



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 (ML) algorithms and Driverless AI--an expert system designed to bring the power of a Kaggle Grandmaster in a box--allow data scientists and enterprises with limited resources to apply data science techniques simply and efficiently. Customers can experience seamless deployment with production-ready hardware and software and trust that all machine learning models can be interpreted transparently with high accuracy.

Powered by NVIDIA GPUs, H2O's GPU-accelerated algorithms provide up to 40X faster than traditional algorithms. This means that data scientists can use the algorithms to build models much faster and more accurately. Driverless AI is now containerized on NVIDIA® DGX™ Systems and NVIDIA GPU Cloud. This empowers business users with domain knowledge to understand the new features that influence predictions parameters, allow the data analysts to explore the data and identify trends simply, and enable data scientists to work on multiple projects simultaneously.

INDUSTRY CHALLENGES

Model Complexity

- > Huge volume and various forms of data mean complicated model preparation.
- > There is constant model re-training and re-evaluation due to the 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. It's hard to collaborate among 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

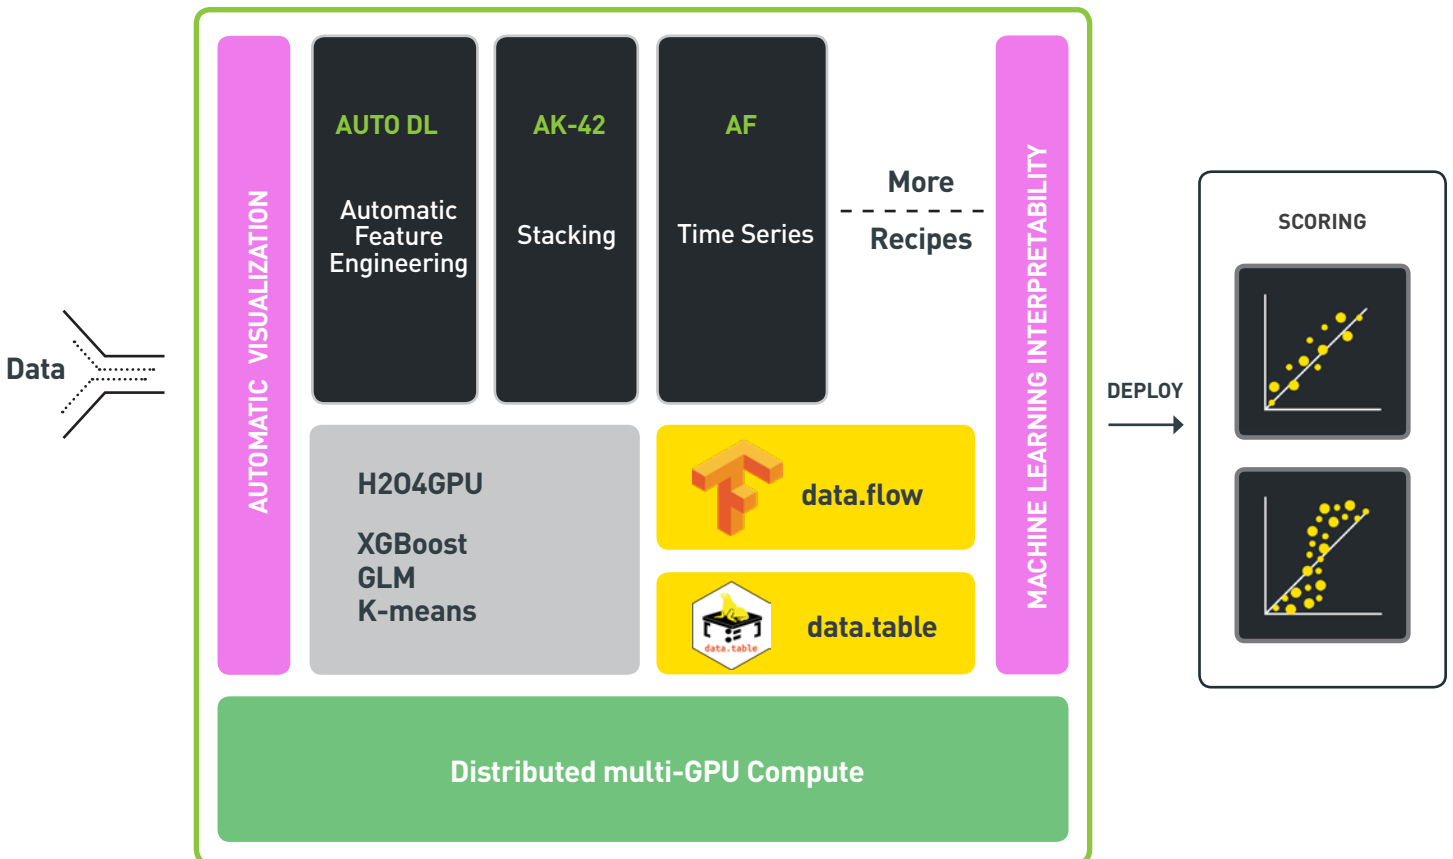
Resource Constraints

- > 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

SPEED AT SCALE	ACCURACY	INTERPRETABILITY
 <p>Discover lightning-fast insights with the fastest and most scalable GPU-accelerated machine learning platform.</p> <p>GPU acceleration to achieve up to 40X speedups vs CPUs</p> <ul style="list-style-type: none"> > GLM: More than 5X faster on GPUs > SVD: More than 5x faster on GPUs > XGBoost: Nearly 10X faster on GPUs > K-Means: More than 40X faster on GPUs 	 <p>Provide the accuracy of a Kaggle Grandmaster right out of the box.</p> <ul style="list-style-type: none"> > Driverless AI built a model ranked in the top 3% of the Kaggle leaderboard within an hour. > Driverless AI performs automatic feature engineering and machine learning out of the box at the level of an expert data scientist. 	 <p>Interpret reason codes and explain predictions transparently for every model.</p> <ul style="list-style-type: none"> > Four dynamic graphs are generated automatically: K-LIME, Variable Importance, Decision Tree Chart, and Partial Dependence Plot; each one helps explore the model output more closely.

H2O Driverless Architecture



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>Banks can use H2O's GPU-accelerated machine learning algorithms to approve a purchase or credit card application, stop a withdrawal or recommend new financial products, or detect payment fraud and money laundering. This is made possible with 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.

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

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

Contact: dgxanalytics@nvidia.com

Partner Webpage: <https://www.nvidia.com/en-us/ai-accelerated-analytics/partners/>

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

Blog: 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

Contact: sales@h2o.ai

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

Blog: www.blog.h2o.ai

