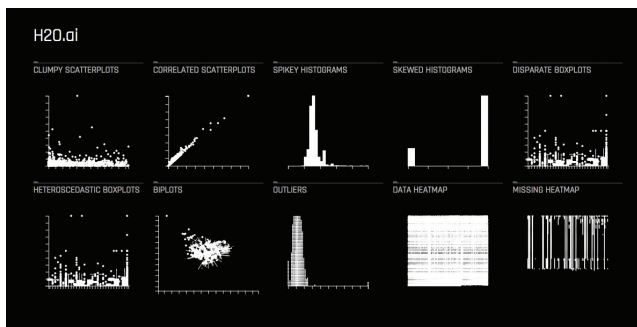
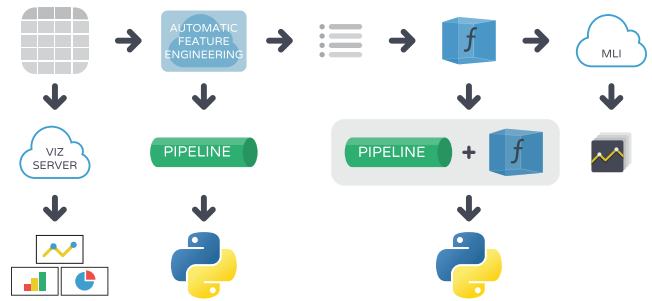


# Driverless AI: Automatic Machine Learning for the Enterprise

Driverless AI speeds up machine learning workflows by automating visualizations, feature engineering, model tuning, ensembling, and model deployment.

Driverless AI takes a raw dataset and automatically visualizes the most interesting patterns for data exploration. It then applies automatic feature engineering to increase accuracy by using recipes for solving a wide variety of use cases. Next, it auto-tunes model parameters and provides the user with the model that yields the best results. Lastly, it gives plain English explanations of model results. The output of the experiments are then available to export as a scoring pipeline for production deployment. Driverless AI enables users of all backgrounds to draw the most value from their data.



## Automatic Visualization

- Allows users to visualize large datasets in the form of various graphs and charts without having to write code
- Takes huge datasets and displays outliers and trends in an interpretable way
- Uses statistics to automatically decide which visualizations to present to the user
- Exploratory tool that presents an overview of the distribution of data

## Automatic Feature Engineering

- It employs a library of feature transformations to automatically engineer new features for a given dataset, based on best practices from expert data scientists.
- A variable importance chart shows the significance of newly engineered features.
- Data scientists can spend less time exploring data and engineering features; as a result, they can work on multiple projects simultaneously and increase their efficiency.



## Model Interpretability

- The Machine Learning Interpretability feature provides users with clear and concise explanations of model results.
- Four dynamic graphs are generated automatically: K-LIME, Variable Importance, Decision Tree, and Partial Dependence; each one helps explore the model output more closely.
- Global interpretability techniques allow the user to broadly understand models at the level of the full dataset.
- Local interpretability techniques explain predictions for individual data points and produce plain English explanations with reason codes.

## Scoring Pipeline

- A scoring pipeline is available after a successfully completed experiment. This pipeline includes a scoring module and a scoring service.
- The scoring service hosts the scoring module as an HTTP or TCP service.
- Currently supports a Python scoring module and a Java based MOJO scoring module for low latency scoring.
- Export the pipelines to operate on new data in production.

---

## Personas & Benefits

### Business Users

- Users with domain knowledge will understand the new features generated by Automatic Feature Engineering and their impact on predictions.
- Analysts who have used visualization tools in the past can easily transition into using the exploratory AutoViz tool and the Machine Learning Interpretability tool to understand the impact of a particular variable on their model.

### Data Analysts

- Using basic data science knowledge, users can better understand features generated by Driverless AI and use outcomes to draw valuable insights for their businesses.
- AutoViz makes data exploration extremely simple so users can identify trends in data through visualizations.

### Data Scientists

- AutoViz simplifies the data exploration process.
- Increase productivity of data scientists by enabling them to work on multiple projects simultaneously.
- The visualizations produced by the Machine Learning Interpretability tool streamline the process of understanding the impact of a model.

## 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 13,000 organizations and 130,000+ data scientists depend on H2O for critical applications like predictive maintenance and operational intelligence. The company accelerates business transformation for 222 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.

Follow us on Twitter @h2oai. To learn more about H2O customer use cases, please visit <http://www.h2o.ai/customers/>. Join the Movement.

## SPECIFICATIONS

### For CPU Installations

#### Minimum Hardware

- x86-64 CPU
- 64 GB RAM

#### Recommended Hardware

- Dual socket Intel Xeon x86-64 or IBM Power 8 CPU's

#### Software

- Docker CE

### For GPU Installations

#### Minimum Hardware

- NVIDIA K80

#### Recommended Hardware

- NVIDIA P100 or higher
- NVLink and NCL enabled

#### Software

- Ubuntu 16.04
- CUDA 8, CUDA 9
- NVIDIA Docker 1.0.1

---

### For Cloud Installations

#### EC2 Instances

- P2 or G3
- Ubuntu 16.04

#### Azure Instances

- NC or NV
- Ubuntu 16.04

#### Google Cloud

- K80 or P100
- Ubuntu 16.04

Driverless AI is also available on Nimble, Paperspace and other cloud environments.

---

## Licensing

Driverless AI is a commercially licensed product. For more information, contact [sales@h2o.ai](mailto:sales@h2o.ai).