


H2O.ai

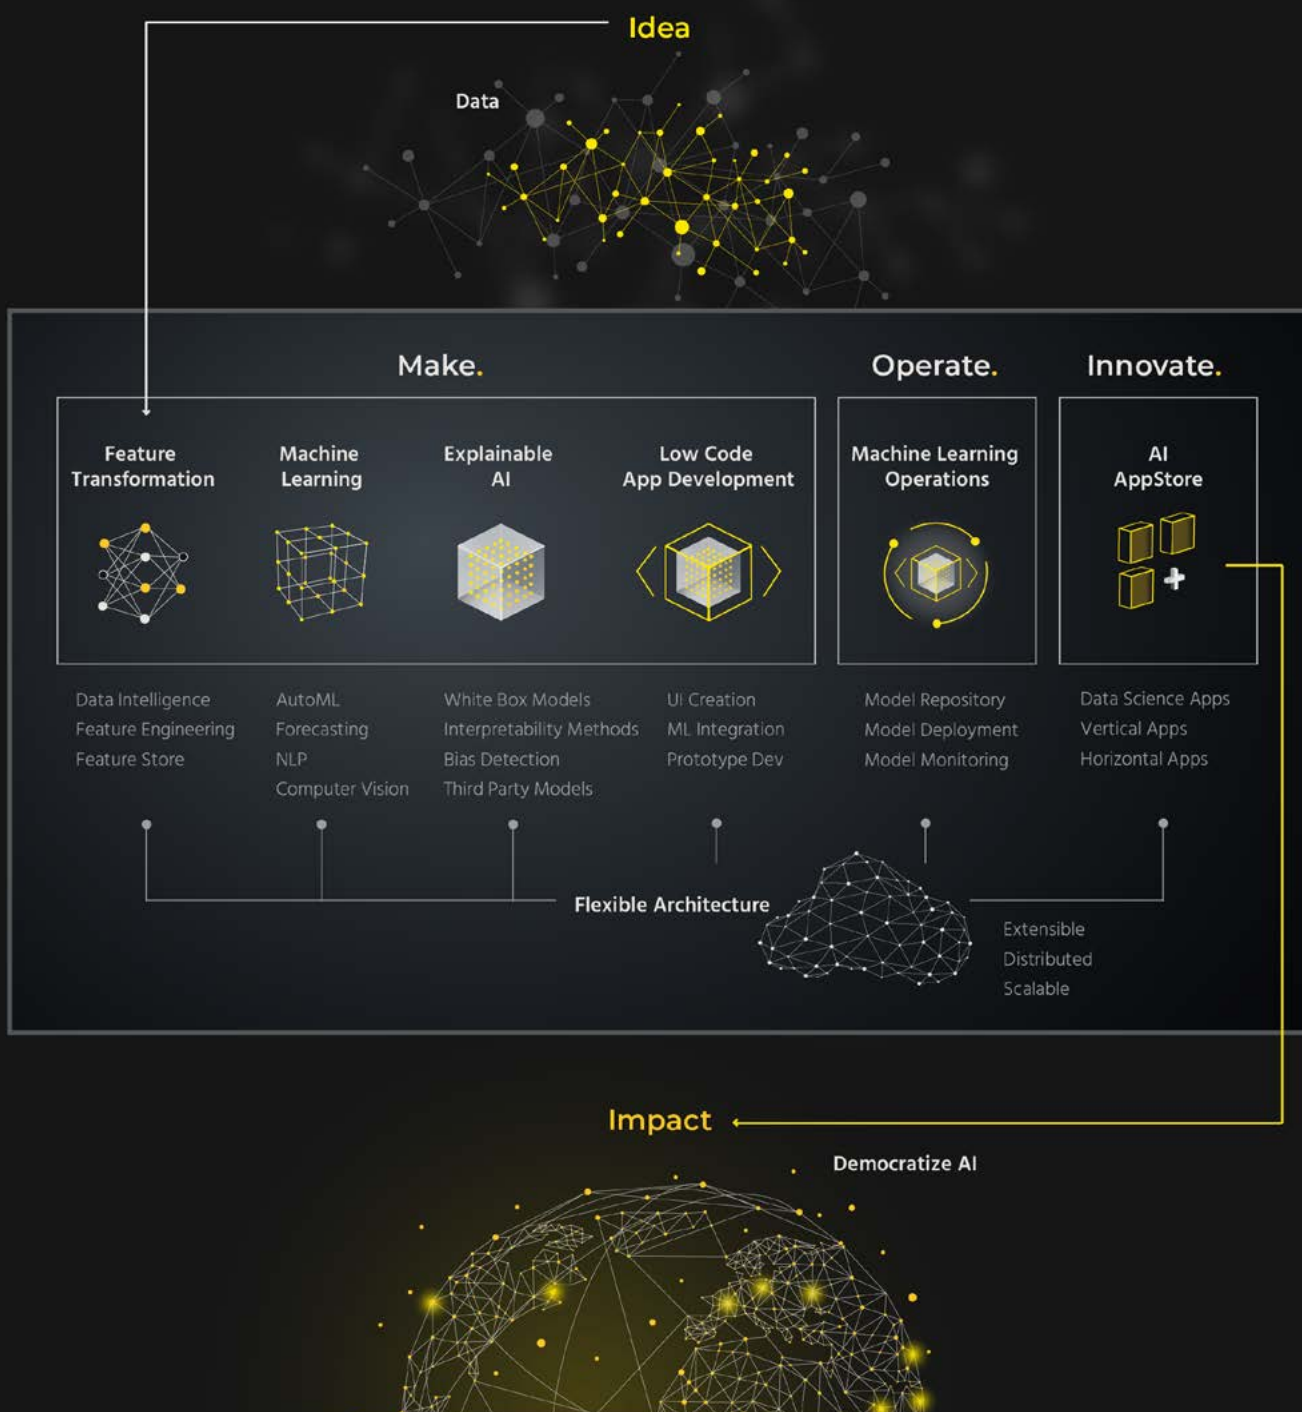
DATASHEET
2022




Key Features of the H2O AI Cloud

The H2O AI Cloud solves complex business problems and accelerates the discovery of new ideas with results you can understand and trust. Our comprehensive automated machine learning (autoML) capabilities transform how AI is created and consumed. We have built AI to do AI, making it easier and faster to use, while still maintaining expert levels of accuracy, speed and transparency.

We deliver on our mission to democratize AI by moving people from idea to impact with the confidence to make, operate and innovate with AI on the H2O AI Cloud.




Agility and transparency around the creation and use of AI solutions ensure a cycle of continuous learning and innovation. Enhanced cross-team collaboration improves the overall quality of results, as well as the effectiveness of responses to evolving circumstances and new insights.



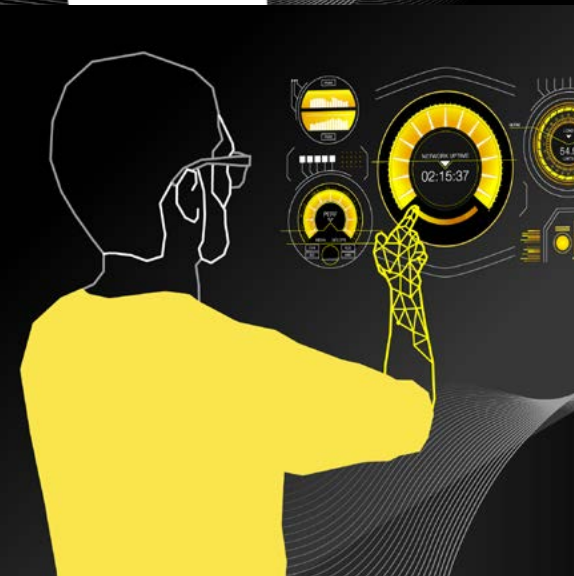
Data Scientists

Advanced autoML capabilities encompass the entire data science lifecycle, optimizing workflows to increase both the quantity and quality of data science projects delivered to your business stakeholders. Robust features for machine learning interpretability and operations give you seamless control over the entire explainability pipeline, resulting in higher accuracy, increased transparency and faster adaptations to changing circumstances.



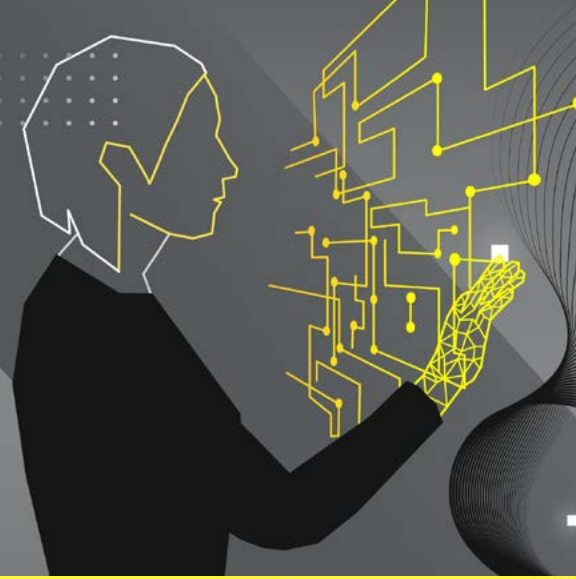
Developers

Automatic conversion of deployment artifacts from Python/R to Java/C++ simplifies the handoff of projects between data science teams and developers, and provides seamless integration of AI capabilities into existing solutions. Our open-source, low code development framework provides flexibility to work with all major libraries and frameworks and accelerates the creation of new AI applications.



Machine Learning Engineers

Guarantee that deployed models are operating as intended with comprehensive model monitoring features designed to support both H2O.ai and third party models. The H2O AI Cloud's comprehensive model management capabilities are accessible through a visual interface or API.





DevOps and IT Professionals

The H2O AI Cloud provides high performance computing with full NVIDIA RAPIDS integration. Easily scale workloads with support for the unprecedented compute and network acceleration of Ampere-based NVIDIA GPUs and latest CUDA runtime. Monitor system usage and provide metrics for resource monitoring and autoscaling of multi-node clusters through our platform performance API.



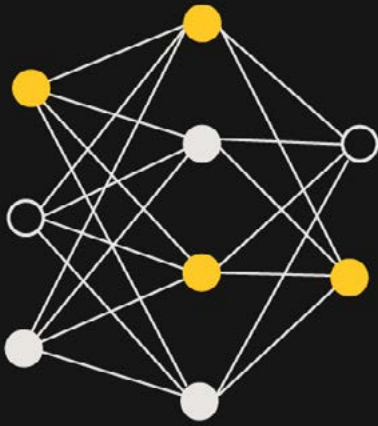
Business Users

The AI AppStore makes it easy and intuitive for business users to find and access AI applications custom built to meet your unique business needs. These solutions are powered by robust machine learning models and provide real-time insights for data driven decisions.



Make.

H2O.ai built AI to do AI. With industry leading automated machine learning (autoML), the H2O AI Cloud gives users more accuracy, speed and transparency throughout the entire machine learning lifecycle, including in the development and deployment of AI applications.



Feature Transformation

Automatically visualize and address data quality issues with advanced feature engineering that transforms your data into an optimal modeling dataset.



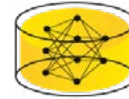
Data Intelligence

- **Data visualization** Visually display interesting statistical properties of your dataset and expose unexpected data quality issues like outliers, correlations or missing values.
- **Automatic data insights** Visual and text descriptions for automatically detected trends and insights including topics in text, correlations, and outliers.
- **Pre-processing transformers** Automatically include custom data preparation as part of your final deployed machine learning pipeline.
- **Dataset splitting** Save time and improve validation with a variety of built-in splitting techniques, including splitting randomly, by time, with stratification and with full customization via live code.
- **Missing value handling** Produce higher accuracy and better generalization with end to end support for missing values in all parts of the machine learning pipeline.
- **Outlier detection** Expose issues or irregularities in data with better accuracy delivered through various proprietary algorithms.



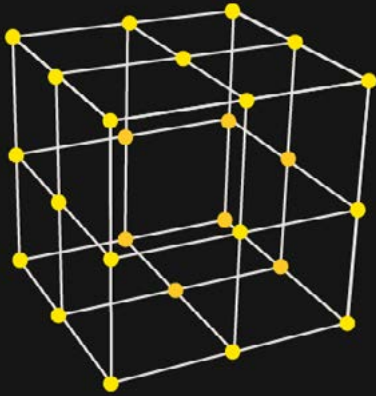
Feature Engineering

- **Automated feature engineering** Increase accuracy and ROI with our proprietary feature engineering that automatically extracts non-trivial statistical information from your data.
- **Feature encoding** Convert mixed data types (numeric, categorical, text, image, date/time, etc.) in a single dataset for use by machine learning algorithms.
- **Feature transformation recipes** Apply your domain knowledge to refine automated feature engineering outputs with fully customizable Python recipes.
- **Per-Feature Controls** Disable feature engineering and feature selection for certain columns in your dataset, and pass them as-is to the model to satisfy your compliance requirements.
- **Automated validation and cross validation** Improve accuracy, robustness and generalization with a multitude of proprietary validation techniques, statistical methods and moving windows.



Feature Store

- **Data pipeline and integrations** Consistent and unified storage for features and feature metadata allow data scientists to have access to the latest, best features.
- **Categorization and search** Metadata on the available feature sets allows users to find the features they need for each dataset. Projects and tags allow users to identify features which are commonly used together to improve their model.
- **Governance and access management** Versioned features ensure that a model trained on a specific combination of features will always lead to the same results.



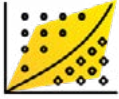
Machine Learning

Quickly create and test highly accurate and robust models with state-of-the-art automated machine learning that spans the entire data science lifecycle and can process a variety of data types within a single dataset.



Automated Machine Learning (autoML) AutoML is pervasive across the entire H2O AI Cloud. Powering everything from feature transformation to model selection, monitoring and deployment, robust autoML capabilities are the engine behind our ability to deliver AI that does AI.

- **Automated feature selection (dimensionality reduction)** Reduce model complexity, produce faster inference time and better model interpretability with a multitude of proprietary feature selection techniques that automatically select the most predictive features for your dataset.
- **Automatic Feature Engineering** Increase accuracy and ROI with our proprietary grandmaster-level feature engineering that automatically extracts predictive statistical information from your data for highest accuracy and for gaining actionable insights into the causal nature of the data.
- **Hyperparameter Autotuning** Increase accuracy, ROI and time savings with optimization across all components of the machine learning modeling pipeline delivered through a mix of our proprietary, genetic algorithm, Monte Carlo, Particle Swarm and Bayesian methods.
- **Champion/Challenger model selection** Speed up testing and validation with autoML that finds the best combination of features and models and automatic selection of the best machine learning model to fit your dataset.
- **Model ensembling** Multiple levels of both fully automatic and easily customizable ensembling to increase accuracy and ROI.
- **Automatic label assignment** Reduce error rates and save time with automatic labeling that predicts the class for every scored record, in addition to returning the per-class probabilities.
- **Automated model documentation** Provide insight into the machine learning process with automatically generated documentation that describes the experiment process, model tuning results, variable importance, model importance, model performance and detailed settings for reproducibility.
- **Machine learning interpretability (MLI)** AutoML powers a robust MLI toolkit to include explanations, visualizations and customizations.
- **Model validation** Assess model robustness and mitigate risks in production by obtaining a holistic view of the models and preventing failures on new data.
- **Unsupervised automatic machine learning** Immediately get new insights on your unlabeled data with unsupervised techniques such as clustering to automatically group topics, outlier detection to identify irregularities in your data, and dimensionality reduction to reduce model overfitting and complexity.
- **Wide dataset handling** Handle datasets with millions of columns with GPU acceleration.
- **Imbalanced dataset handling** Improve the accuracy in imbalanced use cases with access to special, proprietary algorithms which emphasize accuracy of rare classes over the more frequent but less valuable classes.
- **Model leaderboards** Speed up model development, model selection and model validation with automatic leaderboards comparing best-of-breed autoML modeling approaches.



Time Series Forecasting Easily fit and solve forecasting problems with unique feature engineering and autoML capabilities specifically designed to handle time series data.

- **Diagnostics** Automatically show model performance across groups and how predictions were generated with the ability to build forecasts across many categories (individuals skus, product hierarchy, etc.).
- **Leaderboard for forecasting** Save time getting an optimized forecasting model with a new leaderboard mode specific to time series experiments. Automatically design and run multiple experiments with varying amounts of pairs (train-test gap, forecast horizon, etc.) to help with model selection.
- **Time series machine learning interpretability** Time series models work with various MLI explainers, including Sensitivity Analysis, Disparate Impact Analysis, Partial Dependence/Individual Conditional Expectation, Naïve and Kernel Shapley, Surrogate Models and all feature importance techniques.



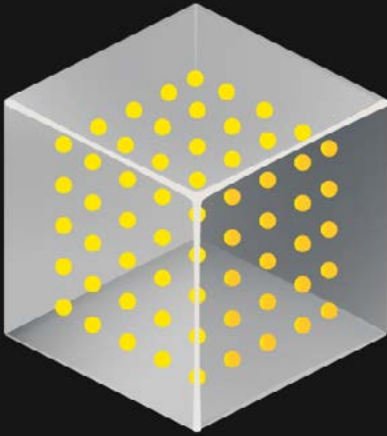
Natural Language Processing (NLP) Extract insights from unstructured text data to discover trends, create more accurate and relevant information retrieval and create personalized recommendations.

- **Text preprocessing** Remove noise from text and improve downstream NLP tasks by creating a pipeline of preprocessing techniques like stop word removal, stemming/lemmatization, emoji extraction and more.
- **Text classification** Tag and categorize similar terms and concepts with a range of supported algorithms from simplistic TFIDF to state-of-the-art BERT transformers.
- **Topic clustering** Gain insights for emerging trends and early detection of rising problems by automatically extracting relationships in text, with the option to fine tune clusters with human feedback to improve model accuracy.
- **Named Entity Recognition (NER)** Categorize terms into common groupings such as people, places, organizations, dates, currency and more.
- **Sentiment analysis** Detect trends in perspectives regarding specific topics such as products and services by identifying positive, negative and neutral connotations.
- **NLP machine learning interpretability** Understand how token importance varies between classes with newly added multinomial support. Understand the impact of tokens on outcomes with LOCO (leave one covariate out) 2.0 and Vectorizer + Linear Model (VLM). Calculate the average outcome of a model when a text token is included versus not included with Partial Dependence for Text Tokens. Select between TF-IDF or the newly added Vectorizer + Linear Model (VLM) when it comes to generating tokens for Surrogate models.



Computer Vision Use image data for modeling with the ability to combine additional data types with your image data to include text, tabular and audio data.

- **Image processing** Recognize objects in images and categorize them with high accuracy using all the latest convolutional neural networks (CNNs) out of the box with GPU accelerated training.
- **Video processing** Input videos are automatically transformed into a set of frames, which are then handled using image processing capabilities.
- **Audio processing** Input audio data is automatically transformed into MEL-spectrograms, which are then handled using image processing capabilities.
- **Explanations for computer vision models** Understand what part of an image or video led to a specific prediction using techniques such as Grad-CAM.



Explainable AI

Easily understand the ‘why’ behind model predictions to build better models and provide explanations of model output at a global level (across a set of predictions) or at a local level (for an individual prediction).



White Box Models

- **Generalized Linear Models (GLM)** GLMs are an extension of traditional linear models. They are highly explainable models with the flexibility of the model structure unifying the typical regression methods (such as linear regression and logistic regression for binary classification).
- **Generalized Additive Models (GAM)** A GAM is a Generalized Linear Model (GLM) in which the linear predictor depends on predictor variables and smooth functions of predictor variables.
- **Generalized Additive Models with two-way interaction terms (GA2M)** GA2M is an extension of GAM which selects the most important interactions between features and includes functions of those pairs of features in the model.
- **Explainable neural networks (XNN)** These neural networks consist of numerous subnetworks, each of which learns an interpretable function of the original features.
- **Rulefit** This algorithm fits a tree ensemble to the data, builds a rule ensemble by traversing each tree, evaluates the rules on the data to build a rule feature set, and fits a sparse linear model (LASSO) to the rule feature set joined with the original feature set. Available in Driverless AI by default.
- **Skopes Rules** This algorithm learns a simple set of rules for performing classification.
- **Decision tree and linear model combination** This model segments the data using a decision tree and then fits a GLM to the data at each leaf node.



Interpretability Methods

- **Partial Dependence Plot (PDP)** Plot that shows how a column affects predictions at a global level, with the ability for users to explore columns and how they affect predictions.
- **Feature importance** Calculate which features are important for the model’s decision making, both naive and with transformed features.
- **Shapley reason codes** Provide model explainability at a record level for non-linear models for global and individual records.
- **Surrogate Decision Trees** Identify the driving factors of a complex model’s predictions in a very simple, visual and straightforward way.
- **Leave One Covariate Out (LOCO)** Identify features that are important to the Surrogate Random Forest predictions from an aggregated or row level view.
- **Individual Conditional Expectation (ICE)** Plot that shows how a column affects predictions at an individual level, with the ability to drill down to any row of choice and compare/contrast with average partial dependence.
- **k-LIME reason codes** Generate novel reason codes at a record level, subsets of the dataset or at an aggregated level for the entire dataset.





Bias Detection

- **Disparate Impact Analysis** Identify areas in your data where a model shows bias across various metrics with a dashboard that shows disparity between groups in the dataset.
- **Sensitivity analysis** Provide model explanation at a very granular level without much overhead with a tool that allows you to score a trained model on a single row, multiple rows or an entire dataset and compare the model's new outcome to the predicted outcome on the original dataset.




Third Party Models

- **Explanations for time series models** Time series models work with various MLI explainers, including Sensitivity Analysis, Disparate Impact Analysis, Partial Dependence/Individual Conditional Expectation, Naïve and Kernel Shapley, Surrogate Models and all feature importance techniques.
- **Explanations for NLP models** Understand how token importance varies between classes with newly added multinomial support. Understand the impact of tokens on outcomes with LOCO (leave one covariate out) 2.0 and Vectorizer + Linear Model (VLM). Calculate the average outcome of a model when a text token is included versus not included with Partial Dependence for Text Tokens. Select between TF-IDF or the newly added Vectorizer + Linear Model (VLM) when it comes to generating tokens for Surrogate models.
- **Explanations for computer vision models** Understand what part of an image or video led to a specific prediction using techniques such as Grad-CAM.
- **External model interpretability** Gain insight into models that were built outside of the H2O AI Cloud with model agnostic explainers and feature importance calculations.
- **Bring your own recipes** Seamless control over the entire explainability pipeline is enabled with API access to customize all explainability components with Python.



Low Code Application Development Framework

Rapidly build AI prototypes and applications with a low code framework (Python/R) that makes it easy to deliver innovative solutions by seamlessly integrating backend machine learning capabilities and front end user experiences.




- **Application Development**
Reduce the time it takes to build applications with access to best practices and a library of prototypes spanning multiple industries, and simplify management of solutions with the ability to see changes to the UI in real time as you develop.



- **Low Code User Interface (UI) Creation** Easily develop consumer facing, interactive applications with a Python/R framework that supports integrated UI and AI development.




- **Machine learning integration**
API access to best-in-class data science capabilities makes it easy to integrate machine learning functionality into new or existing solutions.



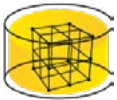
Operate.

H2O.ai provides a comprehensive suite of capabilities surrounding machine learning operations that support data scientists and machine learning engineers in the deployment, management and monitoring of their models in production. Additionally, the H2O AI Cloud provides an incredibly flexible architecture with distributed processing, optimized compute efficiency and the ability to deploy in the environment of your choice. Customization is well supported with easy integration of your own transformers, recipes and models.



Machine Learning Operations

Automatically monitor models in real-time and set custom thresholds to receive alerts on prediction accuracy and data drift and guarantee deployed models are operating as intended.



Model Repository

- **Model management** Create a central place to host and manage all experiments and its associated artifacts, across the entire organization. Register experiments as models, including both auto-generated and custom metadata to have a centralized view of all models.
- **Model versioning** Register experiments as new model versions and maintain a transparent view of all deployed versions.
- **3rd party model support** Manage models trained on any 3rd party framework, including scikit-learn, PyTorch, TensorFlow, XGBoost, LightGBM and more, just like your native Driverless AI models.




Model Deployment


- **Target deployments** Build once and deploy to any scoring environment.
- **Deployment modes** Deploy models within the production environment in different modes, including multi-variant (A/B), champion/challenger and canary.
- **Scoring** Models can be scored in real-time (hosted RESTful endpoint), in batch (supported source and target datastores), asynchronously or as streaming data.



Model Monitoring

- **Data and Concept Drift** Maintain model oversight and know if your models are scoring on data they were not meant to or trained on.
- **Feature importance** Receive local explanations on which features are contributing greatest/least to the prediction value, along with the scoring result.
- **Leaderboard** View and sort experiments and models by key evaluation metrics.
- **Alerts** Receive alerts and notifications for all monitored metrics with the ability to set custom thresholds.





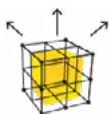
Flexible Architecture

The H2O AI Cloud is environment agnostic so any company, regardless of their existing infrastructure, can incorporate H2O.ai technologies into their machine learning pipelines.



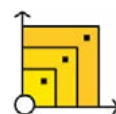
Extensible

- **Platform agnostic** Support for Google Cloud Platform, Amazon Web Services, Microsoft Azure, or on premises.
- **Custom recipe architecture** Benefit from the latest versions of Python, RAPIDS/CUML, PyTorch, TensorFlow, H2O, XGBoost, LightGBM, datatable, datatable, sklearn, pandas, and many more packages. And gain full control over them and any other Python package with our built-in custom recipe architecture.
- **Flexible model support** Train and deploy any H2O.ai or third party model and customize it with Python.
- **Multiple programming languages** Covers the majority of the data science user base with clients for Python, R and Java.




Distributed

- **Multi-node training** Scalable, distributed machine learning backends can handle any data size by scaling out to multiple worker nodes.
- **Multi-CPU/GPU training** Train models faster across multiple CPUs/ GPUs.
- **Kubernetes-based deployment** Simplifies infrastructure scalability and maintenance by automating cloud resource allocations.



Scalable

- **NVIDIA RAPIDS integration** High performance computing delivered through full NVIDIA RAPIDS integration.
- **Ampere-based GPUs** Easily scale workloads with support for the unprecedented compute and network acceleration of Ampere-based NVIDIA GPUs through the use of the latest CUDA runtime.
- **Platform performance API** Visualize and monitor system usage with a publicly available API that provides platform metrics for resource monitoring and autoscaling of H2O AI Cloud multinode clusters.




Innovate.

Democratizing AI across entire organizations requires the merging of powerful machine learning models with intuitive experiences that users of all technical abilities can consume. The H2O AI Cloud provides an AI AppStore designed to simplify the delivery and consumption of complex solutions, meaning more people can access and participate in innovation efforts.

AI AppStore

The H2O AI Cloud supports rapid prototyping and solution development while also fostering collaboration between technical teams and business users. With comprehensive machine learning capabilities, a robust explainable AI toolkit, a low code application development framework and integrated machine learning operations, the H2O AI AppStore is the catalyst needed to move you from big ideas to tangible impact.



Cross-Industry Data Science Apps

- **AutoInsights** Automated interactive, engaging, and actionable insights summarized as easy-to-understand natural language narratives.
 - Explore new datasets to find important business findings
 - Find insights without writing code or a deep analytics background such as unusual data points, topics in text data, and correlations between features
 - Get interactive visualizations and narratives for efficient storytelling
- **Model Validation** Analyze the robustness and stability of your machine learning pipelines to reveal any weaknesses or vulnerabilities in data and models.
 - Understand if your data is changing overtime using techniques like drift and adversarial similarity
 - Use backtesting to understand how a Machine Learning model would have performed historically at different points in time
 - Visualize data and model validation techniques to better understand if a model can be trusted to make business critical decisions
- **Model Monitoring** Monitor your production models for feature drift to understand if they can still be safely used in production.
 - Integration with H2O MLOps to get data science insights on deployed models
 - Understand drift for external models, such as those deployed in SQL database
 - Use custom thresholds to integrate your business knowledge of when drift is a concern or not

Financial Service Apps

- **LIBOR Clause Detection** A comprehensive solution for processing financial contracts and documents using automated NLP pipelines to identify sentences of interest that mentions LIBOR, insights about the context of LIBOR sentences, and a BERT based sentence clustering algorithm to segment different sentences in different categories.
- **Explainable Anomaly Detection** Identify which data points are highly unusual and what about them is different from the rest of the data; this can be helpful for creating new rules to find fraudulent activities.
- **Credit Card Risk** A prototype of how to build your own machine learning model and app front end for predicting whether or not a customer will pay back a credit card on time.

Healthcare Apps

- **Covid-19 Forecast** The application uses the latest Covid-19 data to provide geography-based forecasts of confirmed cases and deaths for up to one month in the future.
- **Covid-19 Vaccine NLP** NLP algorithms that identify sub-topics and the evolving sentiment within opinions expressed in social media about vaccinations.
- **Covid-19 Hospitalization Predictor Name** Plan for future hospital admissions (ED, ICU, other inpatient) based on the latest information Covid-19 prevalence and vaccination rates at the hospital's area.
- **Medical Appointments No-Show** Improve patient care and patient engagement using AI to obtain the likelihood of appointment no-show and gain insights on factors that play an important role when a patient does not honor a scheduled appointment.
- **Employee Churn Prediction** AI playing a central role as part of an overall plan to minimize employee attrition in a prototype solution that visualizes predictions of employee departure, forecast churn rates, and identify relevant factors contained in employee data.
- **Medicine Reviews** AI- and NLP-driven assessment of the perceived effectiveness of medicines based on uploaded or individually entered patient reviews.
- **Chest Abnormality Detection** Computer vision prototype for the detection and identification of different types of anomalies from X-ray images, aimed at both data scientists and clinicians.
- **Provider Claims Fraud** A solution for payers powered by AI that evaluates the trustworthiness of providers, giving explanations on the reasons that drive the evaluation.

Customer Service Apps

- **Customer Churn Risk** An example of how to build your own machine learning model and app front end for real-time predictive analytics on identifying if a customer is likely to leave for a competitor. Customer Churn is repeatedly marked as a top use case for almost all industries. Knowing which customers are most likely to leave us, and more importantly, why, can give our support staff the right tools to help reduce customer loss.
 - Explore Customer Data
 - Allows the end-user to look at a few rows of training data
 - Shows historic trends in the data using a third-party plotting library
 - View Callers
 - Sends simulated new data to a model REST endpoint every 3 seconds
 - Shows the end-user the latest predictions
 - Shows distribution of predictions by state and overall
 - View the App Code
 - Explore the Python code used to make this app
- **Market Basket Analysis** Build and explore market basket algorithms to create product recommendations for your customers. Market Basket Analysis helps in understanding the relationship between various products in a catalog and generating recommendations using historic transactional data. It is popularly used for scenarios like:
 - Recommending similar products to a particular product for better engagement
 - Recommending products relevant to an existing cart to add during checkout
 - Recommending combination of products that can be sold together as an upselling offer
 - Recommending substitute products of a particular product to manage inventory
- **Explainable Hotel Ratings**

Use model interpretability to understand what about a text in a review makes it look positive or negative in sentiment. This Hotel Review Analysis demo allows a business user to explore satisfaction predictions made by a machine learning model. The user can compare the key words in the average prediction to various subsets of the data to understand how different guests enjoyed different types of hotels.

Business Operations Apps

- **Predictive Maintenance** Get real-time information from sensor data to understand which parts are likely to fail, when, and why. The Fourth Industrial Revolution is the next iteration of automation of traditional manufacturing and industrial practices using modern smart technology. This includes the concepts of interconnectivity, digitisation, and automation. In this context, prognostics and health management in the manufacturing industry, also known as Predictive Maintenance, aim at minimizing maintenance costs by the assessment, prognosis, diagnosis, and health management of industrial equipment using intelligent and automated processes. This application provides a prognosis and diagnosis of manufacturing equipment using real-time sensor data.
- **Demand Sensing** Combine traditional demand forecasting with external factors such as weather or consumer buying behavior to improve the supply chain.
 - Demand Sensing is a methodology and technology that leverages new mathematical techniques and current real-time information to create an accurate near-future forecast (whatever may be of hours or days depending on how the dynamic supply chain is) of demand, based on current realities of the supply chain
 - Demand Sensing is fundamentally different in forecasting as it uses a much broader range of demand signals including current data from the supply chain and various mathematical models to create a more accurate forecast that responds to a real-world event such as Market Shifts, Weather Changes, Natural Disasters, Consumer Buying behaviours etc
- Demand sensing uses huge amounts of data and recognizes patterns, so supply chains have actionable signals and can make accurate decisions. It helps companies understand consumer behaviour and variables by synthesizing big data in real-time
- **Employee Churn** Predict the risk of employee departure and identify relevant factors. This app demonstrates the use of machine learning as part of an overall plan to minimize employee attrition. Users can view predictions of employee departure, forecast churn rates, and identify relevant factors contained in employee data.

Try the H2O AI Cloud Today.

Visit www.h2o.ai/free



H2O.ai