



## NVIDIA AND H2O GPU-ACCELERATED MACHINE LEARNING

Machine learning and deep learning are transforming the way enterprises make decisions. This is especially true in industries like financial services, insurance, healthcare, and telecom where companies are quickly replacing rules-based systems with AI-accelerated applications. However, a big part of machine learning involves detecting complex interactions among features, learning low-level features from minimally processed raw data, working with high-cardinality class memberships, and working with unlabeled data much more efficiently.

### INTEGRATED SOLUTION

H2O's GPU-accelerated machine learning algorithms, powered by NVIDIA® DGX™ Systems and NVIDIA GPUs, allow enterprises to apply the data science techniques. These solutions are designed to solve your enterprise problems with speed, and at scale. Together, customers can experience seamless deployment with production-ready hardware and software and trust that all machine learning models can be interpreted transparently and with high accuracy.

### INDUSTRY CHALLENGES

#### Model Complexity

- > Complicated model preparation due to huge volume and various forms of data
- > Constant model re-training and re-evaluation due to ever changing nature of customer data and structure, and the need to improve model prediction accuracy
- > Model comparison and collaboration need to take place across organizations, making it hard to collaborate with data scientists, software engineers, and DevOps

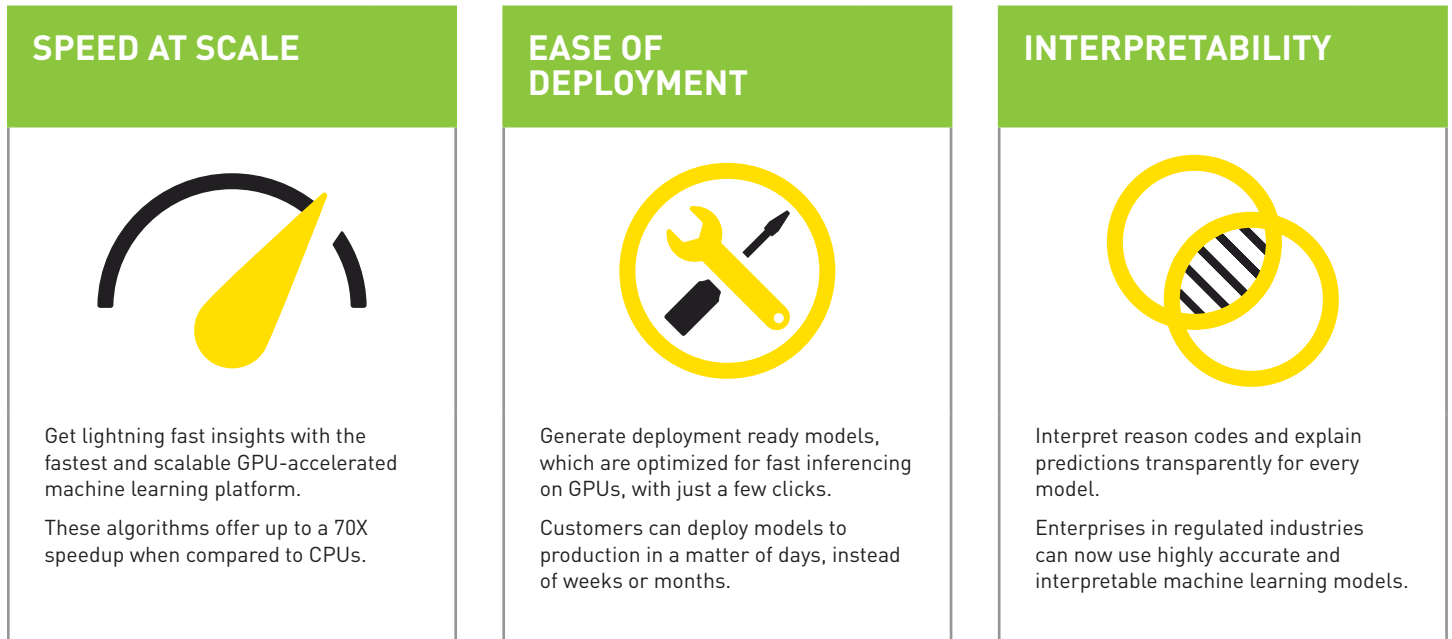
#### Slow Processes

- > Fast deployment of newly trained models needed to provide business value
- > Updating models is an offline batch process, and thus business cannot be agile in applying new models
- > Slow compute means the time to train is long and customers want results faster

#### Resource Constraints

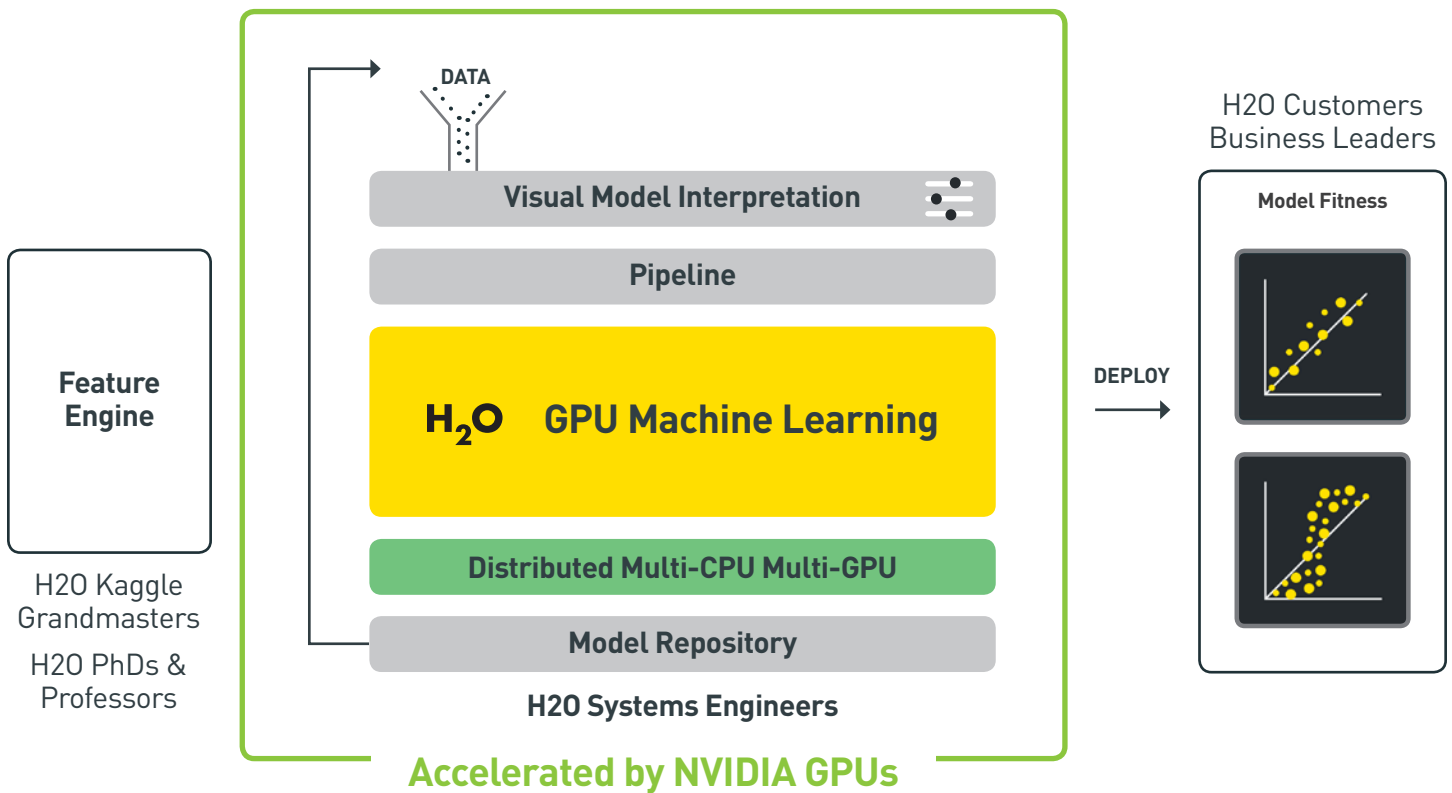
- > Most data science solutions are out of reach for teams without data science expertise
- > Endless waiting for dedicated resources (e.g. red tape and security concerns) means further delay in getting business insight

## Together, NVIDIA and H2O Deliver



## H2O Marketecture

### Driverless AI for the Digital Brain — Enabled by Fast Model Training



## Industry Insights

Customers everywhere are using massively parallel graphics processors to provide higher throughput for compute-intensive workloads and achieving significant performance gains without the hidden cost of scale-out architecture.

Finance	Healthcare
<p>Leverage H2O's GPU-accelerated machine learning algorithms to approve a purchase or a credit card application, to stop a withdrawal or recommend new financial products, or to detect payment fraud and money laundering. This is made possible with data like personal profile data, spending history, social media data, transaction data, asset information, and demographic data for a more holistic view, instead of just past credit scores.</p>	<p>Hospitals can use machine learning to calculate the probability of a rapid deterioration in patients who need unplanned transfer to the ICU, and to detect cancer and other fatal diseases early and accurately. With scans, patient tests, and other sensor information, hospitals can offer personalized healthcare recommendations from the wealth of data available.</p>
Insurance	Telco
<p>Insurance companies analyze myriad data sources (click streams, Customer Relationship Management systems, social media, and weblogs) from one holistic view. With machine learning algorithms, they are provided with better insurance recommendations, more cross-sell and upsell opportunities, and better classification of risk groups.</p>	<p>Telecom companies analyze customer usage data with GPU-accelerated machine learning algorithms to fix issues like dropped calls, connectivity issues, network latency and avoid expensive in-person technician visits, saving millions of dollars.</p>



“Future advancements in machine learning will unlock opportunities for us to create breakthrough consumer experiences in ways that we can't even imagine today,” said Adam Wenchel, VP of AI and Data Innovation at Capital One. “As users of the H2O.ai and NVIDIA platforms, we see GPU-acceleration of machine learning as a transformative development for the enterprise distributed machine learning community.”

## Recommended NVIDIA Hardware

NVIDIA datacenter GPUs are available in servers, DGX Systems, and cloud platform around the world. You can now get end-to-end accelerated Machine Learning solutions powered by NVIDIA GPUs with supporting software technologies and support from NVIDIA experts.

<b>TESLA</b> SERVERS IN EVERY SHAPE AND SIZE		<b>DGX SYSTEMS</b> AI TOOLS FOR INSTANT PRODUCTIVITY		<b>CLOUD EVERYWHERE</b>	
      				 	 

## Find Out More

NVIDIA GPUs for Machine Learning: Helping customers effectively analyze, visualize, and unleash the power of AI to transform their digital business into an AI enterprise.

Website: [www.nvidia.com/analytics](http://www.nvidia.com/analytics)

Partner Webpage: [www.nvidia.com/dgx-apps](http://www.nvidia.com/dgx-apps)

Twitter: [@NvidiaAI](https://twitter.com/NvidiaAI)

Blog: [blogs.nvidia.com](http://blogs.nvidia.com)

H2O.ai is the maker behind H2O, the leading open source machine learning platform for smarter applications and data products.

Website: [www.h2o.ai](http://www.h2o.ai)

Contact: [sales@h2o.ai](mailto:sales@h2o.ai)

Twitter: [@h2oai](https://twitter.com/h2oai)

Blog: [www.blog.h2o.ai](http://www.blog.h2o.ai)

