


H2O.ai

Building a Data Driven Government

with Responsible AI



Building a Data Driven Government with Responsible AI

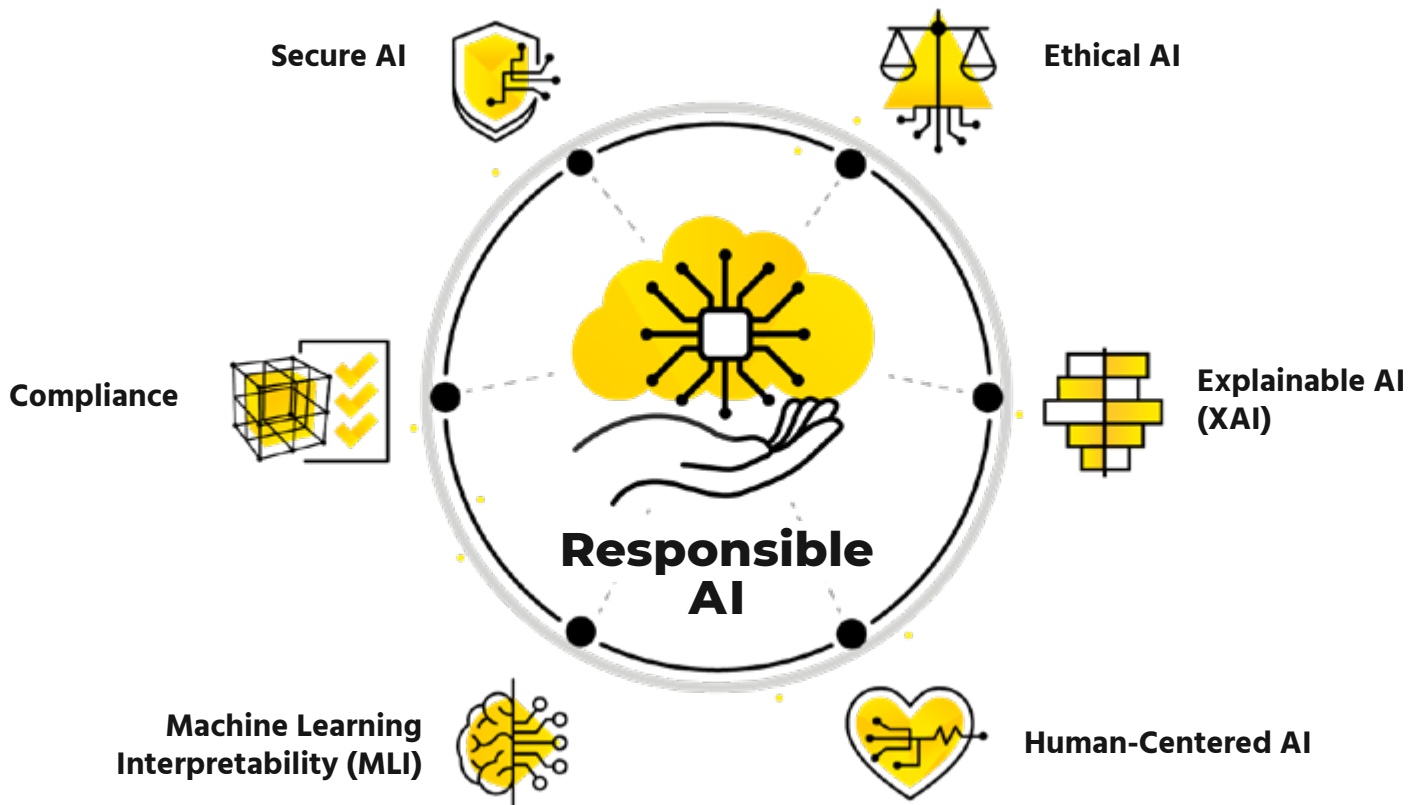


Artificial Intelligence (AI) enables agencies to make better decisions, streamline operations, and reduce risk. AI can have a significant impact across departments in areas such as fraud, waste and abuse, supply chain and logistics, cybersecurity and mission specific objectives. The technology is powerful in its ability to process massive amounts of data, deliver insights and inform key decisions

But, like humans, AI is known to make errors and exhibit bias. Because of the types of information AI processes for government organizations, it's important to understand and track how algorithmic output influences human decision making. This is where the requirement for Responsible AI comes into play.

Responsible AI incorporates best practices for people, processes, and technology. Beyond regulatory compliance, agencies operationalizing AI solutions have legitimate interests in understanding their models and trusting that those models will perform as expected.

The requirement to create and manage Responsible AI programs is complex with many different technical factors playing a key role.



- **Ethical AI:** Sociological fairness in machine learning predictions (i.e., whether one category of person is being weighted unequally).



- **Explainable AI (XAI):** The ability to explain a model after it has been developed.



- **Human-Centered AI:** User-centric design experiences that support human interactions with AI and machine learning systems.



- **Machine Learning Interpretability (MLI):** Transparent model architectures and increasing how intuitive and understandable machine learning models can be.



- **Compliance:** Ensuring AI systems meet relevant regulatory requirements (ie: GDPR, CCPA, FCRA, ECOA, etc.).



- **Secure AI:** Debugging and deploying machine learning models with similar counter-measures against insider and cyber threats as seen in traditional software.

AI is critical to the evolution and transformation of government services. Results from the technology must be explained, understood, and trusted. Combining the technical requirements to deliver those expectations from AI with appropriate human skill sets and procedural guidelines is the basis for establishing a Responsible AI culture.

Transparency and Governance


Explainable AI helps data scientists easily understand the ‘why’ behind model predictions. This allows them 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).

Explainable AI with H2O.ai



White Box Models

Competition winning modeling methods simultaneously enable model transparency and robust post-hoc interpretability methods for explaining and understanding your models. Additionally, H2O.ai’s autoML provides virtually endless constraints and parameter controls to ensure your model is as simple or as complex as you need it to be.




Interpretability Methods

The H2O AI Cloud comes with one of the most robust and dynamic explainable AI toolkits that enables customers to leverage dozens of post-hoc explanation methods and understand why your machine learning model came to the decisions that it did.




Bias Detection:

One of the largest concerns for agencies adopting machine learning models is the possibility of having the models perpetuate bias from the given dataset. Built from our leading research on fairness in AI, we provide multiple methods to identify, explain and debug bias in machine learning models.




Model performance can easily be captured and recorded with automatic model documentation. This provides 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 operations allow agencies to automatically monitor models in real-time and set custom thresholds to receive alerts on prediction accuracy and data drift. This helps guarantee deployed models are operating as intended.




Machine Learning Operations with H2O.ai



Model Repository:

Create a central place to host and manage all experiments. Maintain a view of all deployed versions with complete, integrated model management capabilities that are accessible by both an easy-to-use web interface and an API. You can also manage models trained on any 3rd party framework.



Model Deployment:

You can build once and deploy to any scoring environment with target deployments. Deploy in different modes, including multi-variant (A/B), champion, challenger and canary. Models can be scored in real-time, in batch, asynchronously or with streaming data.



Model Monitoring:

Maintain model oversight and know when data drift occurs. Feature importance delivers local explanations as to which features are contributing the most or least to prediction values. You can set custom thresholds to receive alerts and notifications for all monitored metrics.

Operate Models.



Data Scientists

The H2O AI Cloud makes it easy for data scientists to quickly and seamlessly hand over their models to machine learning engineers. This allows data scientists to focus on discovering new insights in additional data sources, increasing the accuracy and performance of machine learning models and driving further systematic innovation efforts.

Operate Deployment.



Machine Learning Engineers

H2O.ai makes deployment easy with real-time, customizable monitoring and alert systems. The H2O AI Cloud offers a multitude of capabilities for backtesting, challenging and validating your models over time. Easily incorporate multiple ongoing Responsible AI and fairness metrics into your ongoing monitoring programs.

Operate Access.



DevOps/IT Professionals

The H2O AI Cloud simplifies the provisioning of software for all parts of the data science lifecycle, from data access all the way through to AI application deployment. Self-service is enabled through a centralized deployment environment. Resource monitoring and cost controls allow IT professionals to optimally balance cost and performance.


Transformational Change



As governments look to transition their AI projects beyond small experiments, the ability to scale AI operations depends heavily on the ability to provide transparency and oversight across every stage of making, operating and innovating with AI. There is an inherent goal of connecting users and stakeholders across teams to ensure machine learning models are accurately built, easily managed and ultimately delivering mission value.

The way to operationalize and deliver this value is defined in H2O.ai's approach to strategic transformation.


Strategic Transformation Process





Define

The merging of leadership, technology and collaboration begins with your agency's larger purpose around transformation and the tangible objectives you want to explore in support of that purpose. We partner with you to solve complex business problems and accelerate the discovery of new ideas, including partnering with you to work on methodology, application, deployment or net new technology development.



Experiment

Our autoML supports rapid development and the ability to adapt and make adjustments quickly. Use your own Python scripts or choose from the many accelerators we provide for models, transformers and scorers. Champion model selection, feature engineering, hyperparameter autotuning, model validation and model documentation are all automated for you. And we have a massive open source community to provide support and guidance as well as some of the world's top data scientists, Kaggle Grandmasters, to help in this phase as well.



Integrate

Merge your AI models and applications into existing operations with our platform agnostic, distributed, extensible, and scalable architecture. Comprehensive machine learning operations support the deployment, monitoring and management of machine learning models in production. Customization is well supported with easy integration of your own transformers, recipes and models to help in this phase as well.



Enable

We lower the adoption barrier with the option to deliver machine learning capabilities through applications that can be accessed through the AI AppStore. These applications, and the AI AppStore itself, can easily be customized based on end user needs. The ability for all departments of business to access AI solutions from one central location streamlines expectations and experiences, while still offering the flexibility to meet unique needs across the company.



Evolve

With AI and machine learning capabilities built into the core of agency operations, you create a dynamic organization capable of learning and adapting to real world needs with relevance and agility. You can develop new or improve existing plans and evolve your intelligence discovery and analysis process with new insights to make smarter decisions.

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. H2O.ai gives you the freedom to innovate at scale with results you can trust.



Request a Demo of the H2O AI Cloud

[Request Demo](#)

H2O.ai