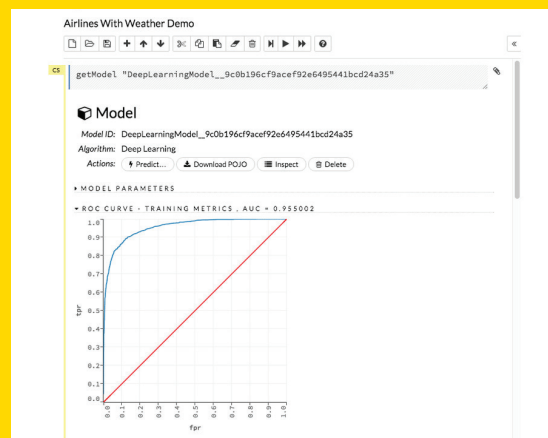
- Go to h2o.ai/downloads. Download H2O. This is a zip file that contains everything you need to get started.
- From your terminal, run:

```
cd ~/Downloads
unzip h2o-3.10.2.2.zip
cd h2o-3.10.2.2
java -jar h2o.jar
```

- Point your browser to <http://localhost:54321>

H₂O FLOW

- H2O provides an interactive web interface called H2O Flow that seamlessly blends an command line, text based shell with a modern notebook style GUI.
- With Flow, data scientists can import data, inspect it, join with another frame, split train & test, build models with grid search, evaluate results and export POJO, all without writing a single line of code



About H2O.ai

H2O.ai is focused on bringing AI to businesses through software. Its flagship product is H2O, the leading open source platform that makes it easy for financial services, insurance and healthcare companies to deploy machine learning and predictive analytics to solve complex problems. More than 8,500+ organizations and 75,000+ data scientists depend on H2O for critical applications like predictive maintenance and operational intelligence. The company accelerates business transformation for 107 Fortune 500 enterprises, 8 of the world's 12 largest banks, 7 of the 10 largest auto insurance companies and all 5 major telecommunications providers. Notable customers include Capital One, Progressive Insurance, Transamerica, Comcast, Nielsen Catalina Solutions, Macy's, Walgreens, Kaiser Permanente, and Aetna.

Follow us on Twitter [@h2oai](https://twitter.com/h2oai). To learn more about H2O customer use cases, please visit <http://www.h2o.ai/customers/>. **Join the Movement.**

