



INDUSTRY CHALLENGES

Credit Risk Scoring: Model Complexity

- > Banks and credit card companies use credit scores to evaluate potential risk when lending money or providing credit. But this type of data requires sophisticated data-visualization techniques.
- > AI can create individualized credit scores based on factors such as income, employment opportunity, and credit history. The accuracy of the machine learning models that generate these scores depends on comprehensive data transformations to create engineered features.
- > Moreover, the ever-changing nature of customer data requires improving model-prediction accuracy. AI adapts to new problems, like credit card churners and regulatory requirements, and accounts for these factors in credit decisions.

Fraud Detection: Slow Processes

- > Fraud is a growing problem in the banking industry. Accelerated compute and fast deployment of newly trained models is needed to tackle this problem and boost business value.
- > AI can quickly, accurately, and in near real time detect fraud with algorithms that identify outliers and analyze patterns. These models can even anticipate the likelihood of fraud, enabling investigators to prioritize their work on the biggest threats.

NVIDIA AND H2O: GPU-ACCELERATING THE FINANCIAL SERVICES INDUSTRY

Intelligence. Optimization. Innovation.

With artificial intelligence, industries like financial services are streamlining operations, deriving actionable insights from smart analytics, and creating new business models. By giving them the ability to draw on their deep troves of data, AI empowers businesses to find competitive market advantages.

But critical steps in machine learning workflows like feature engineering and model validation, tuning, selection, and deployment are complex and time-consuming. That's where H2O's AI enterprise platform makes a difference.

Integrated Solution

H2O's machine learning algorithms run on its award-winning automatic machine learning platform, Driverless AI, to seamlessly train and deploy machine learning models with production-ready hardware and software.

Powered by NVIDIA GPUs and CUDA-X™ AI, H2O's GPU-accelerated algorithms are up to 40X faster than traditional algorithms running on legacy hardware, resulting in faster and more accurate models. Driverless AI is containerized on NVIDIA® DGX™ Systems and available from NVIDIA NGC for supported NVIDIA GPU platforms. It gives business users robust interpretability of machine learning models to explain modeling results, allows data analysts to explore data and identify trends, and enables data scientists to achieve the highest predictive accuracy quickly.

Industry Challenges

With strong partnerships among leading financial institutions, including Wells Fargo, Citigroup, Capital One, PayPal, Discover, Dun & Bradstreet, and Equifax, H2O solves key industry challenges. On the left, view a snapshot of the top three challenges H2O's AI solution addresses and the opportunities it creates.

Customer Experience: Resource Constraints

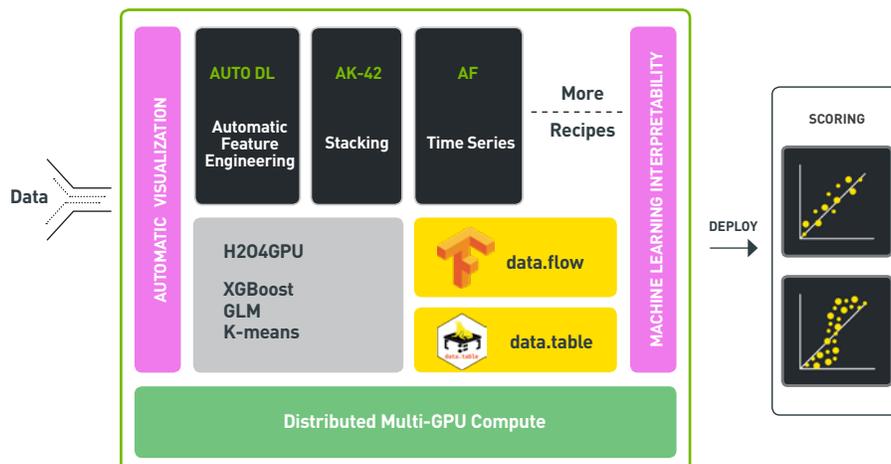
- > Most data science solutions are out of reach for teams without niche expertise, which makes it difficult to harness decades' worth of information and deliver the best customer experience.
- > Using smarter analytics by AI, financial firms are creating strategies and business models to tap into new revenue streams through a better understanding of their customers. As a result, product recommendations are personalized and in real time.

Together, NVIDIA and H2O Deliver

TIME	TALENT	TRUST
 <p>Achieve up to 40X speedups over CPUs with the fastest and most scalable GPU-accelerated machine learning platform.</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 expertise 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 percent 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>Use reason codes to explain predictions transparently for every record or transaction.</p> <ul style="list-style-type: none"> > Explain the behavior of models with charts that are generated automatically: K-LIME, Shapely, variable importance plot, decision tree chart, and partial dependence plot. Each one helps explore the model output more closely.

H2O Driverless Architecture in Action

“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, vice president 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 data center GPUs are available in servers, DGX Systems, and cloud platforms 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.

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



Find Out More

NVIDIA GPUs for machine learning are helping customers effectively analyze, visualize, and unleash the power of AI to transform their digital business into an AI enterprise.

Website: nvidia.com/finance

Email: finance@nvidia.com

Data Analytics Partners: nvidia.com/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.h2o.ai/blog/