DOING LARGE SCALE AI FASTER

Miguel Martínez - Solution Architect

NEURAL NETWORKS ARE NOT NEW

And are surprisingly simple as an algorithm

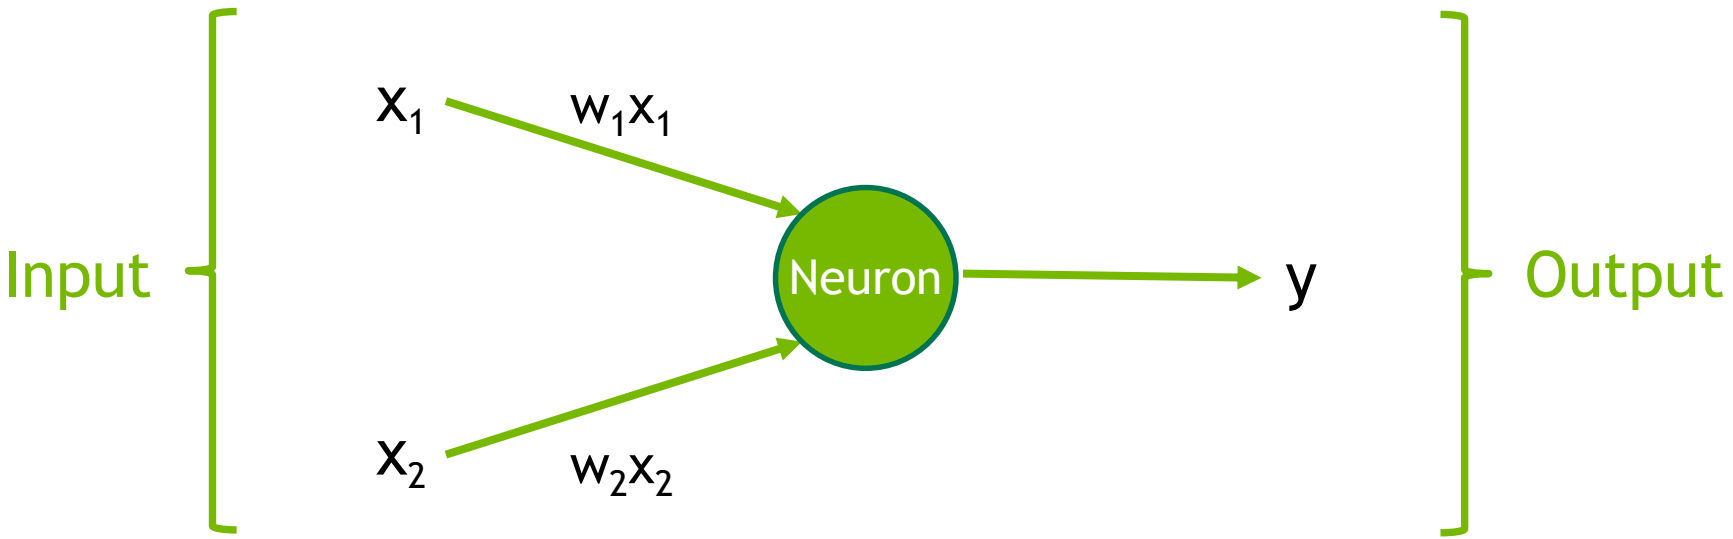

A simple neuron

COMBINING NEURONS

Stacking neurons and layers creates a more powerful model

Additional neurons can be added to create a layer.

Multiple layers can also be added, resulting in input, hidden, and output layers.

Expanding the neural network size creates additional predictive power.

In feed forward neural networks, neurons are fully connected to surrounding layers.

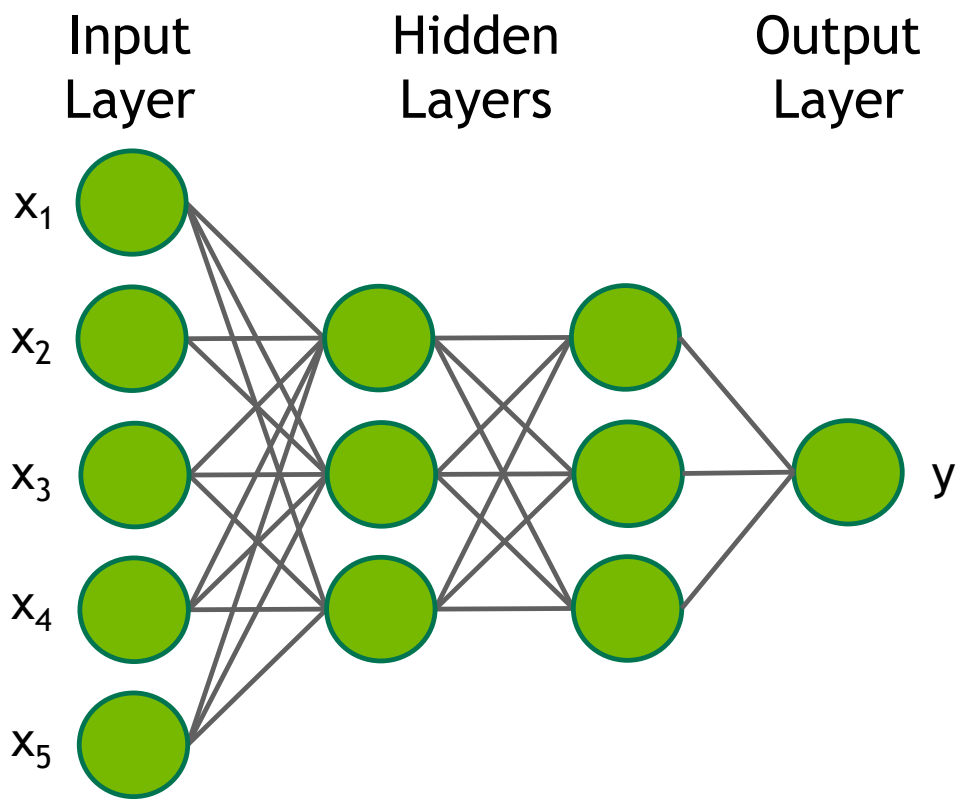

DEEP NEURAL NETWORKS (DNNS)

Neural networks with many layers enable deep learning

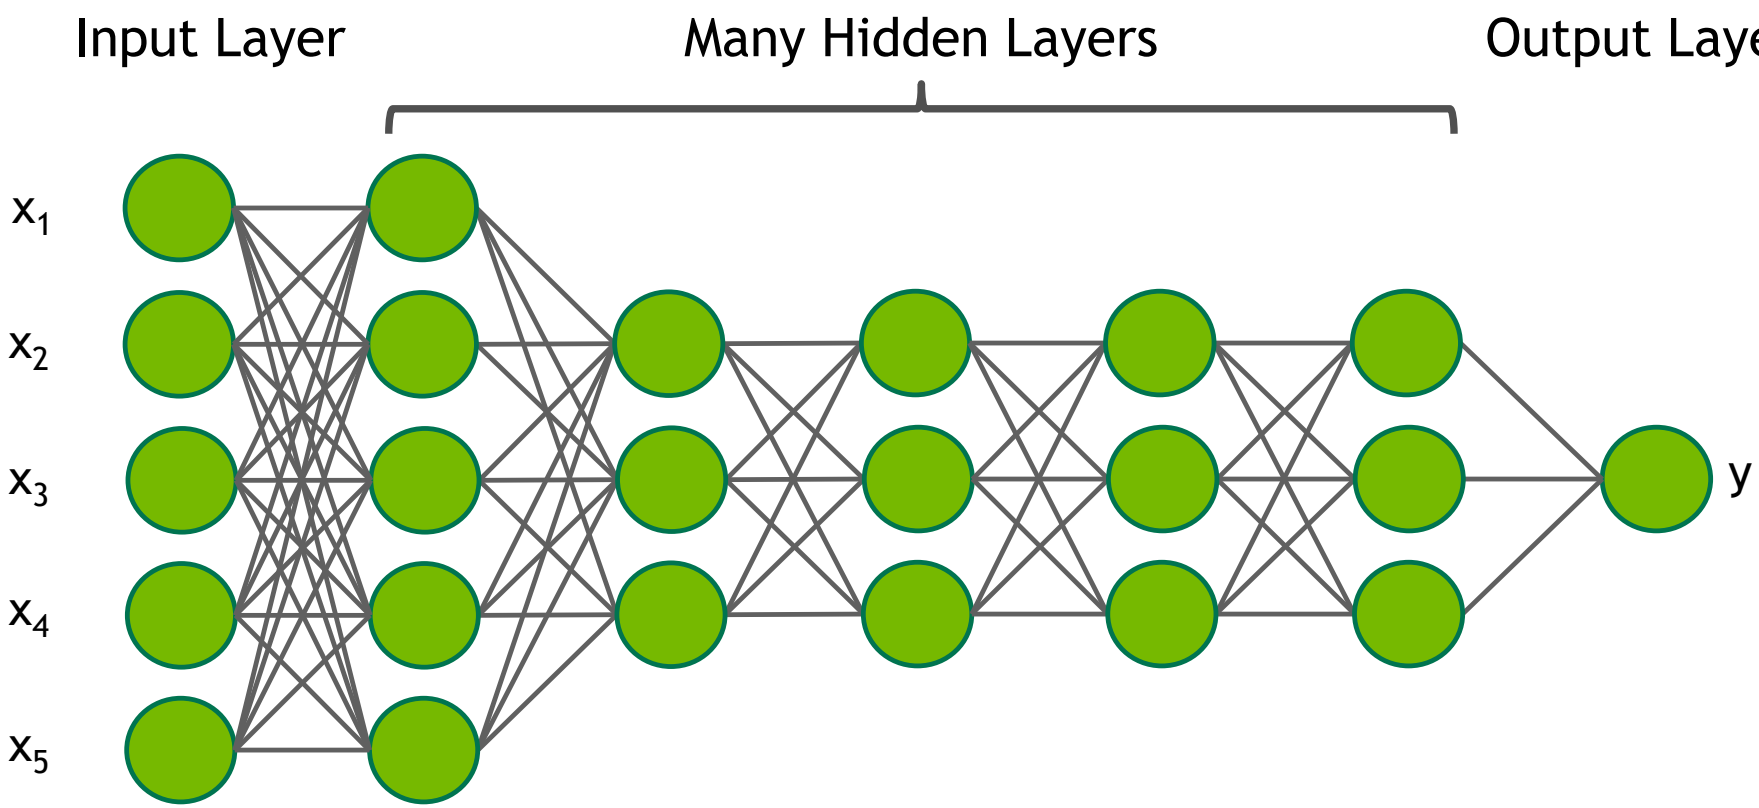

NEURAL NETWORKS ARE NOT NEW

They just historically never worked well

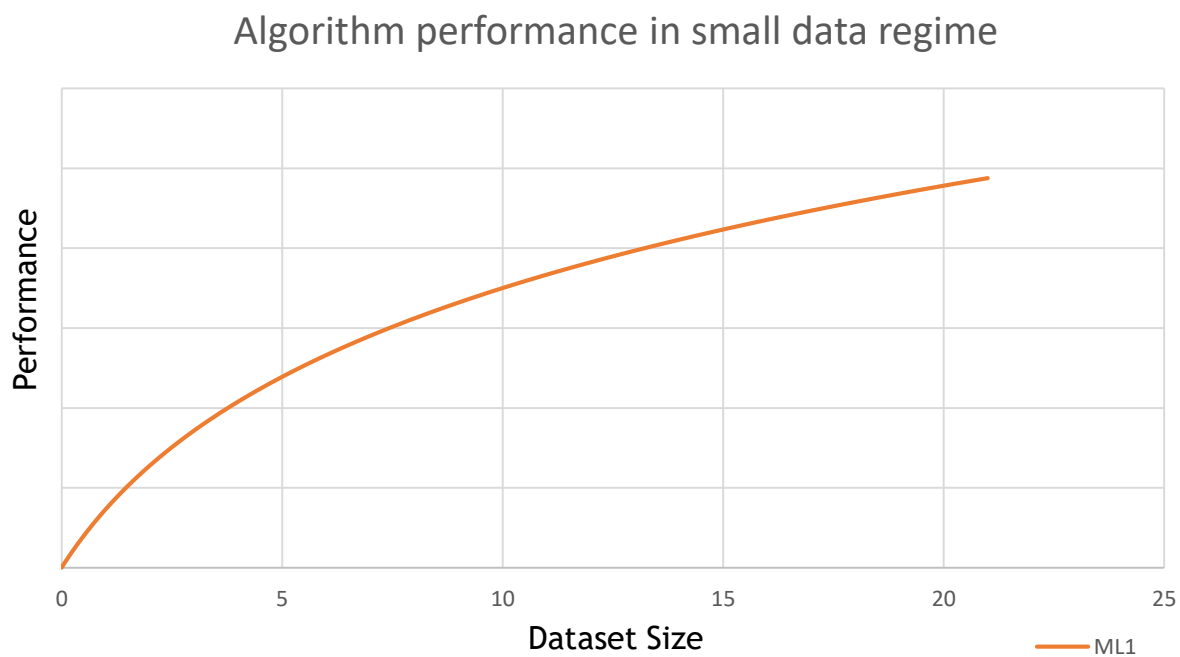

Andrew Ng, “Nuts and Bolts of Applying Deep Learning”, https://www.youtube.com/watch?v=F1ka6a13S9I

NEURAL NETWORKS ARE NOT NEW

They just historically never worked well

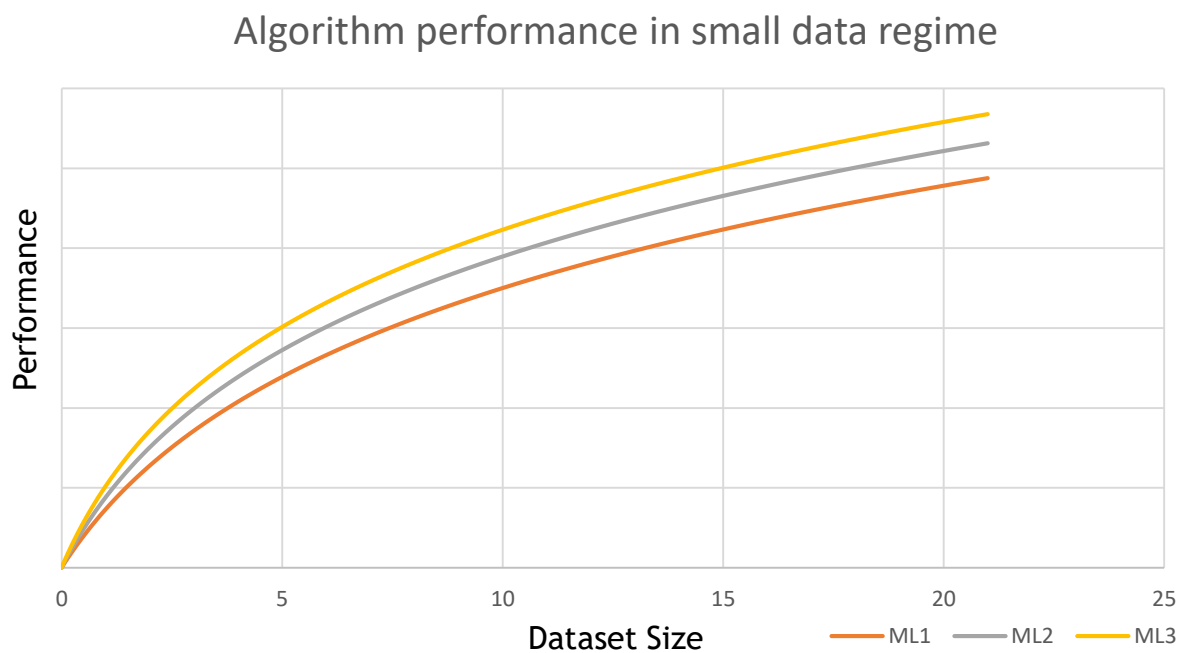

Andrew Ng, “Nuts and Bolts of Applying Deep Learning”, https://www.youtube.com/watch?v=F1ka6a13S9I

NEURAL NETWORKS ARE NOT NEW

They just historically never worked well

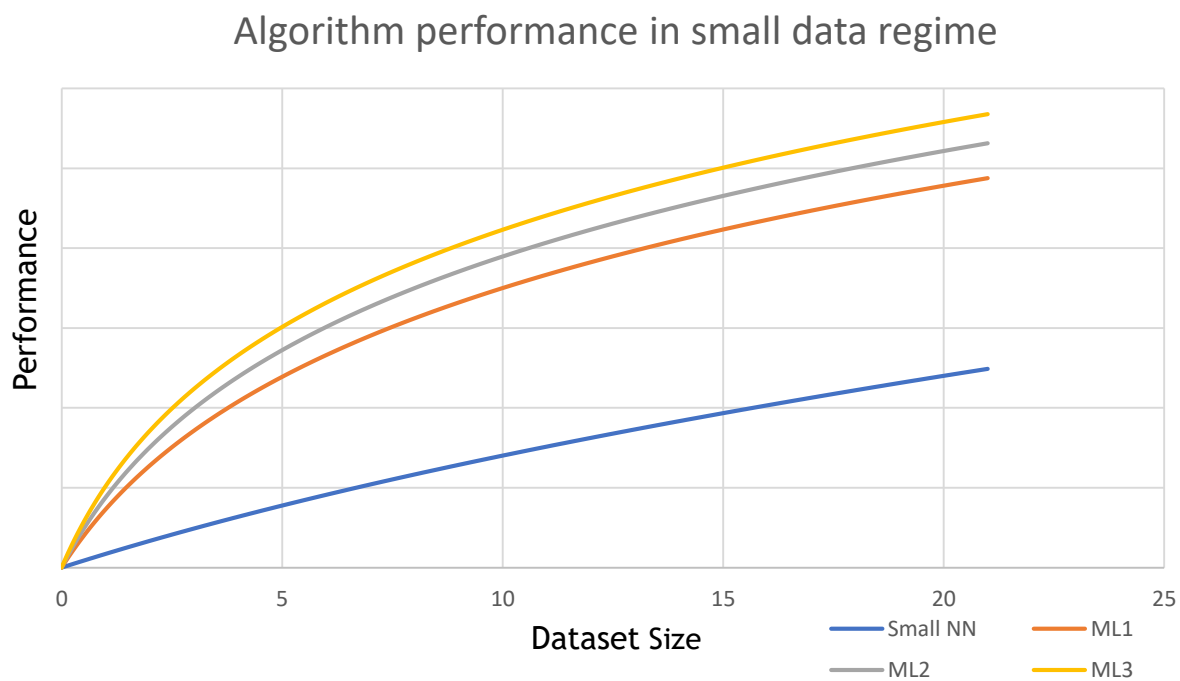

Andrew Ng, “Nuts and Bolts of Applying Deep Learning”, https://www.youtube.com/watch?v=F1ka6a13S9I

NEURAL NETWORKS ARE NOT NEW

But that changed and transformed the way we do ML

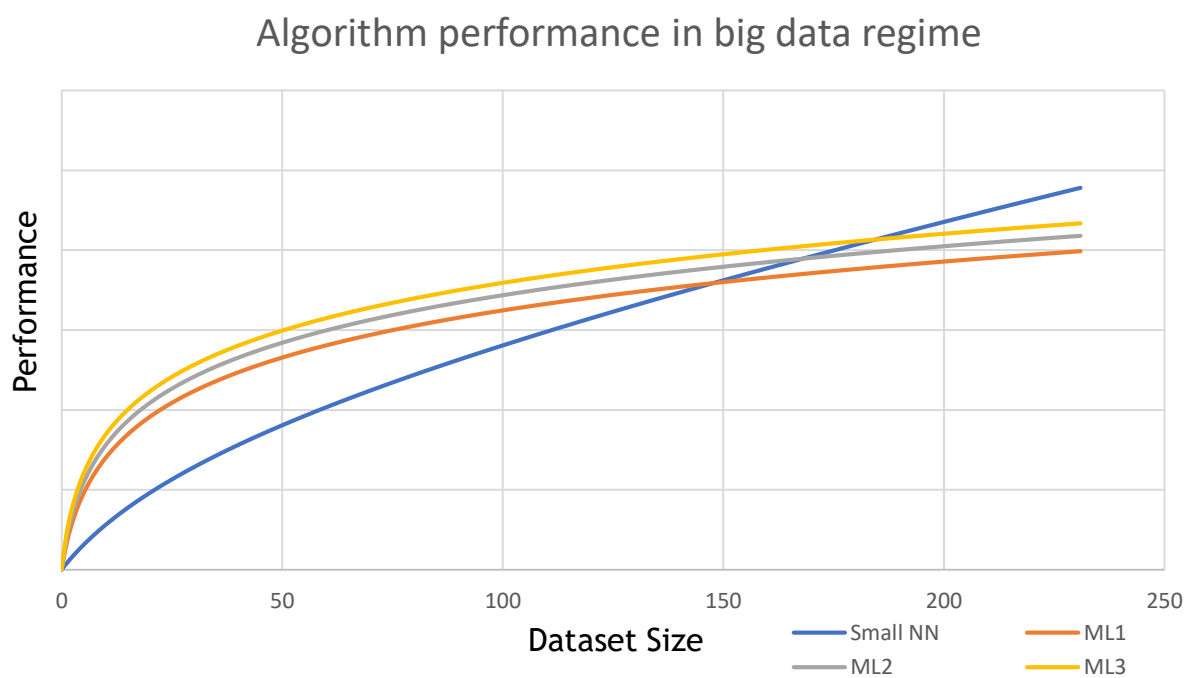

Andrew Ng, “Nuts and Bolts of Applying Deep Learning”, https://www.youtube.com/watch?v=F1ka6a13S9I

NEURAL NETWORKS ARE NOT NEW

Data and model size the key to accuracy

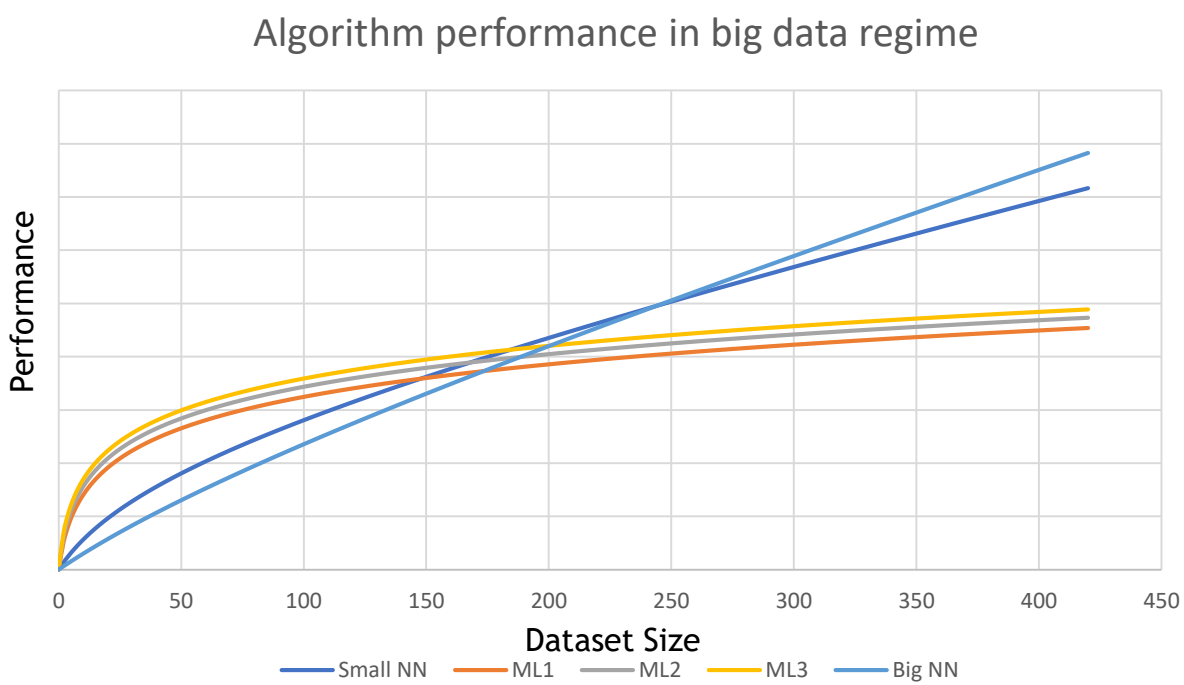

Andrew Ng, “Nuts and Bolts of Applying Deep Learning”, https://www.youtube.com/watch?v=F1ka6a13S9I

NEURAL NETWORKS ARE NOT NEW

Exceeding human level performance

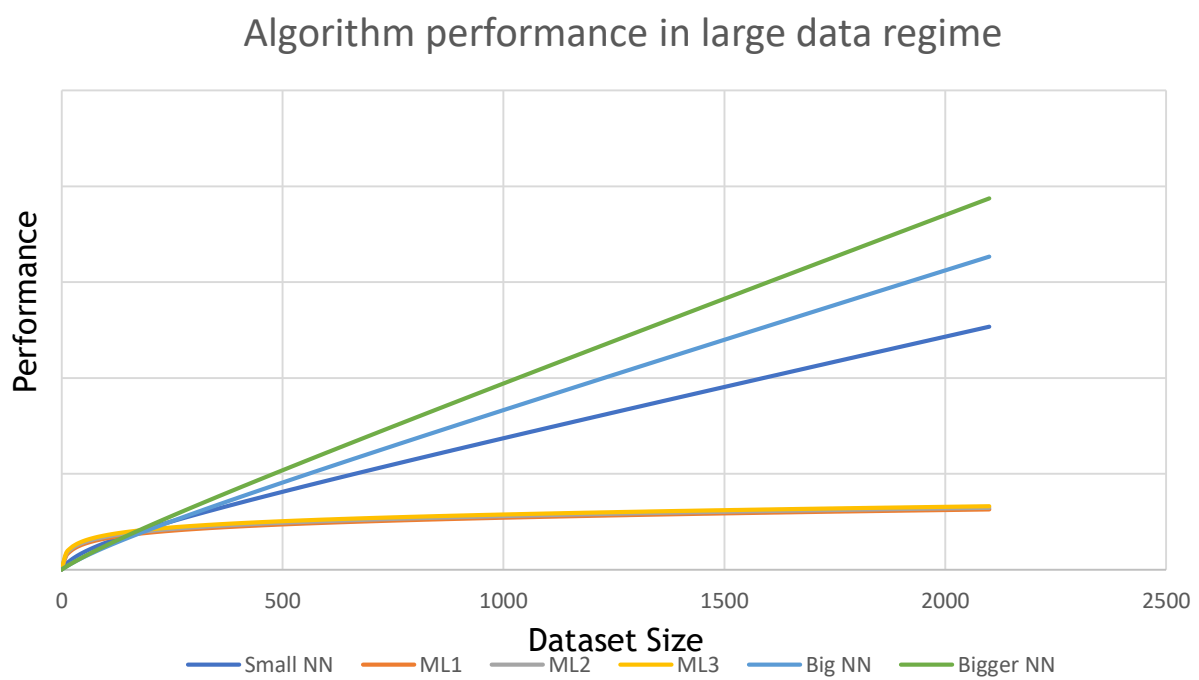

Andrew Ng, “Nuts and Bolts of Applying Deep Learning”, https://www.youtube.com/watch?v=F1ka6a13S9I

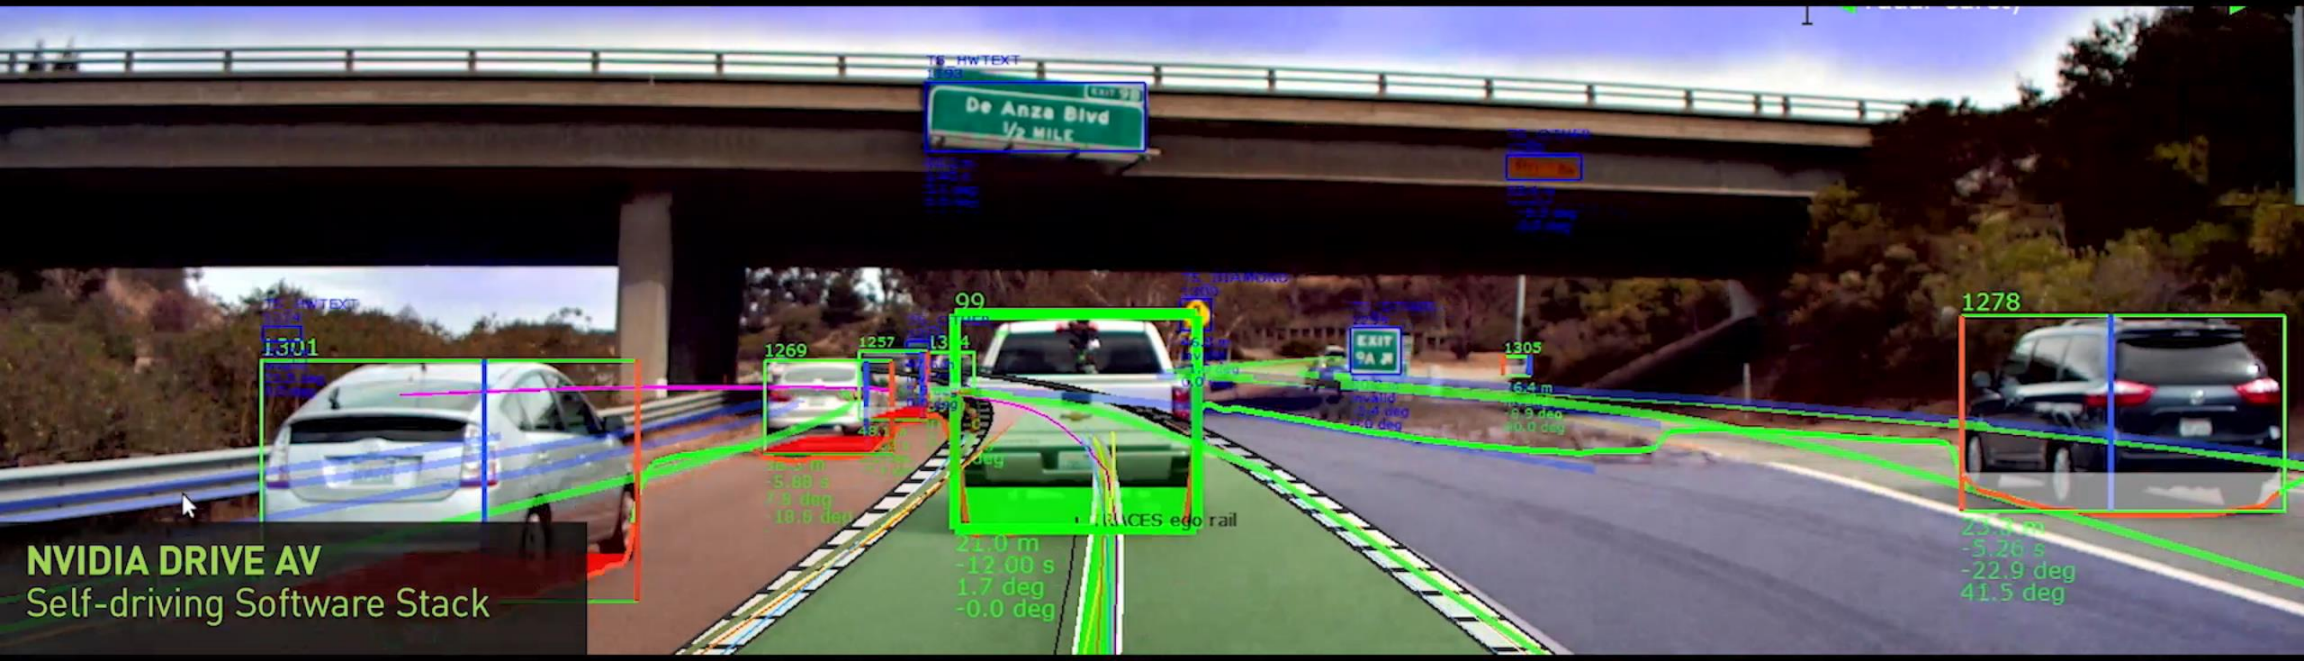
De Anza Blvd
1/2 MILE
1301
1269
1257
99
1305
1278
EXIT
-5.88 s
21.0 m
-12.00 s
1.7 deg
-0.0 deg
-5.26 s
-22.9 deg
41.5 deg
NVIDIA DRIVE AV
Self-driving Software Stack

IMPLICATIONS

Automotive example

Majority of problems are too complex for a single GPU training

	VERY CONSERVATIVE	CONSERVATIVE
Fleet size (data capture per hour)	100 cars / 1TB/hour	125 cars / 1.5TB/hour
Duration of data collection	260 days * 8 hours	325 days * 10 hours
Data Compression factor	0.0005	0.0008
Total training set	104 TB	487.5 TB
InceptionV3 training time (with 1 Pascal GPU)	9.1 years	42.6 years
AlexNet training time (with 1 Pascal GPU)	 1.1 years 	 5.4 years 

IMPLICATIONS

Experimental Nature of DL - Unacceptable training time

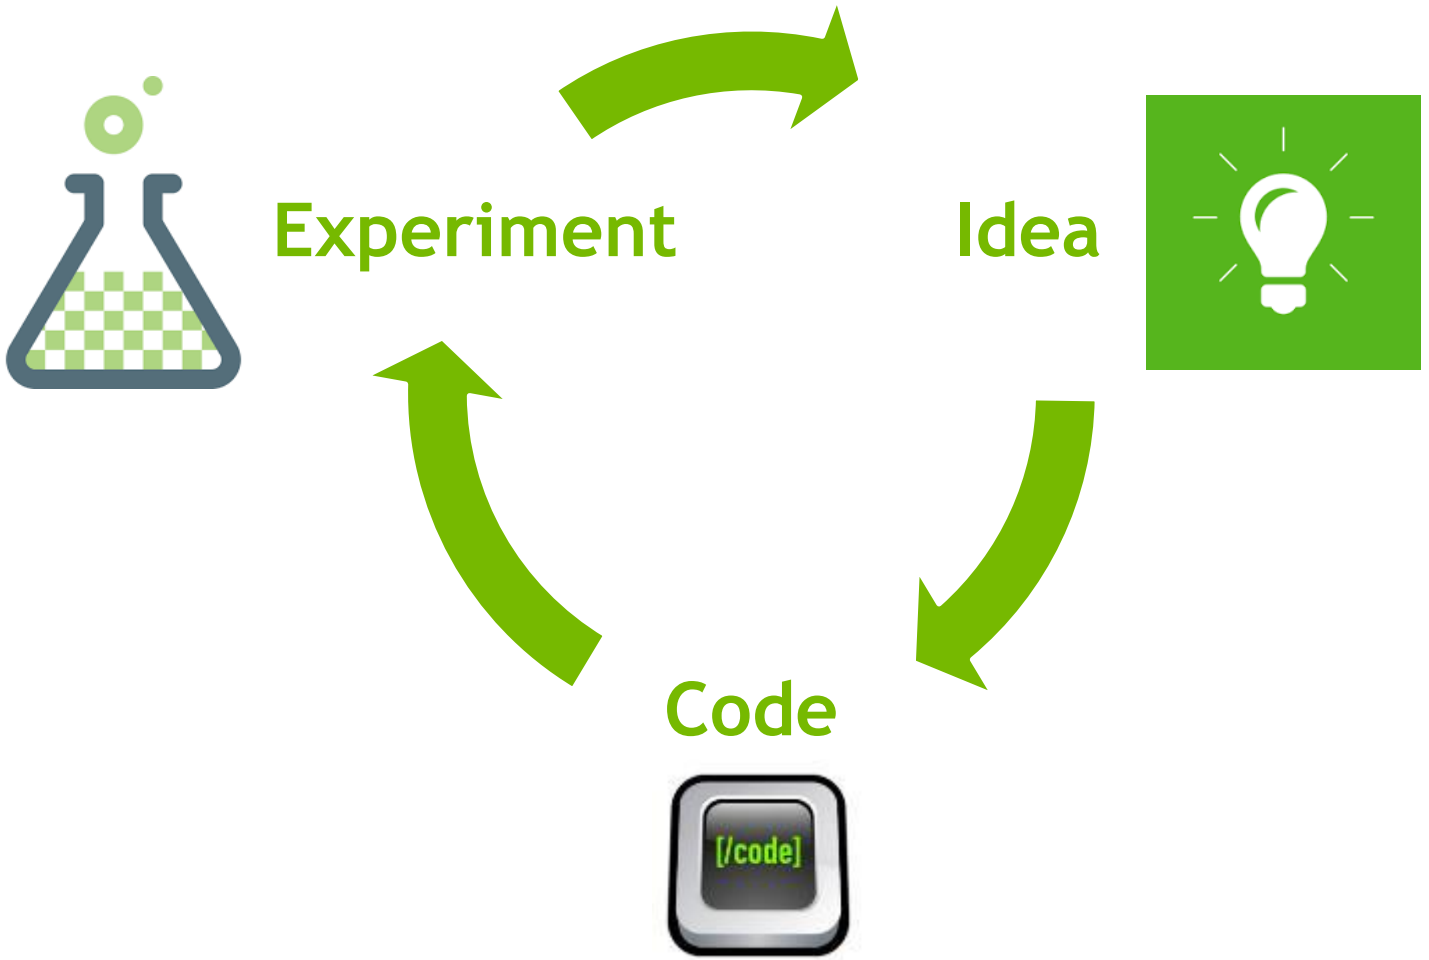

CONCLUSIONS

What does your team in the meanwhile?

CONCLUSIONS

Need to scale the training process for a single job

1 NVIDIA DGX-1

10 NVIDIA DGX-1's

	VERY CONSERVATIVE	CONSERVATIVE
Total training set	104 TB	487.5 TB
InceptionV3 (one DGX-1V)	166 days (5+ **months**)	778 days (2+ **years**)
AlexNet (one DGX-1V)	21 days (3 **weeks**)	98 days (3 **months**)
InceptionV3 (10 DGX-1V's)	16 days (2+ **weeks**)	77 days (11 **weeks**)
AlexNet (10 DGX-1V's)	2.1 days	9.8 days

Training From Months or Years

To Weeks or Days

ITERATION TIME

Short iteration time is fundamental for success

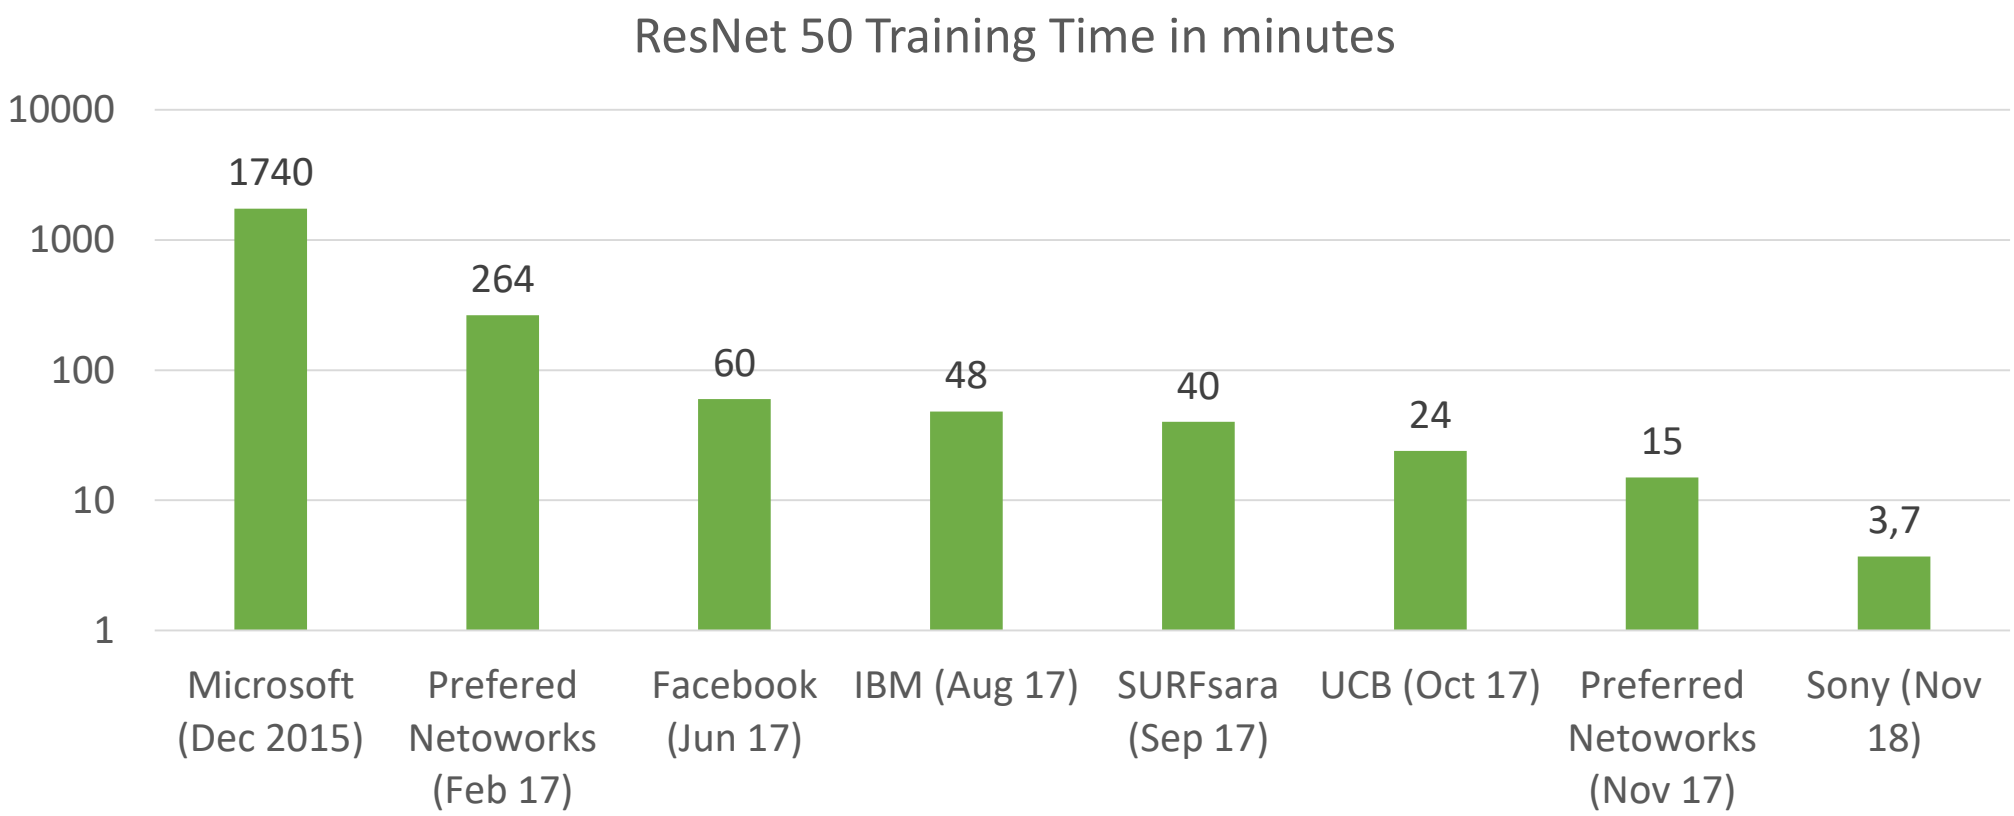

SAMPLE: PREFERRED NETWORKS

Training ImageNet in 15 minutes

- It consists of 128 nodes with 8 NVIDIA TESLA GPUs each, for **1024 GPUs in total**.
- The nodes are connected with two FDR Infiniband links (56Gbps x 2).

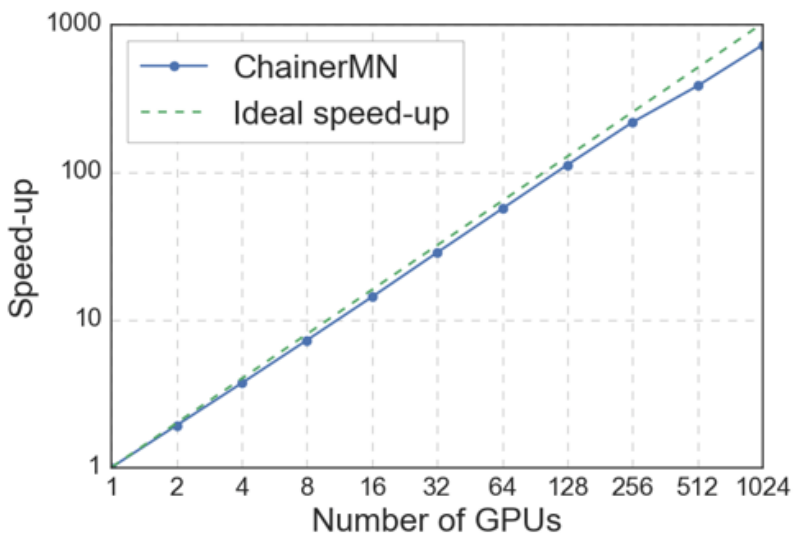

Akiba, T., Suzuki, S., & Fukuda, K. (2017). Extremely large minibatch sgd: Training resnet-50 on imagenet in 15 minutes. *arXiv preprint arXiv:1711.04325*.

DATA vs MODEL PARALLELISM

Comparison

- Data Parallelism
  - Allows to speed up training
  - All workers have the same copy of the model
  - All workers train on different data
  - Neural network gradients (weight changes) are exchanged
- Model Parallelism
  - Allows for a bigger model
  - Parts of the model are distributed across GPUs
  - All workers train on the same data
  - Neural network activations are exchanged

DATA PARALLELISM

How it works

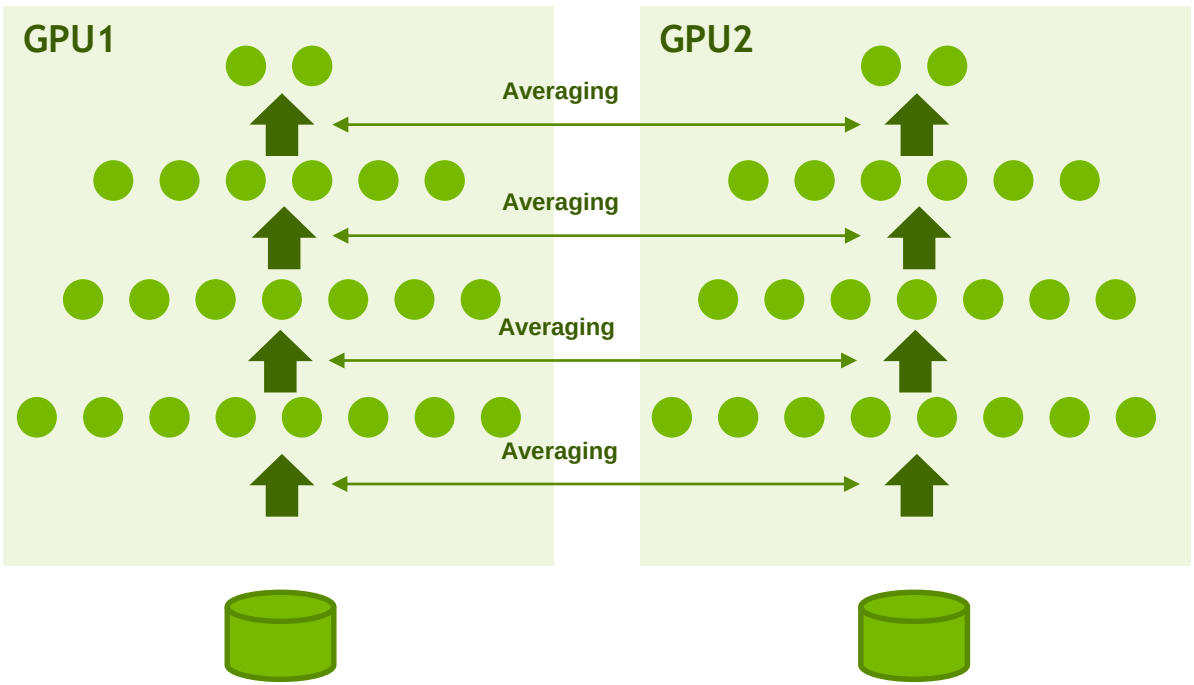

HOW DO OUR TRAINING JOBS LOOK LIKE

- SINGLE GPU CODE is a dying specie
- All our DL code is made for MULTIGPU and scalable:
  - Runs on Single GPU
  - Runs on Multi GPU
  - Runs on Multi Nodes with Multiple GPUs
- Just ONE codebase

SOME HPC AGAIN

Starting MPI jobs - all commands run into NGC containers

- 2 GPUs: mpirun -np 2 --allow-run-as-root -H localhost:2 -bind-to none -map-by slot -x NCCL_DEBUG=INFO -x LD_LIBRARY_PATH -x PATH -mca pml ob1 -mca btl ^openib python test.py
- 8 GPUs: mpirun -np 8 --allow-run-as-root -H localhost:8 -bind-to none -map-by slot -x NCCL_DEBUG=INFO -x LD_LIBRARY_PATH -x PATH -mca pml ob1 -mca btl ^openib python test.py

OK, you got the gist of it, now lets go beyond a single node

- 24 GPUs across 3 DGX-1: mpirun -np 24 --allow-run-as-root -H dgx1:8,dgx2:8,dgx3:8 -bind-to none -map-by slot -x NCCL_DEBUG=INFO -x LD_LIBRARY_PATH -x PATH -mca pml ob1 -mca btl ^openib python test.py

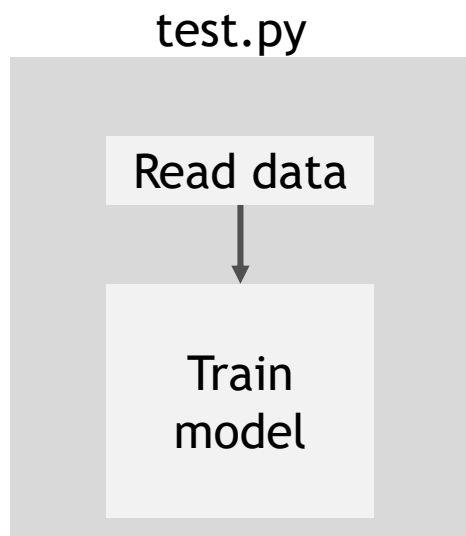

Horovod enabled DL code

BRIEF ARCHITECTURE OVERVIEW

DGX-1

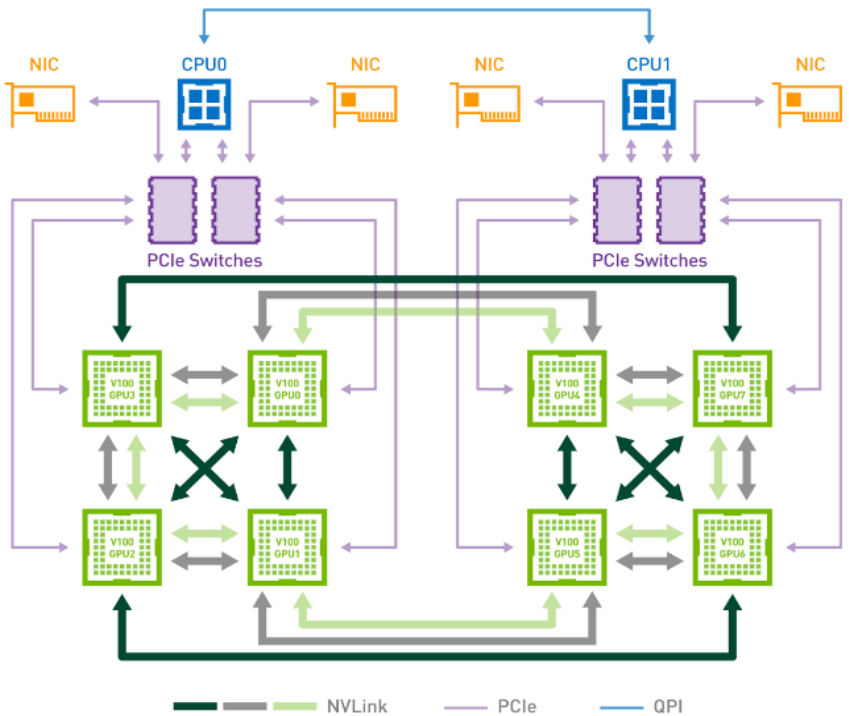

Interconnected GPUs

- Internal: NVLINK/NVSWITCH
- External: IB/RoCE2
- Precondition for effective MultiGPU
- IB/RoCE2 also perfect for Storage

STACKING DGX

One Process per GPU – One Data Pipeline per GPU

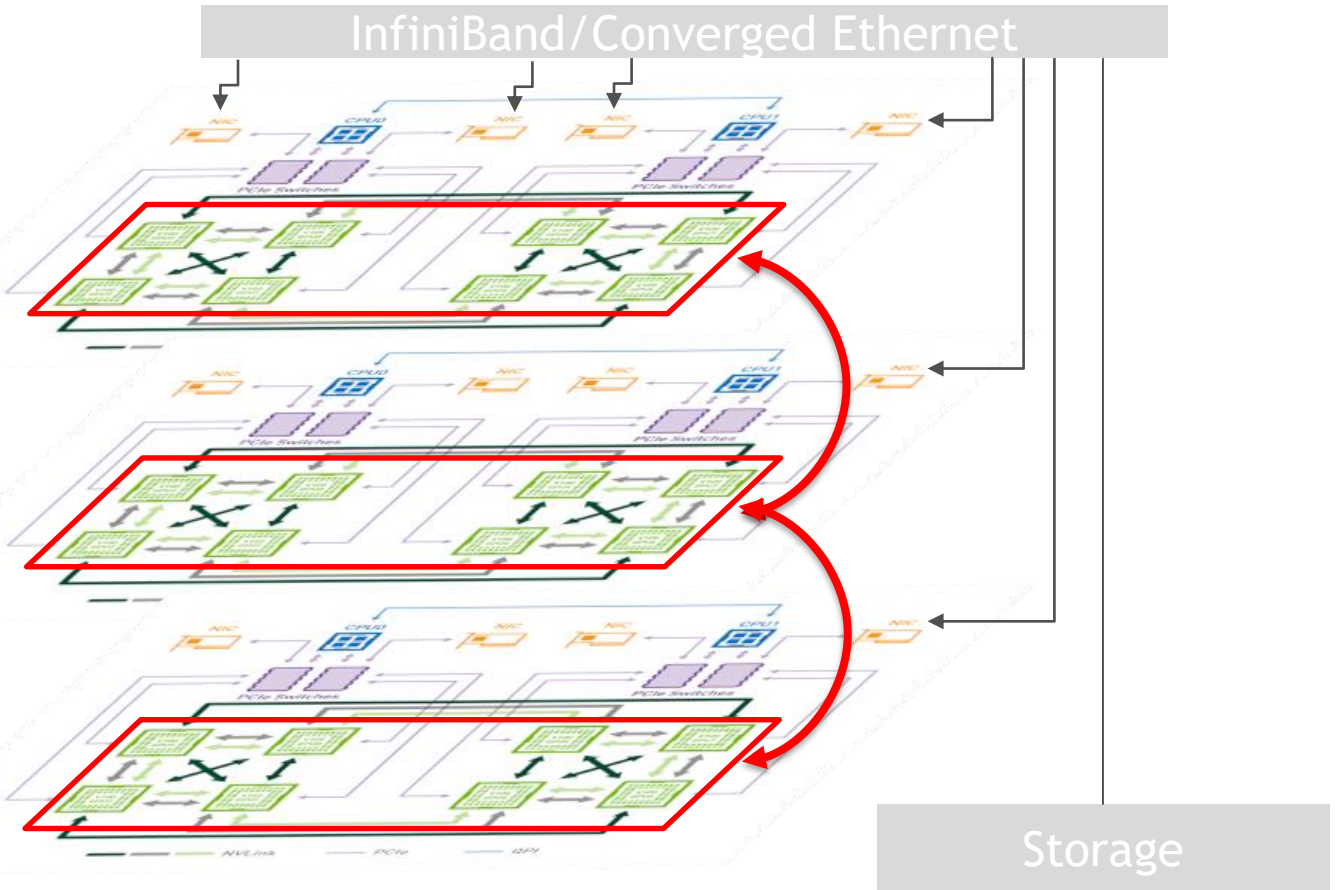

test.py

Read data

Train model

SCALING WITH HOROVOD

One Process per GPU – One Datapipeline per GPU

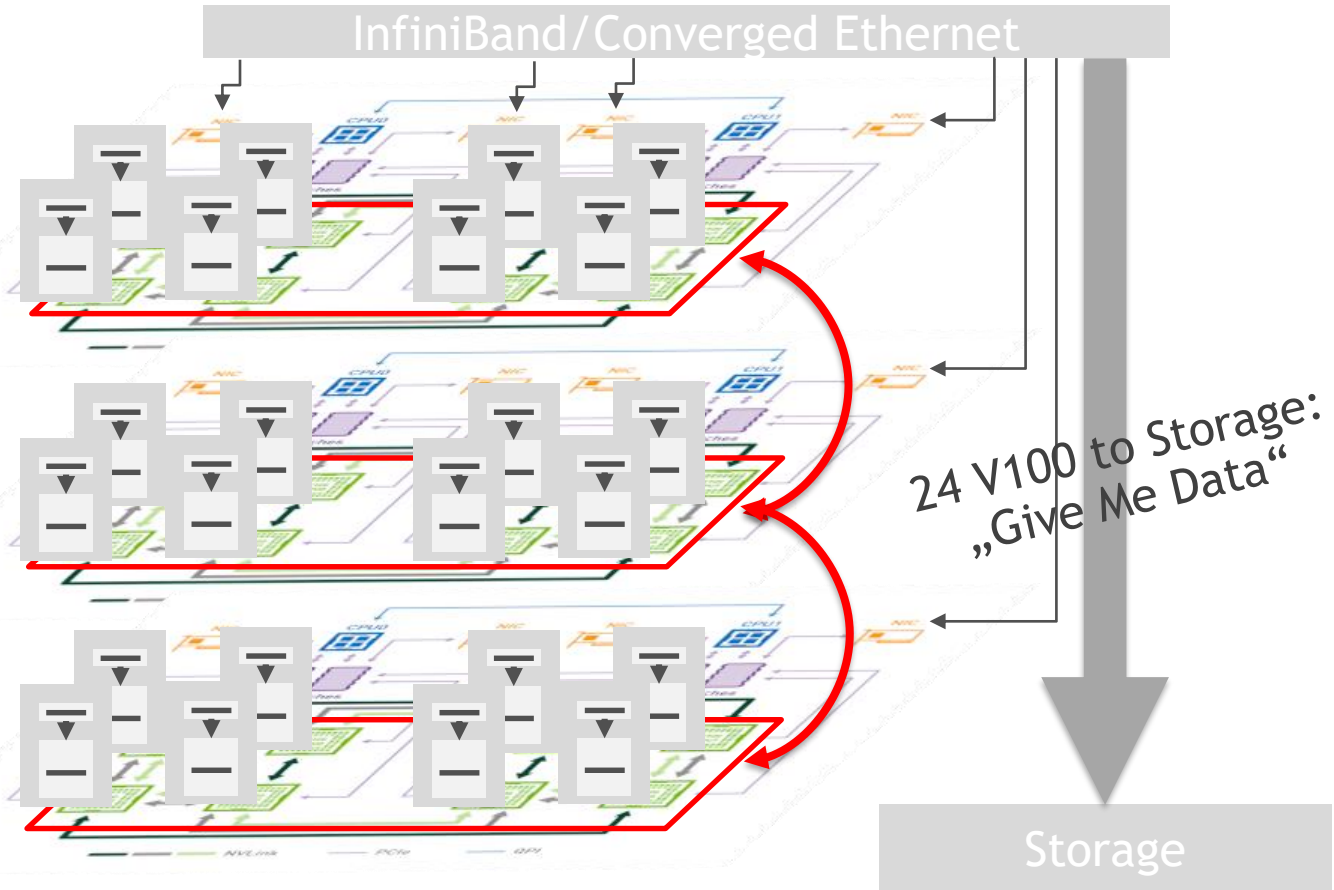

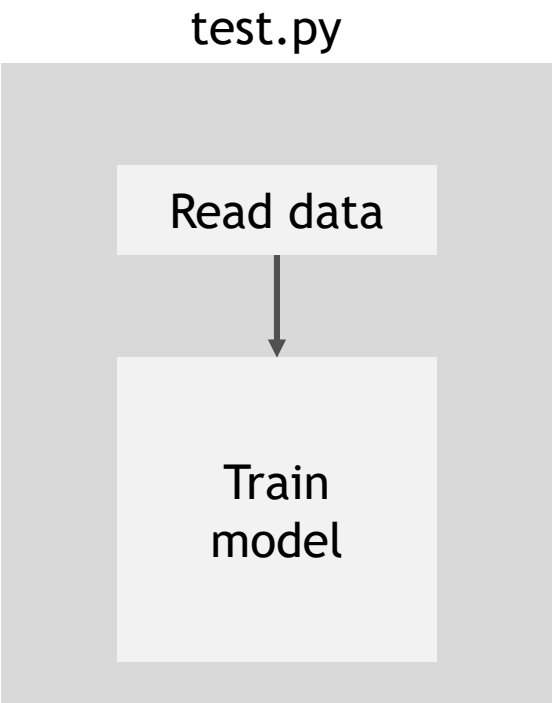

IBM SPECTRUM STORAGE FOR AI WITH NVIDIA DGX

The Engine to Power Your AI Data Pipeline

HARDWARE

- **NVIDIA DGX-1 | up to 9**x DGX-1 Systems
- **IBM Spectrum Scale NVMe Appliance|** 40GB/s per node, 120GB/s in 6RU| 300TB per node
- **NETWORK: Mellanox SB7700 Switch |** 2x EDR IB with RDMA

SOFTWARE

- **NVIDIA DGX SOFTWARE STACK** | NVIDIA Optimized Frameworks
- **IBM:** High performance, low latency, parallel file system
- **IBM:** Extensible and composable

IBM & NVIDIA REFERENCE ARCHITECTURE

Validated design for deploying DGX at-scale with IBM Storage

Download at
https://bit.ly/2GcYbgO

Learn more about
DGX RA Solutions at:
https://bit.ly/2OpXYeC

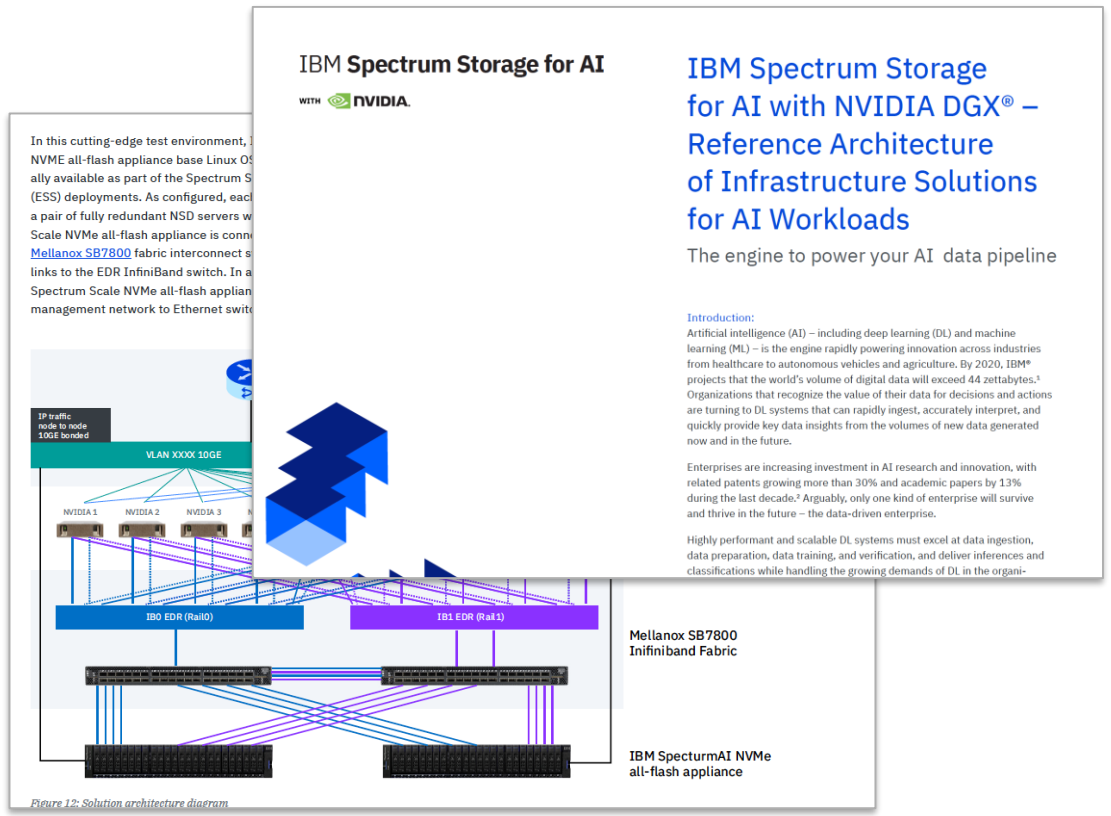

Thank You!

NVIDIA

IBM Spectrum Storage