Deep Water

GPU Enabled Deep Learning on All Data Types Integrating with TensorFlow, MXNet and Caffe

The recent advancements in open-source machine learning platforms and GPU technologies offer a solid foundation for Deep Learning to grow as one of the most prominent fields in Artificial Intelligence. Using complex multi-layer artificial neural networks, Deep Learning helps us derive insights from large unstructured data such as images, videos, sound and text or structured data from transactional databases such as financial data or time series. It enables innovations in domains as varied as medicine, social media, customer service, targeted marketing, automotive safety, security or fraud detection.

Open-source Deep Learning frameworks such as TensorFlow, MXNet and Caffe are optimized for fast training of such models using GPUs. GPUs excel at massively parallel workloads and speed up neural network training by 10-75x compared to conventional CPUs.[1] **Deep Water** brings all these frameworks together under the same user interfaces as the H2O platform. Now, in addition to the original H2O Deep Learning algorithm, users can access TensorFlow, MXNet and Caffe[2] backends in H2O, and build complex deep networks of up to 1,000's of layers with GBs/TBs of data. Processing large datasets becomes orders of magnitude faster.

Bringing all the GPU-optimized frameworks into the H2O platform, Deep Water offers the following benefits:
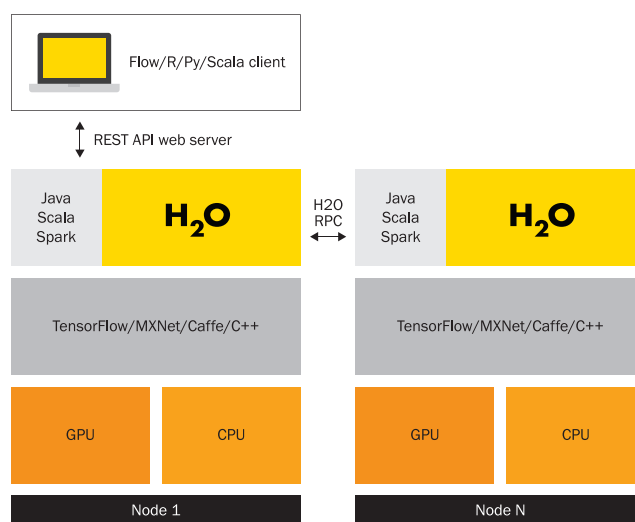

Figure 1: *Deep Water High Level Scalable Architecture*

- **Open Source:** Deep Water is open source. The community has the freedom to use and improve this Deep Learning framework. As of December 2016, H2O has a community of 75,000+ users in 8,500+ organizations.

- **Speed:** Integrated with the state-of-the-art GPU-optimized Deep Learning frameworks – TensorFlow, MXNet and Caffe, H2O Deep Water speeds up Deep Learning on GPUs[3] that are known to be faster than CPUs by 10-75x. This acceleration gives enterprises opportunities to build better models and enable new use cases.

[1]Data quoted from http://www.nvidia.com/object/deep-learning.html
[2]First release of Deep Water includes only MXNet integration. TensorFlow and Caffe integration is pending further announcement.
[3]First release of Deep Water supports training on single GPUs. Distributed training is pending further announcement.

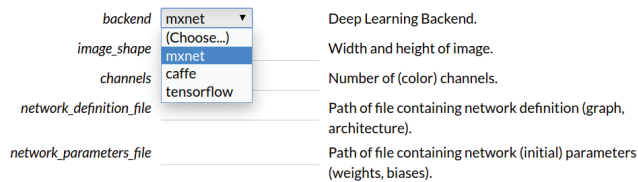| backend | mxnet ▼ | Deep Learning Backend. |
image_shape	(Choose...)	Width and height of image.
	mxnet	
channels	caffe	Number of (color) channels.
	tensorflow	
network_definition_file		Path of file containing network definition (graph, architecture).
network_parameters_file		Path of file containing network (initial) parameters (weights, biases).

Figure 2: Choose GPU-optimized backend using a dropdown in H2O Flow

- **Ease of Use:** Deep Water users can readily access TensorFlow, MXNet and Caffe backends with the interfaces familiar to them: H2O Flow, R, Python, Spark/Scala, Java or the REST API, including model deployment via MOJO and H2O Steam. In H2O Flow, the user can easily switch between the backends from a dropdown menu (see Figure 2).

- **Enterprise Ready:** The H2O platform is known to make model training and deployment easy with its interactive interfaces, distributed algorithms and scalable architecture. Enterprise features such as hyper-parameter optimization, cross-validation and automatic model tuning are fully supported by Deep Water.

To Try Out H2O Deep Water

You can choose to download, compile the source code and configure your own environment or use the pre-built AMI Image:

- Github repository:
  https://github.com/h2oai/deepwater

- GPU system dependencies:
  - Latest NVIDIA Display driver
  - CUDA 8
  - CUDNN 5.1

- Pre-built AMI Image:
  - Refer to **http://docs.h2o.ai**. Deep Water AMI Guide is listed in the Deep Water section in the Getting Started section.

- Deep Water product page:
  http://www.h2o.ai/deep-water/

About H2O.ai

H2O.ai is focused on bringing AI to businesses through software. Its flagship product is H2O, the leading open source platform that makes it easy for financial services, insurance and healthcare companies to deploy machine learning and predictive analytics to solve complex problems. More than 8,500+ organizations and 75,000+ data scientists depend on H2O for critical applications like predictive maintenance and operational intelligence. The company accelerates business transformation for 107 Fortune 500 enterprises, 8 of the world's 12 largest banks, 7 of the 10 largest auto insurance companies and all 5 major telecommunications providers. Notable customers include Capital One, Progressive Insurance, Transamerica, Comcast, Nielsen Catalina Solutions, Macy's, Walgreens, Kaiser Permanente, and Aetna.

Follow us on Twitter @h2oai. To learn more about H2O customer use cases, please visit http://www.h2o.ai/customers/. **Join the Movement.**