NVIDIA FOR DEEP LEARNING

Bill Veenhuis | bveenhuis@nvidia.com

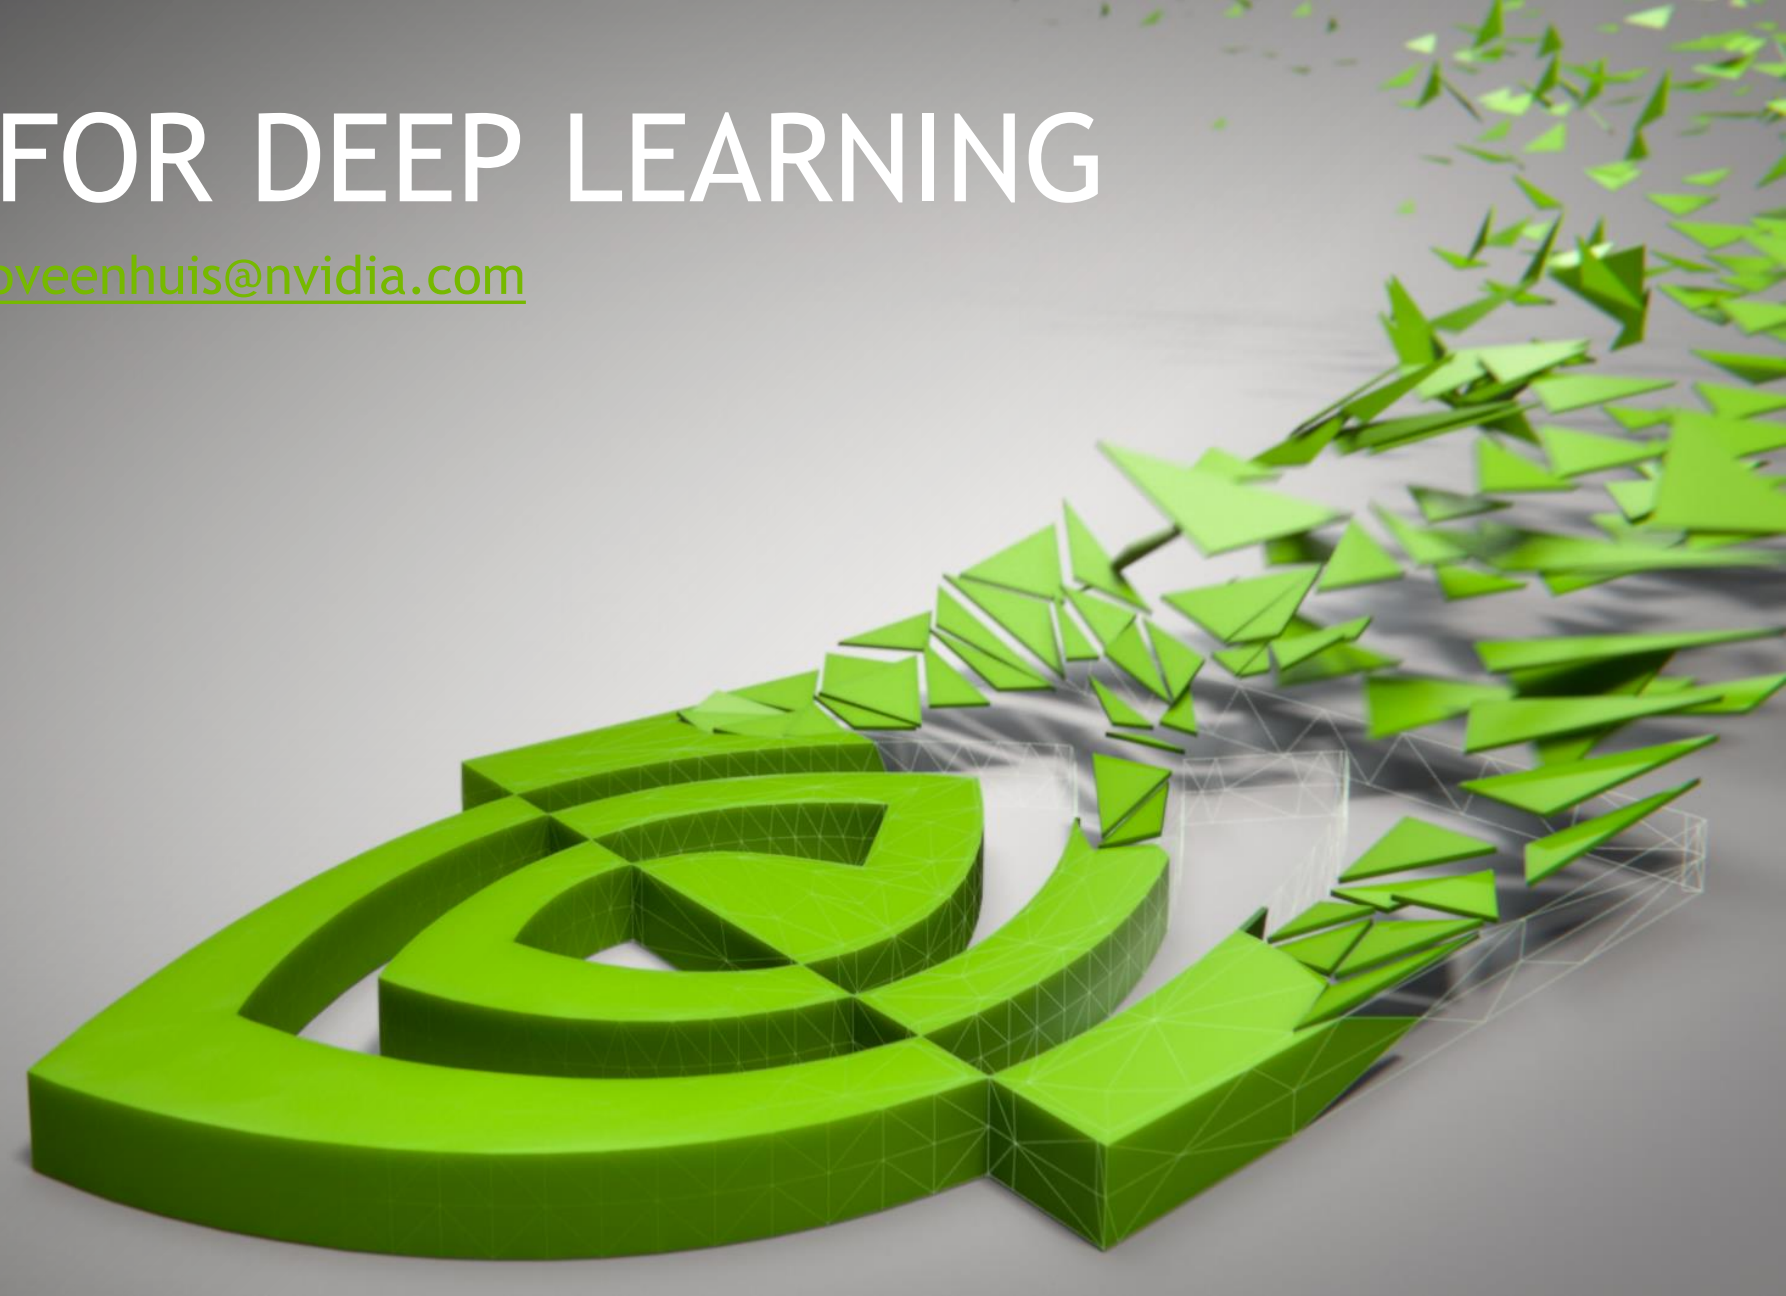

Nvidia is the world’s leading ai platform

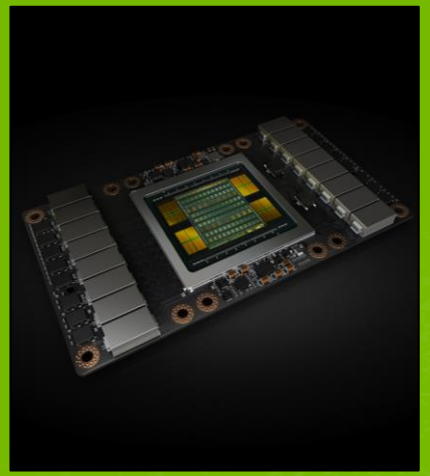


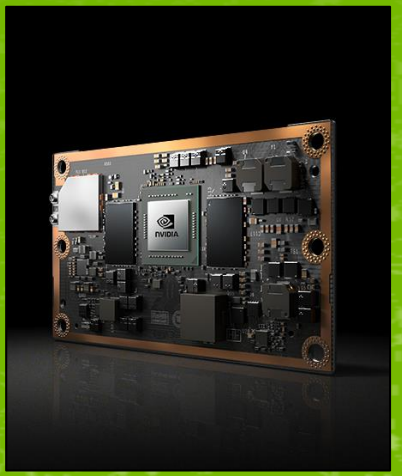

ONE ARCHITECTURE — CUDA

GPU: Perfect Companion for Accelerating Apps & A.I.

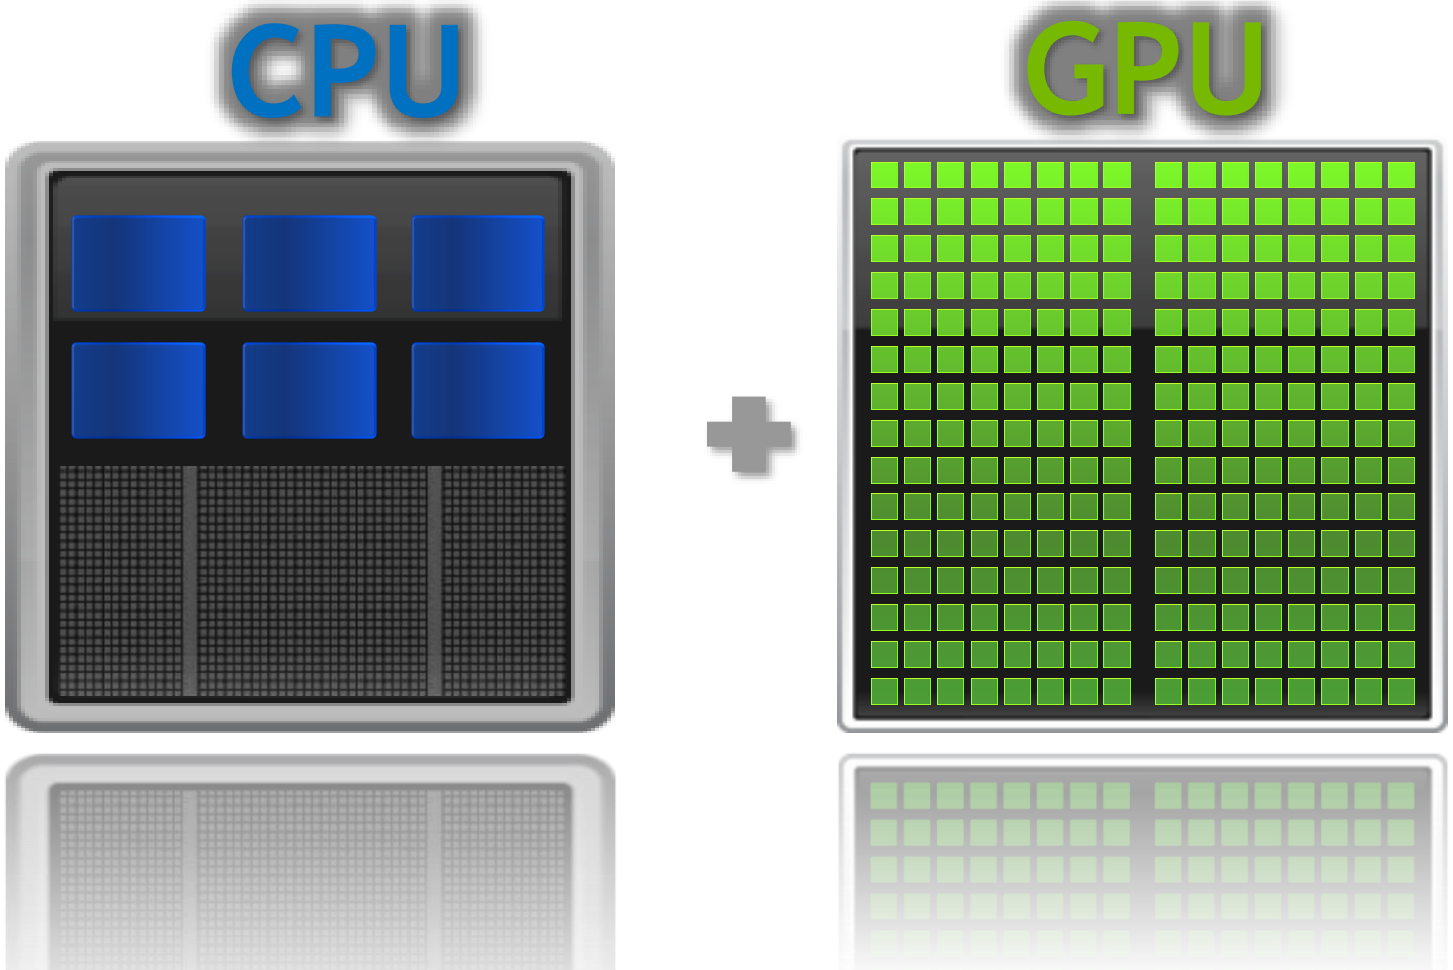

AGENDA & TOPICS

- Intro to AI
- Deep Learning Intro
- NVIDIA’s DIGITS
- Autoencoding Enhancement
- TensorRT

Intro to AI

ARTIFICIAL NEURONS

Biological neuron

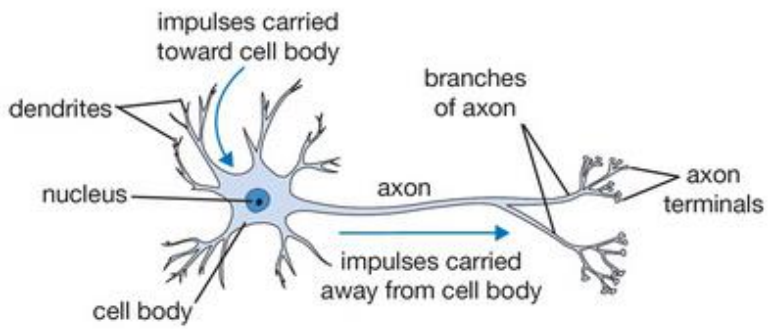

From Stanford cs231n lecture notes

Artificial neuron

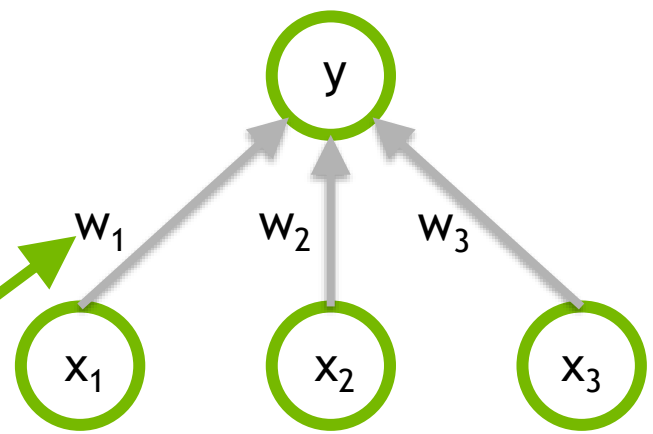

Weights (W_n)
= parameters

$y=F(w_1x_1+w_2x_2+w_3x_3)$

ARTIFICIAL NEURAL NETWORK

A collection of simple, trainable mathematical units that collectively learn complex functions

Hidden layers

Input layer

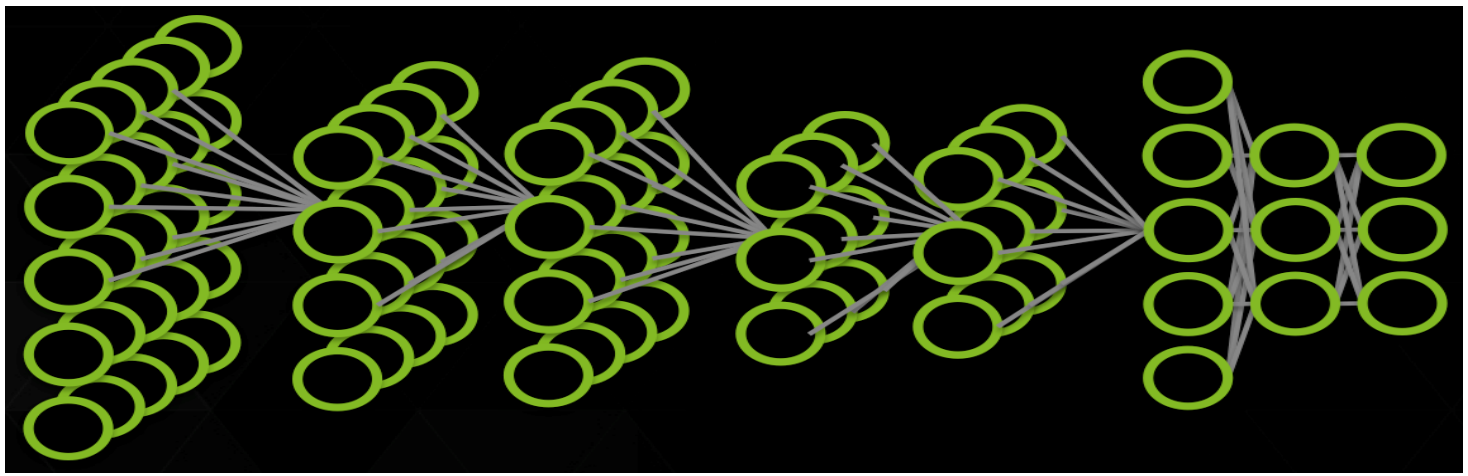

Output layer

Given sufficient training data an artificial neural network can approximate very complex functions mapping raw data to output decisions

DEEP NEURAL NETWORK (DNN)

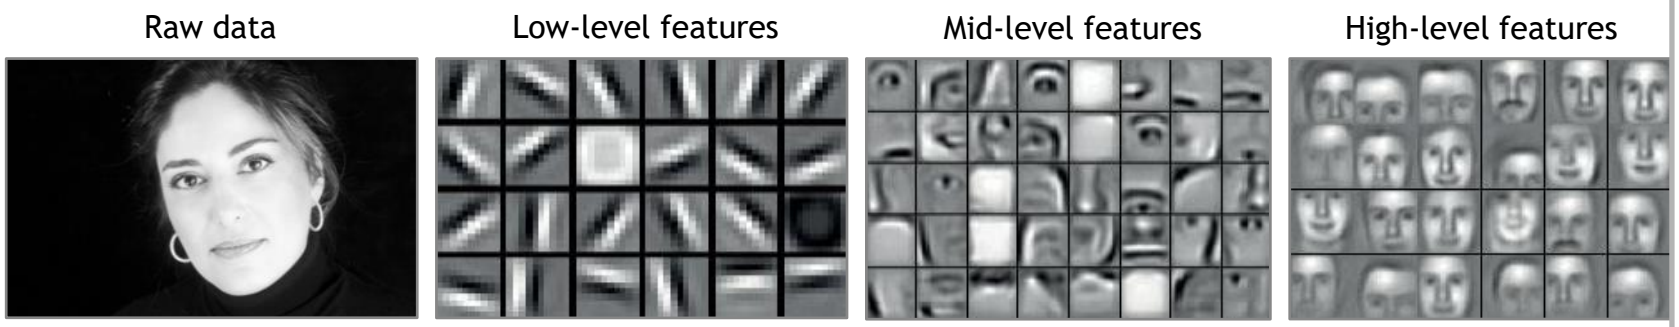

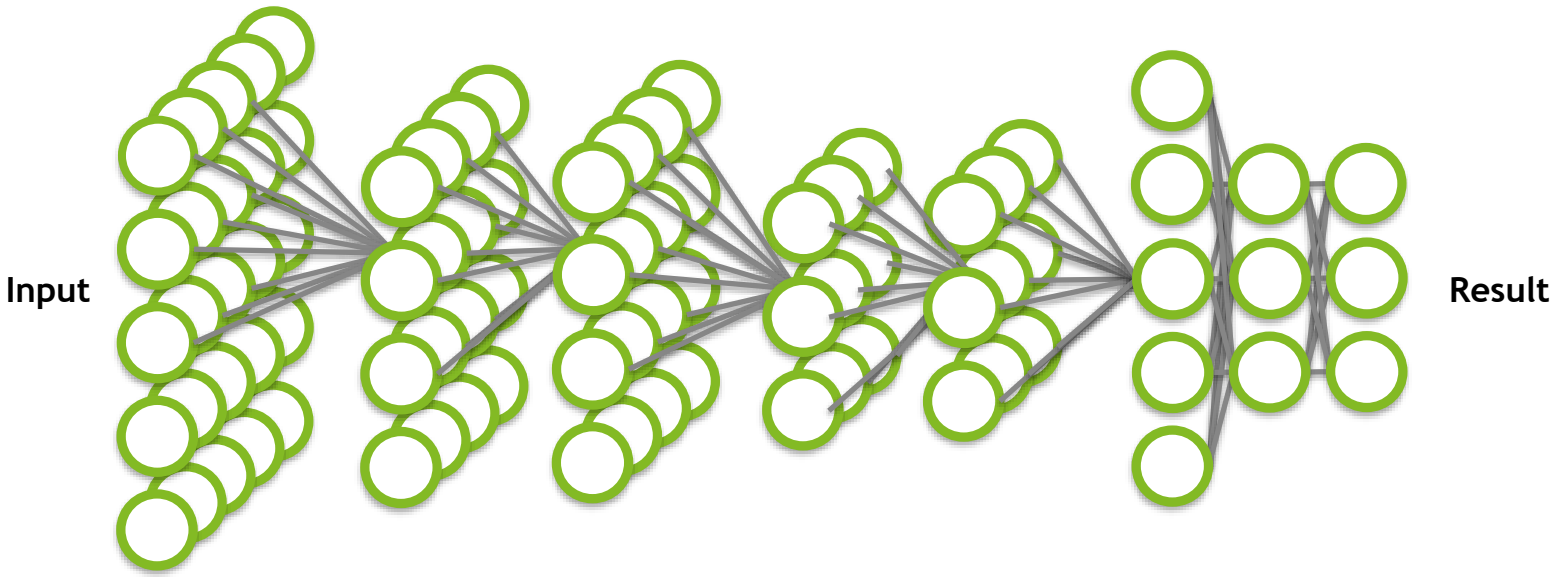

Application components:

Task objective
e.g. Identify face

Training data
10-100M images

Network architecture
~10s-100s of layers
1B parameters

Learning algorithm
~30 Exaflops
1-30 GPU days

WHAT IS DEEP LEARNING?

Accomplishing complex goals

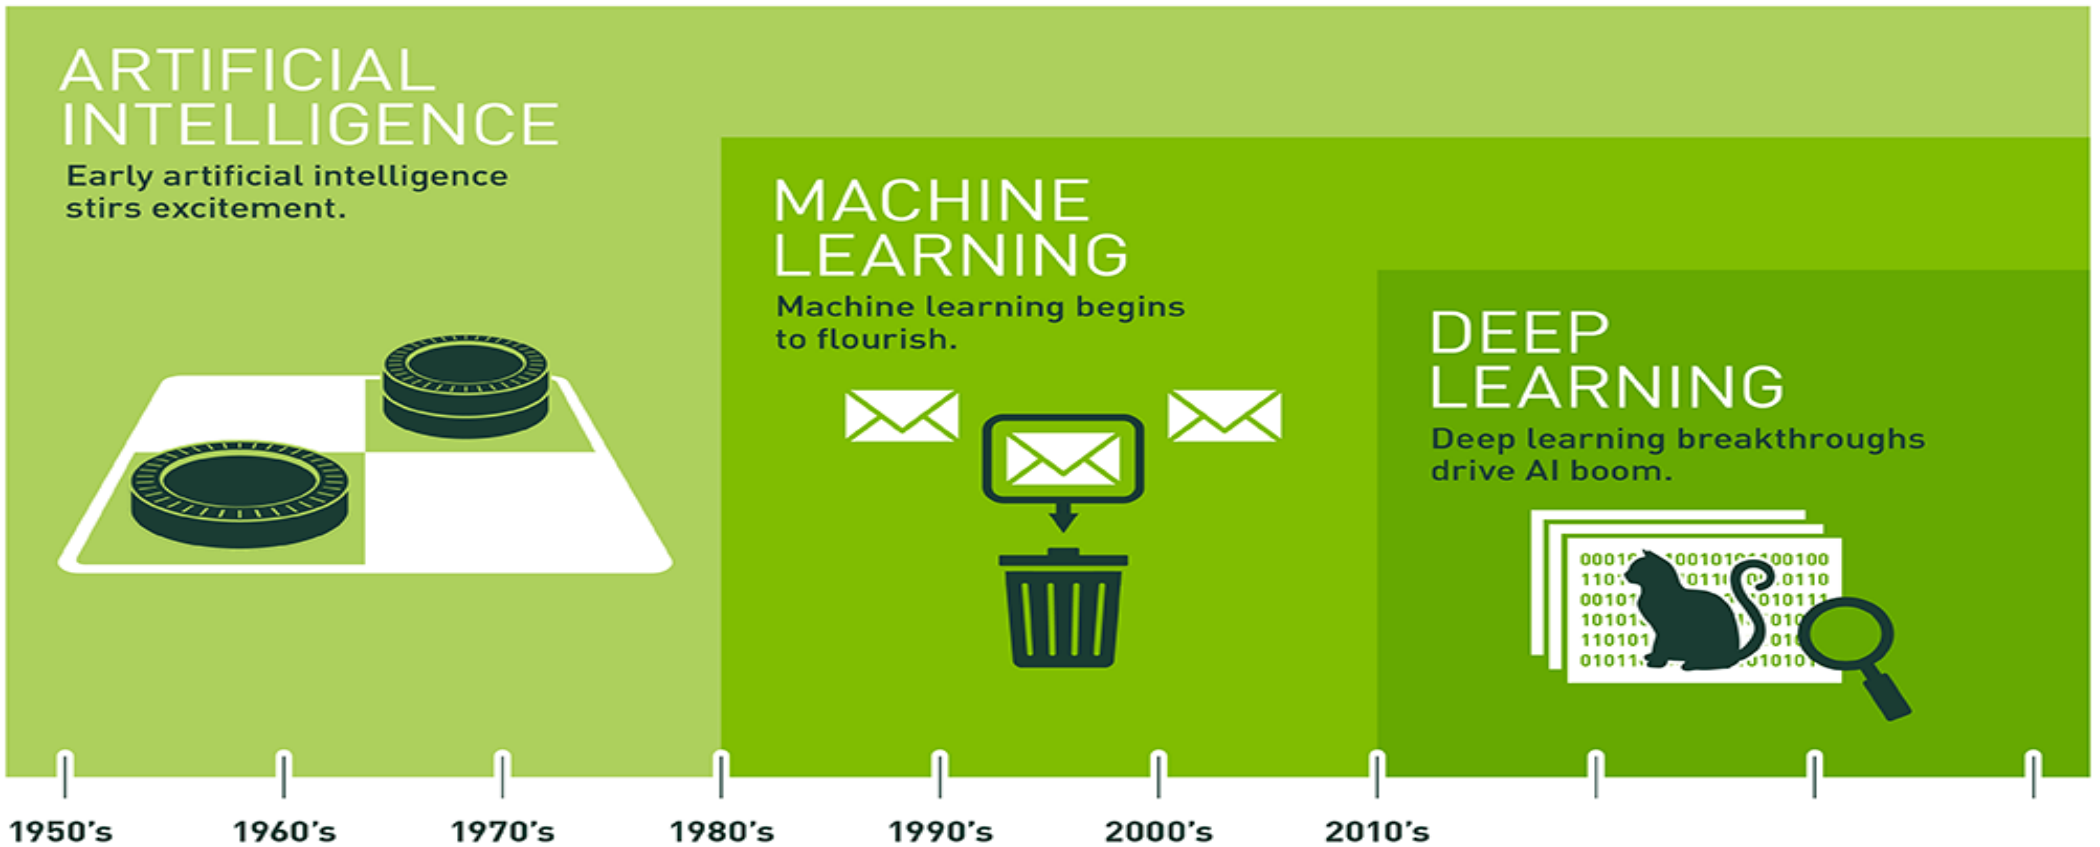

Difference in Workflow

Classic Machine Learning [1990 : now]

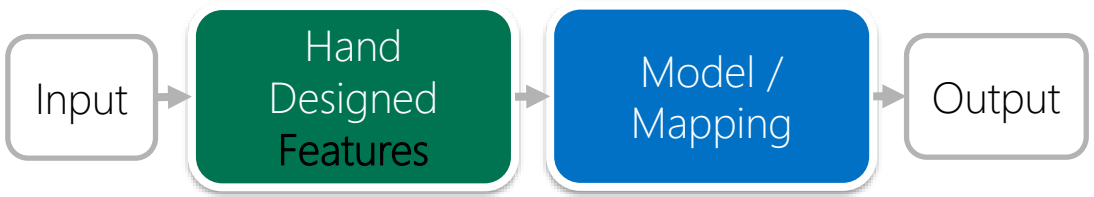

Examples [Regression and SVMs]

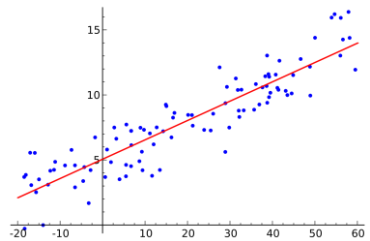

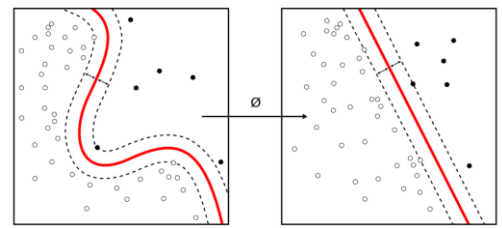

Deep/End-to-End Learning [2012 : now]

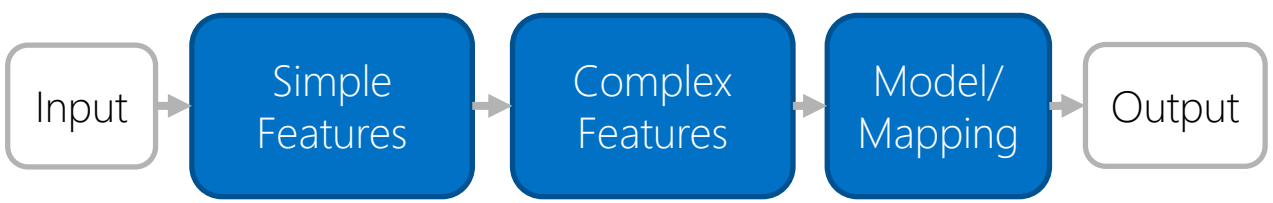

Example [Conv Net]

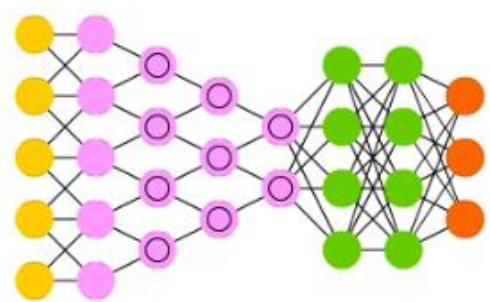

Traditional Workflow

Classic Machine Learning [1990 : now]

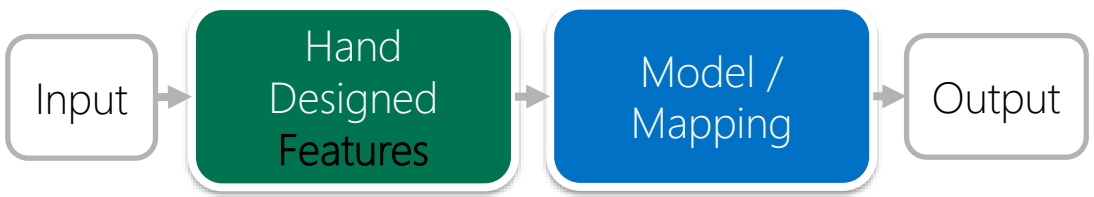

Examples [Regression and SVMs]

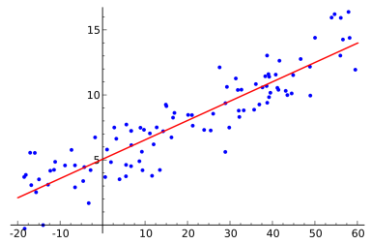

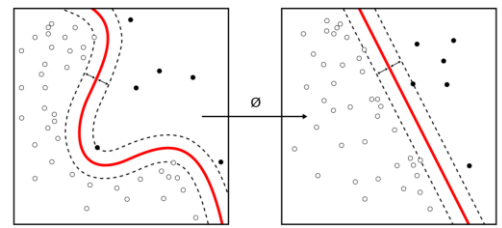

Challenge in Slack channel: How would you describe this image to someone (or something) blind?

Difficult: From it's raw pixels.
Medium: From geometric primitives (lines, curves, colors)
Easy: Using any words that you may know

Deep Learning Workflow

Experience: Trust Neural Network to learn features and model by providing inputs and outputs.

Key Skill: Experience (data) creation

Deep/End-to-End Learning [2012 : now]

Example [Conv Net]

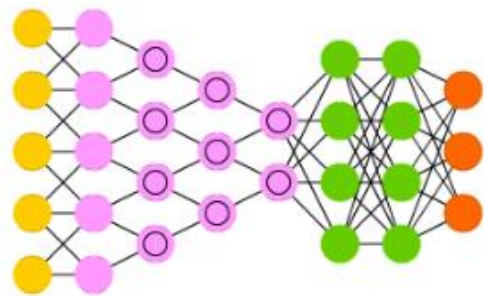

NVIDIA’S DIGITS

NVIDIA'S DIGITS

Interactive Deep Learning GPU Training System

- Simplifies common deep learning tasks such as:
  - Managing data
  - Designing and training neural networks on multi-GPU systems
  - Monitoring performance in real time with advanced visualizations
- Completely interactive so data scientists can focus on designing and training networks rather than programming and debugging
- Open source

NVIDIA'S DIGITS

Interactive Deep Learning GPU Training System

Process Data

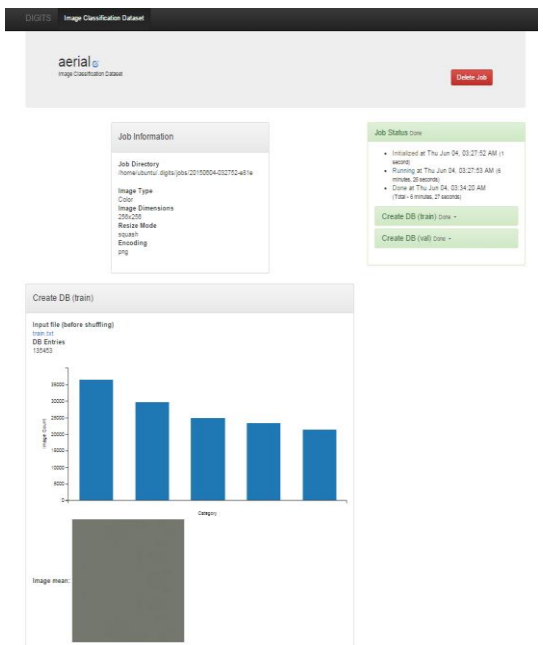

Configure DNN

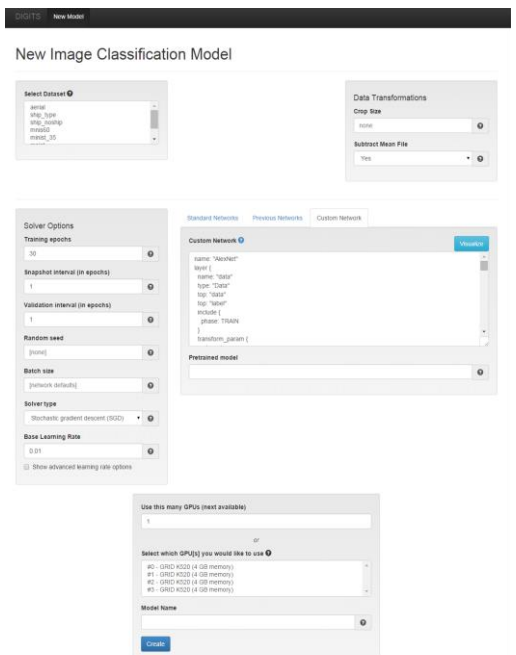

Monitor Progress

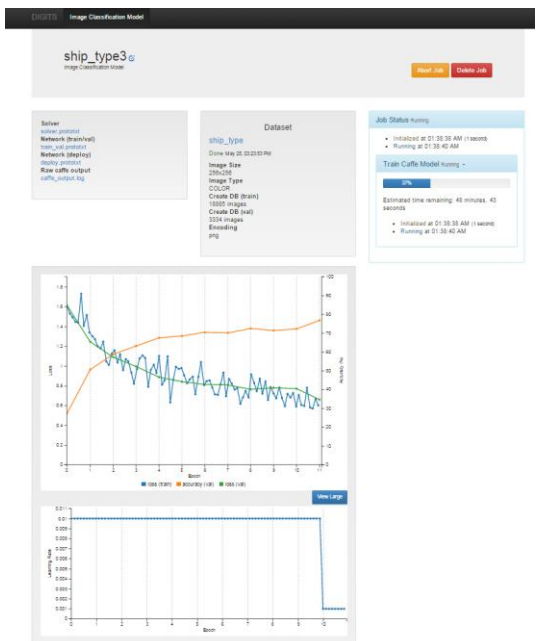

Visualization

DIGITS - MODEL

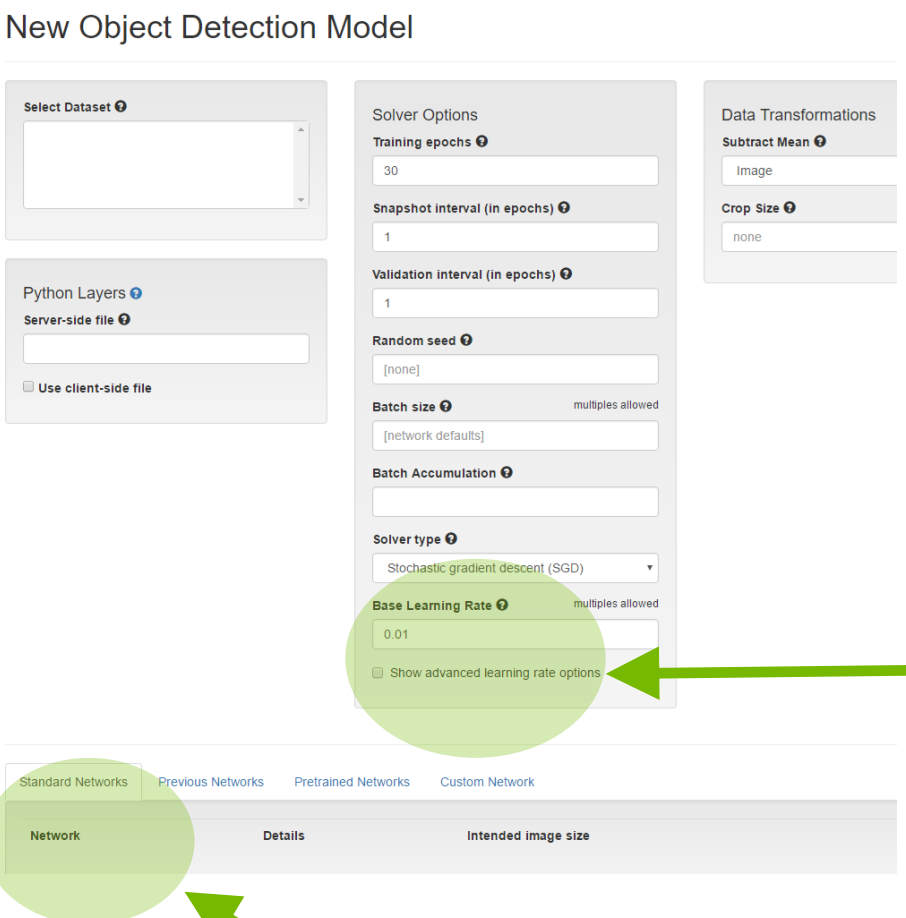

Define custom layers with Python

Can anneal the learning rate

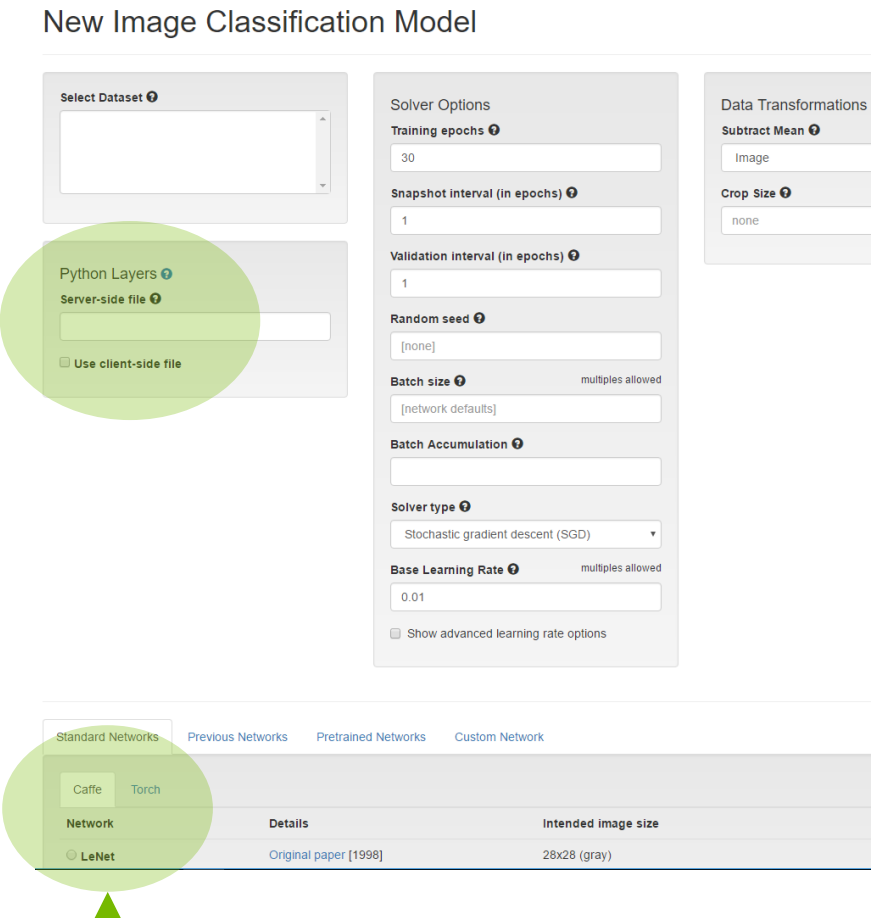

Differences may exist between model tasks

DIGITS – VISUALIZATION RESULTS

Summary

Output visualizations

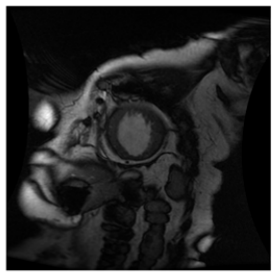

Layer visualizations

Description	Statistics	Visualization
"data" Activation	**Data shape:** [1 256 256] **Mean:** 3.27138 **Std deviation:** 75.5979	
"conv1" Weights *(Convolution layer)* 11,712 learned parameters	**Data shape:** [96 1 11 11] **Mean:** 0.0 **Std deviation:** 0.0	

ENHANCING IMAGES WITH AN AI AUTOENCODER

A great candidate for Deep Learning!

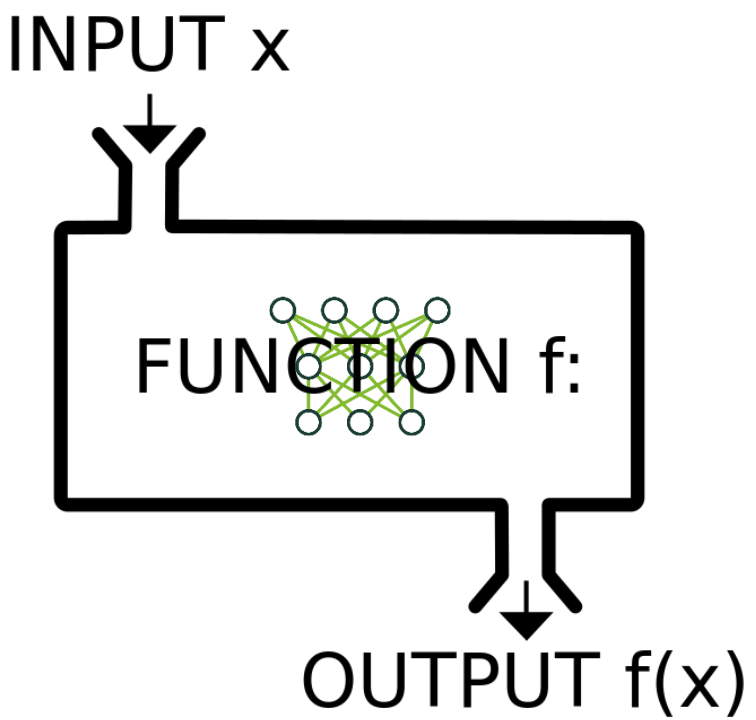

Training Set of images.

1 sample per pixel

- It requires pairs of noisy and noise-free images. The network will learn to remove the noise from the images.
- We can then deploy this trained model to any image we want to denoise. (inference)

Deep Learning for Image Denoising

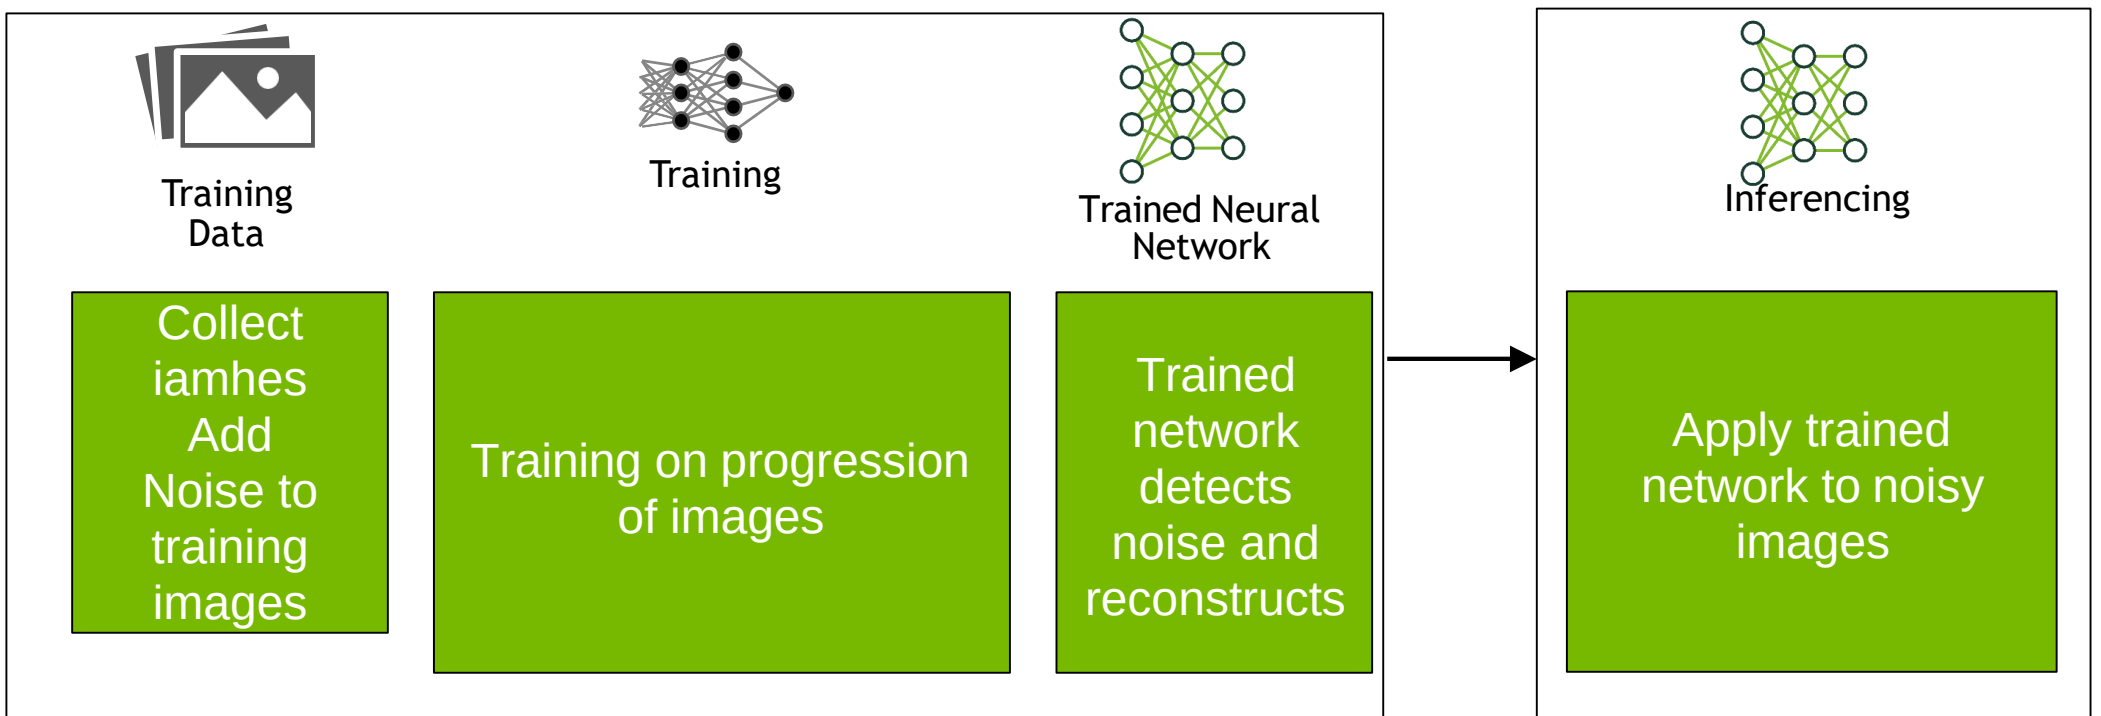

Learning about images (CNN)

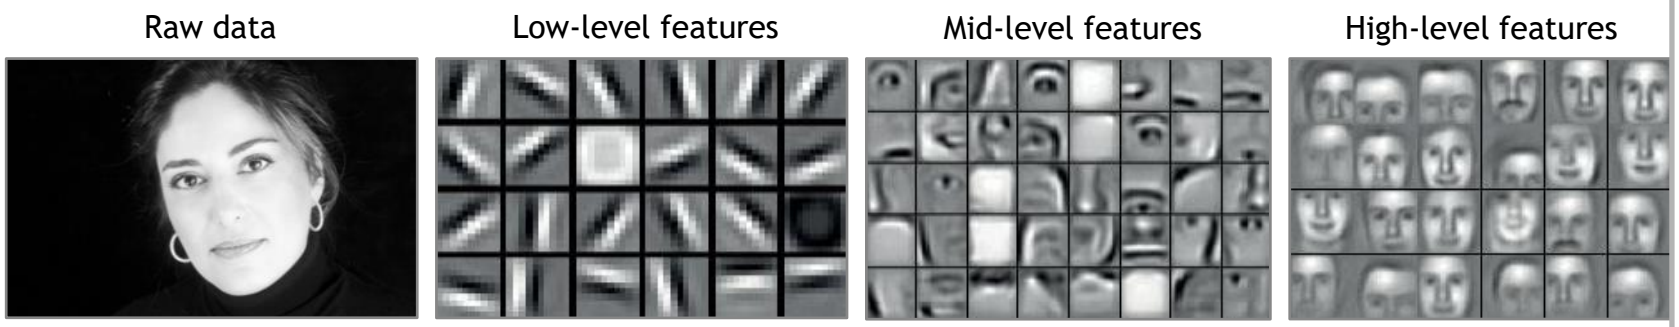

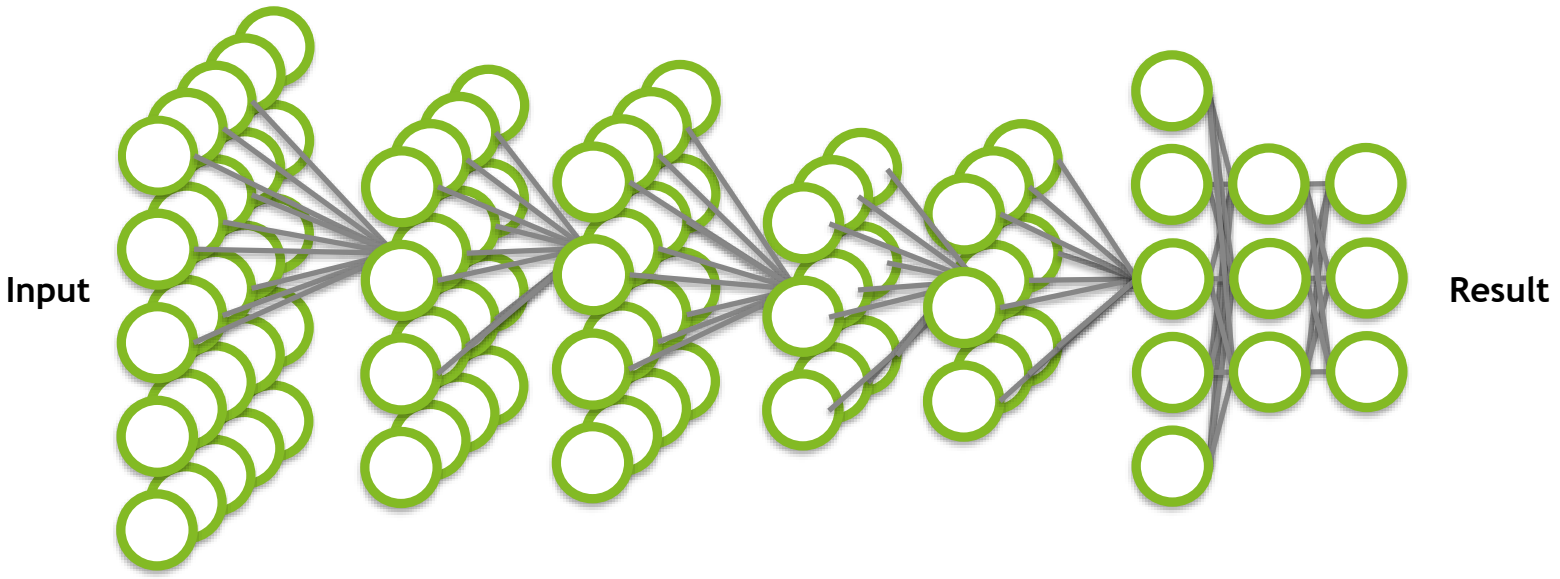

Application components:

Task objective
e.g. Identify face

Training data
10-100M images

Network architecture
~10s-100s of layers
1B parameters

Learning algorithm
~30 Exaflops
1-30 GPU days

Autoencoder – in Action

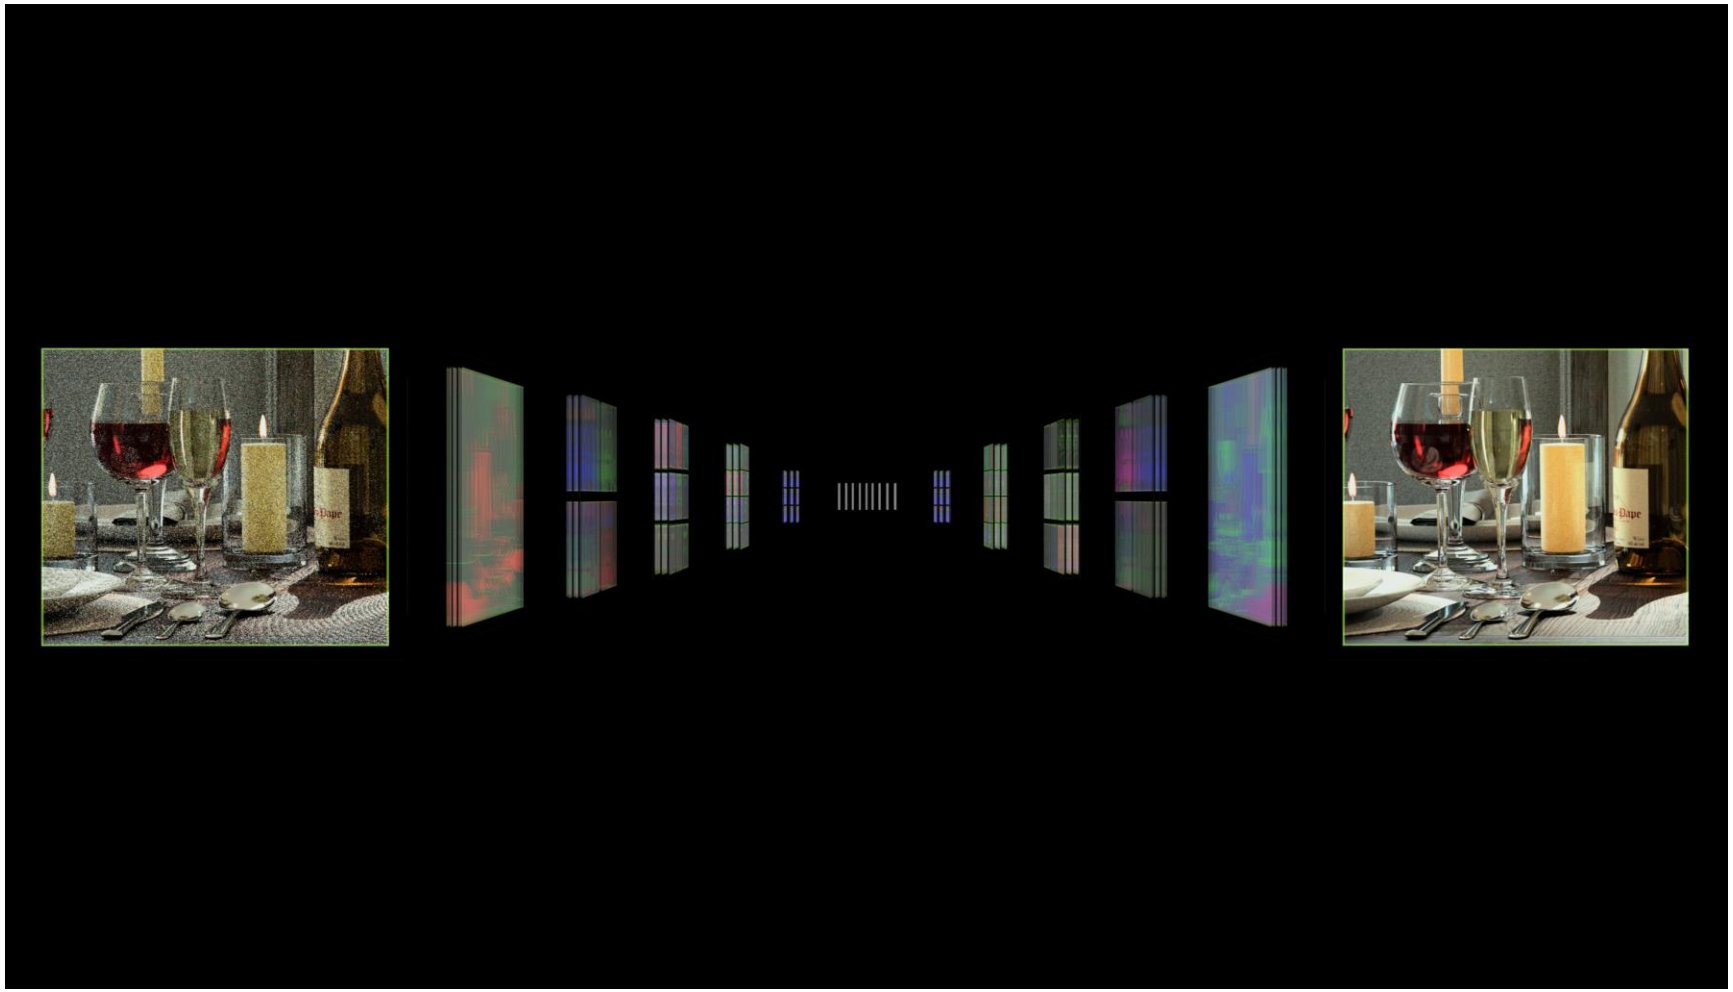


HERMANNVS VENRAED
PARS AV
STRALIS

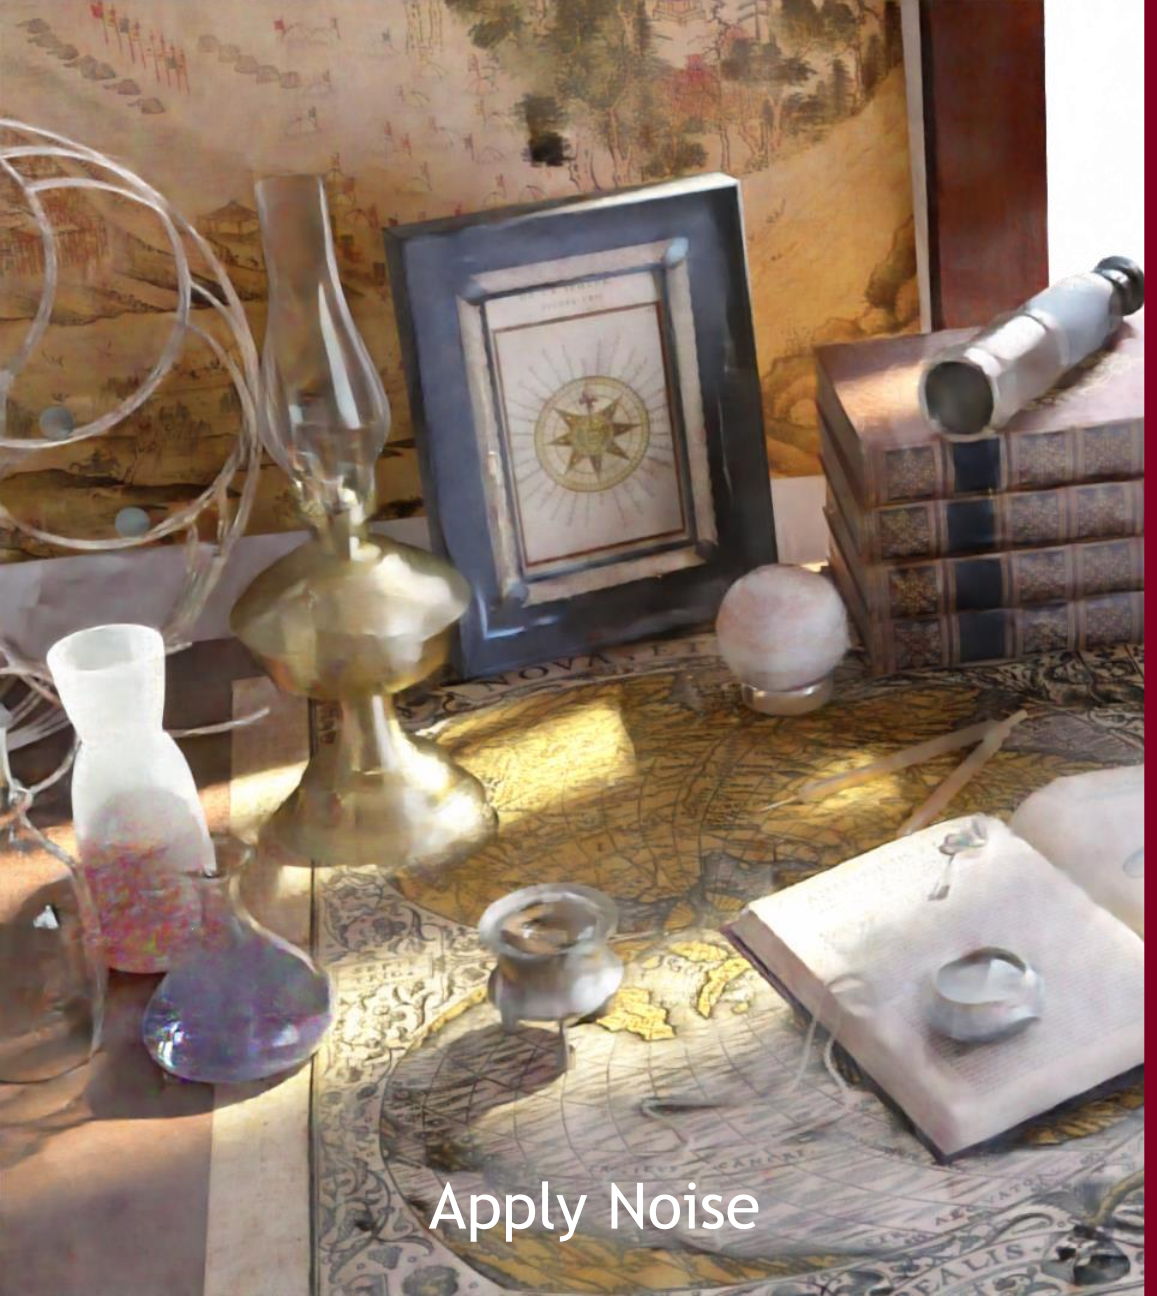
Apply Noise

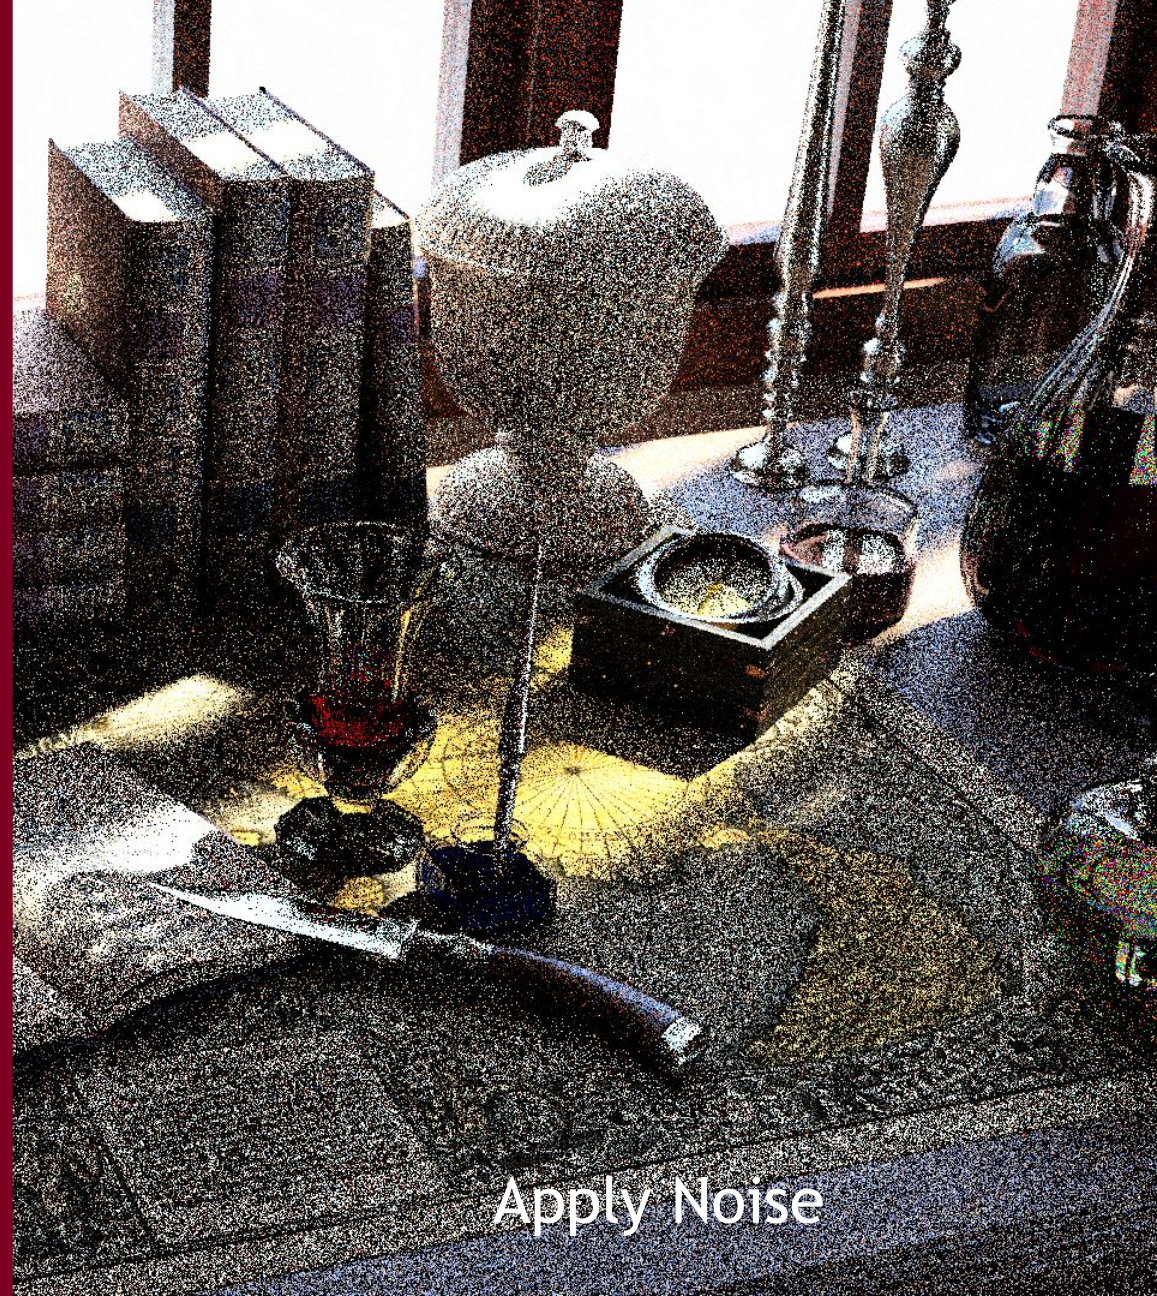
Apply Noise


Provide image to autoencoder enhance


HERMANNVS VENRAED
PARS AV
STRALIS

TensorRT

SOFTWARE INFERENCING PERFROMANCE EHANCEMENT

NVIDIA DEEP LEARNING SOFTWARE PLATFORM

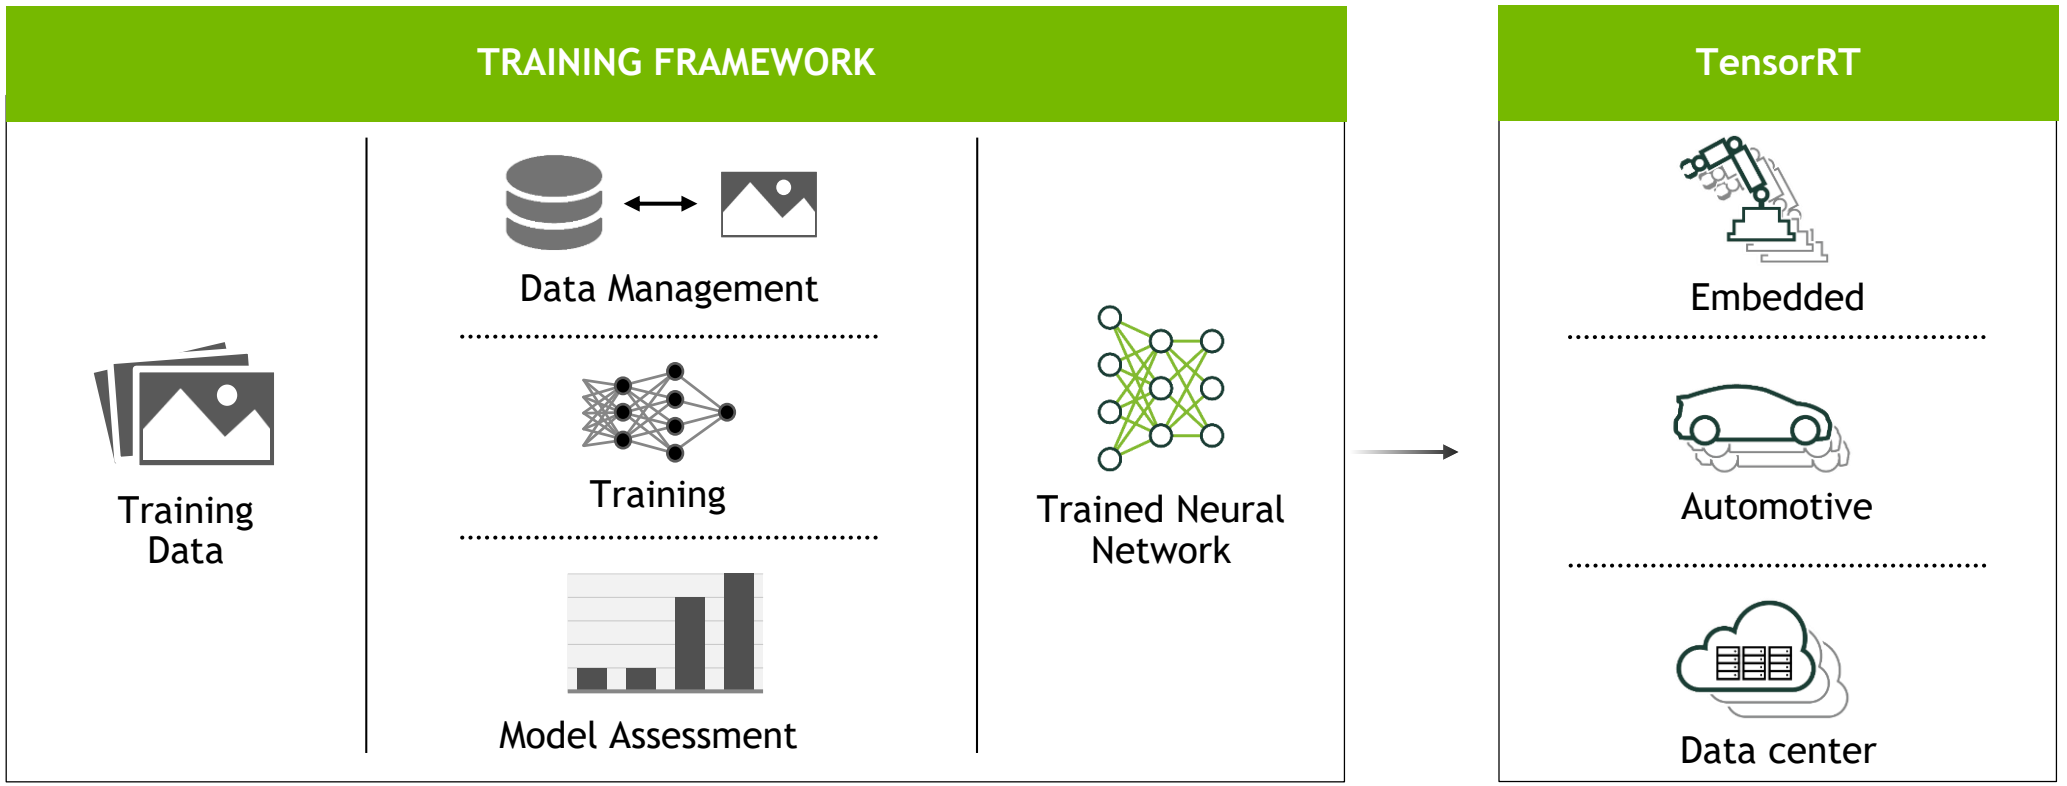

developer.nvidia.com/deep-learning-software

NVIDIA TensorRT

High-performance deep learning inference for production deployment

High performance neural network inference engine for production deployment

Generate optimized and deployment-ready models for datacenter, embedded and automotive platforms

Deliver high-performance, low-latency inference demanded by real-time services

Deploy faster, more responsive and memory efficient deep learning applications with INT8 and FP16 optimized precision support

developer.nvidia.com/tensorrt

Up to 36x More Image/sec

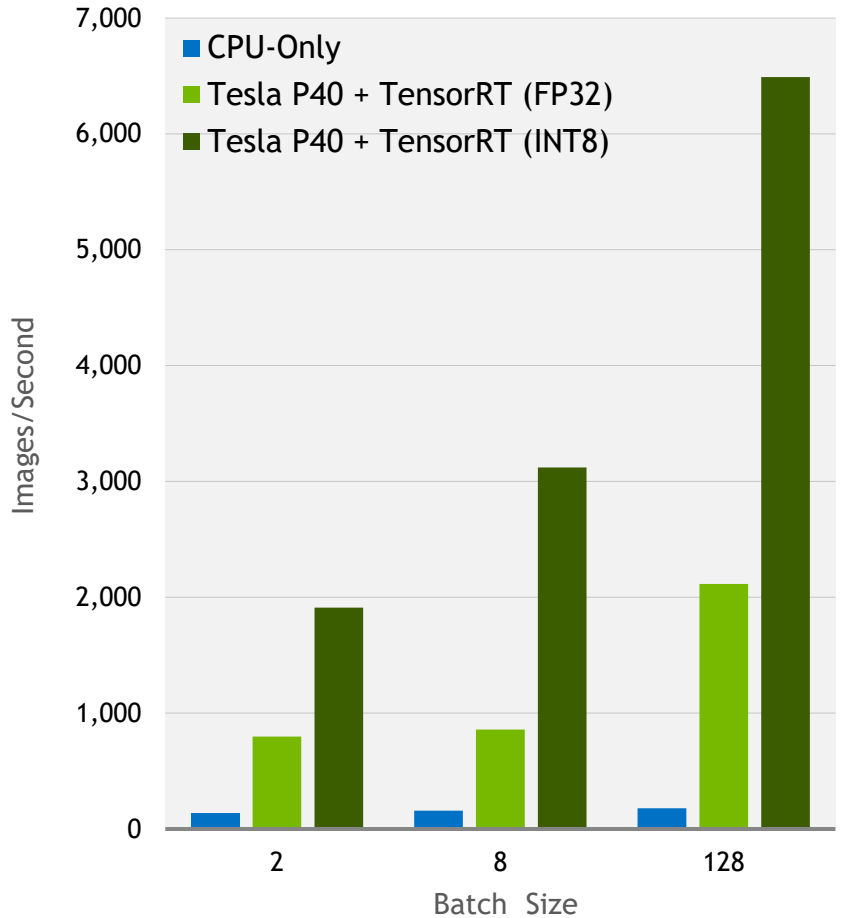

GoogLenet, CPU-only vs Tesla P40 + TensorRT
CPU: 1 socket E4 2690 v4 @2.6 GHz, HT-on
GPU: 2 socket E5-2698 v3 @2.3 GHz, HT off, 1 P40 card in the box

TENSORRT

Networks Supported

- Image Classification (AlexNet, GoogleNet, VGG, ResNet)
- Object Detection
- Segmentation

Not Yet Supported

- RNN/LSTM
- 3D convolutions
- Custom user layers

TENSORRT

Layers Types Supported

- **Convolution:** Currently only 2D convolutions
- **Activation:** ReLU, tanh and sigmoid
- **Pooling:** max and average
- **Scale:** similar to Caffe Power layer (shift+scale*x)^p
- **ElementWise:** sum, product or max of two tensors
- **LRN:** cross-channel only
- **Fully-connected:** with or without bias
- **SoftMax:** cross-channel only
- **Deconvolution**

TENSORRT

Workflow

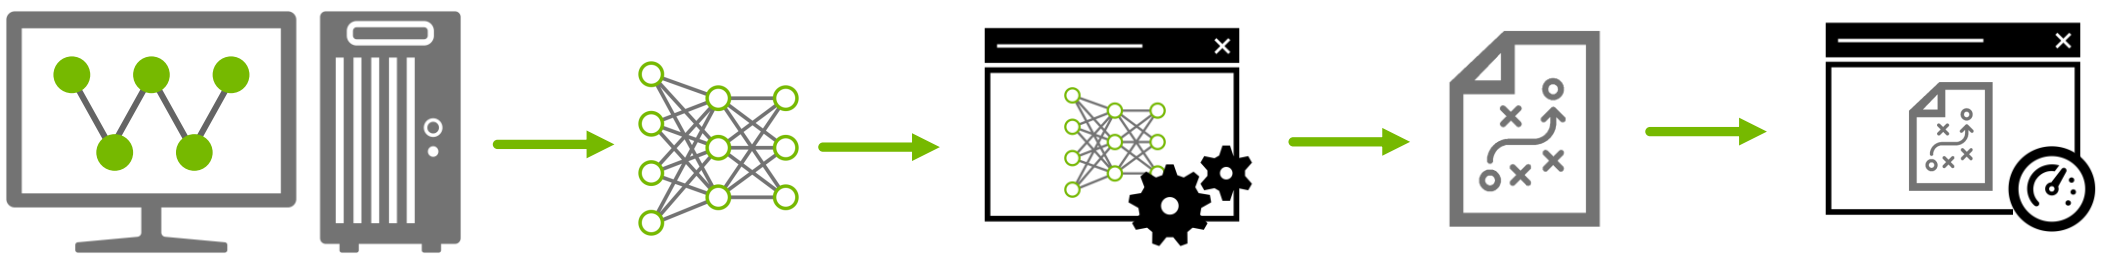

TENSORRT

Optimizations

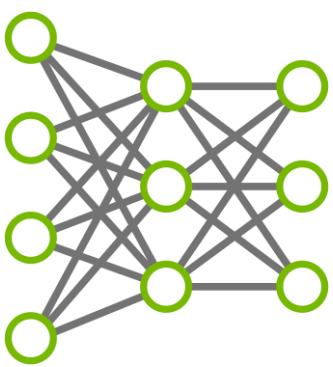

TRAINED
NEURAL NETWORK

- Fuse network layers
- Eliminate concatenation layers
- Kernel specialization
- Auto-tuning for target platform
- Tuned for given batch size

OPTIMIZED
INFERENCE
RUNTIME

GRAPH OPTIMIZATION

Unoptimized network

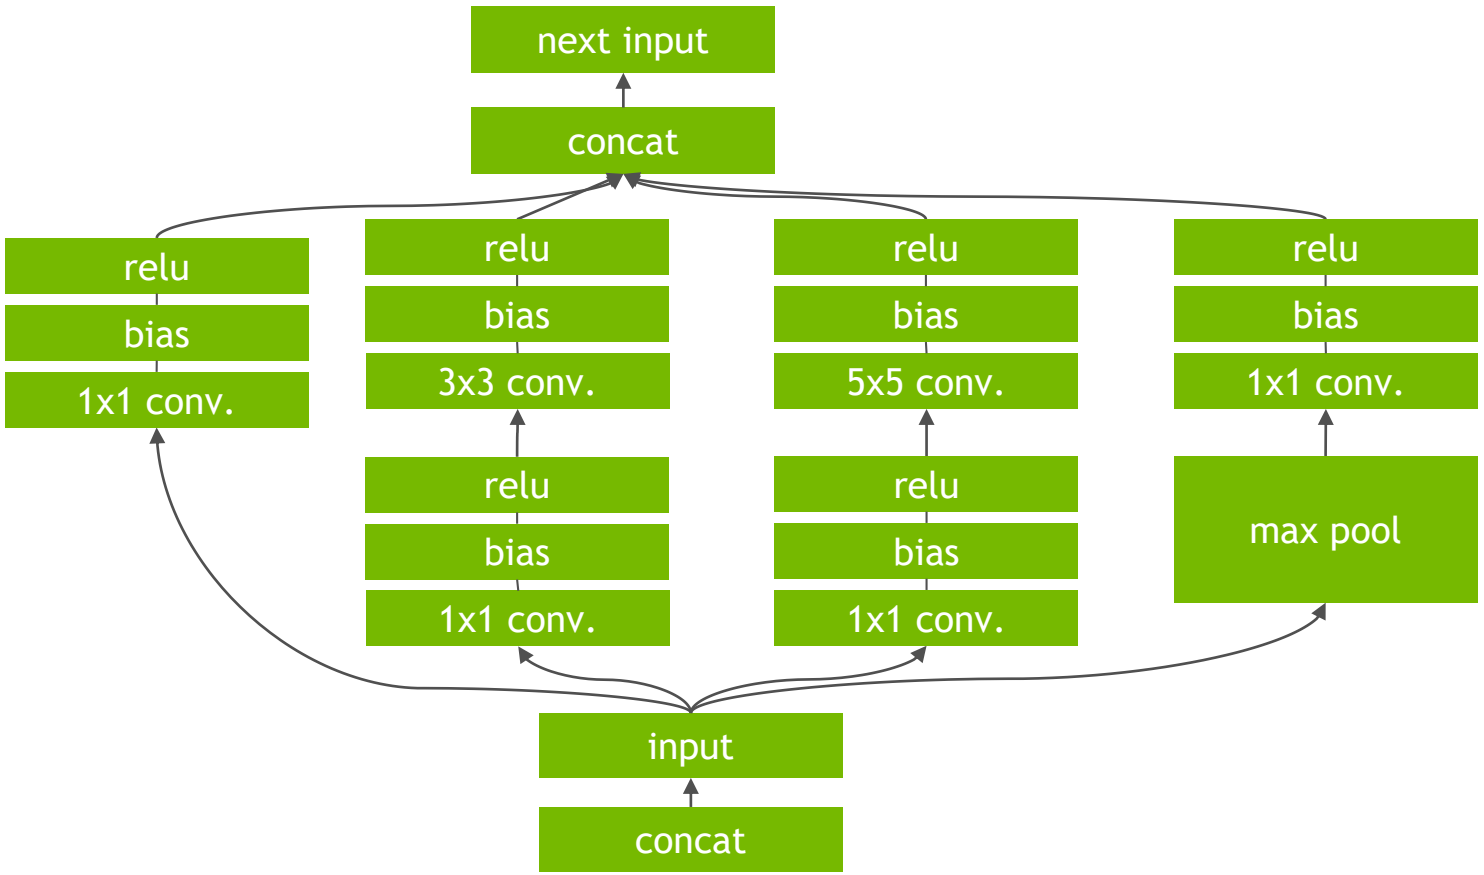

GRAPH OPTIMIZATION

Vertical fusion

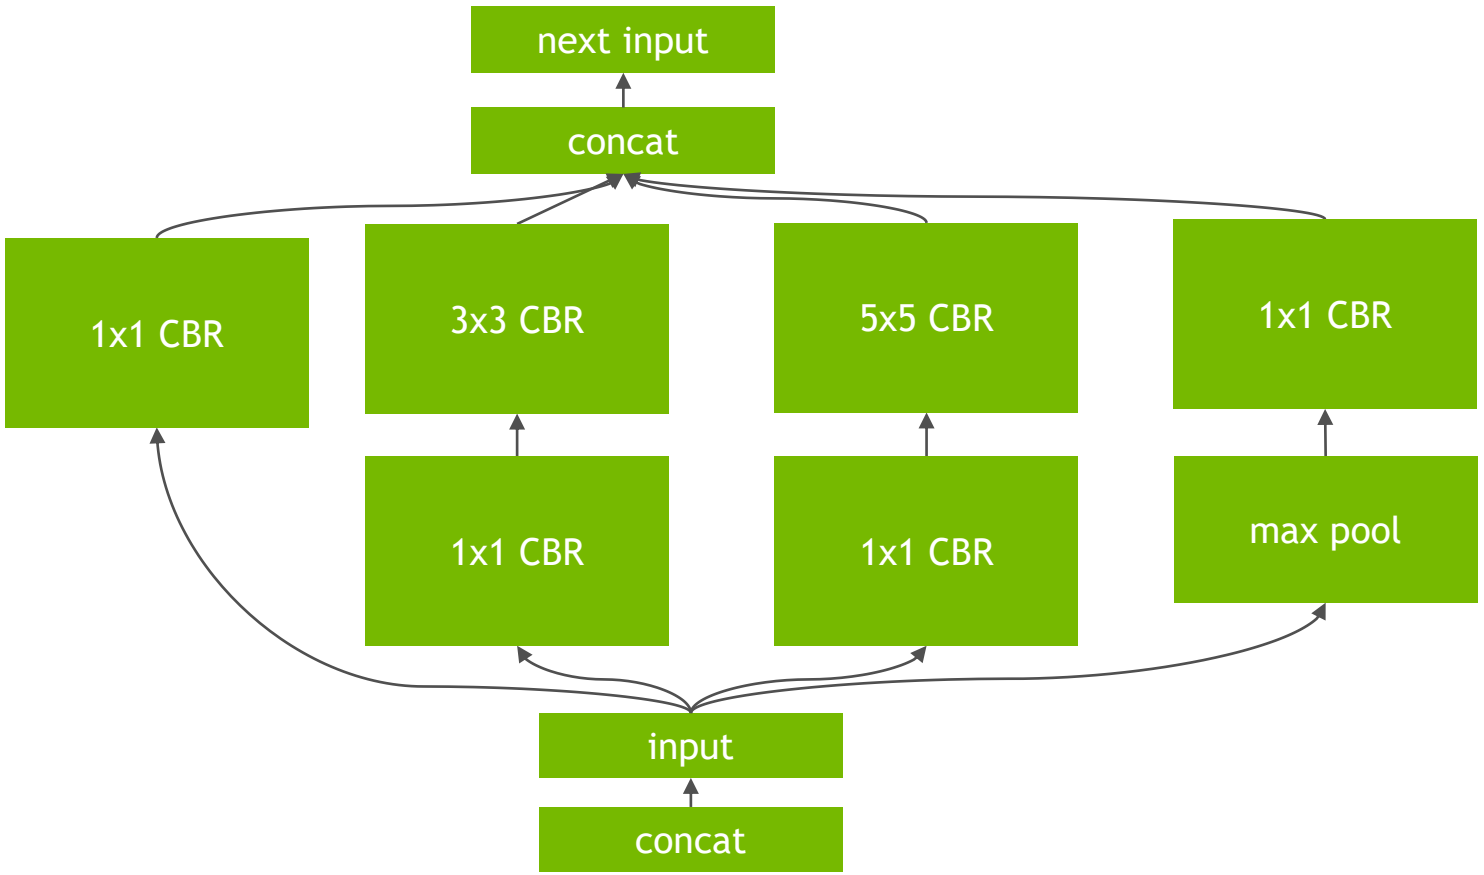

GRAPH OPTIMIZATION

Horizontal fusion

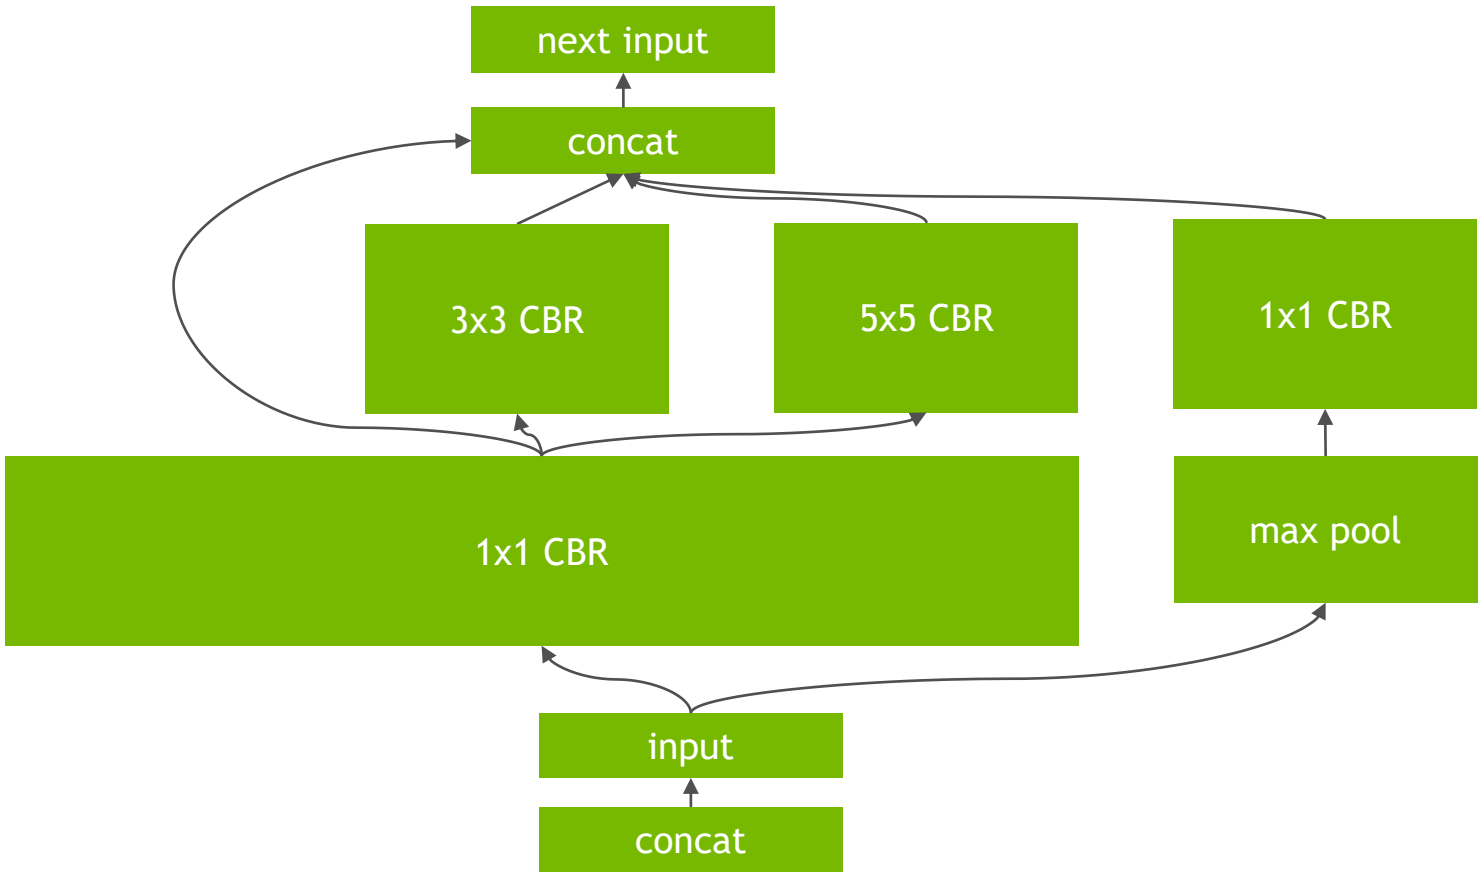

INT8 PRECISION

New in TensorRT

PERFORMANCE

Up To 3x More Images/sec with INT8 Precision

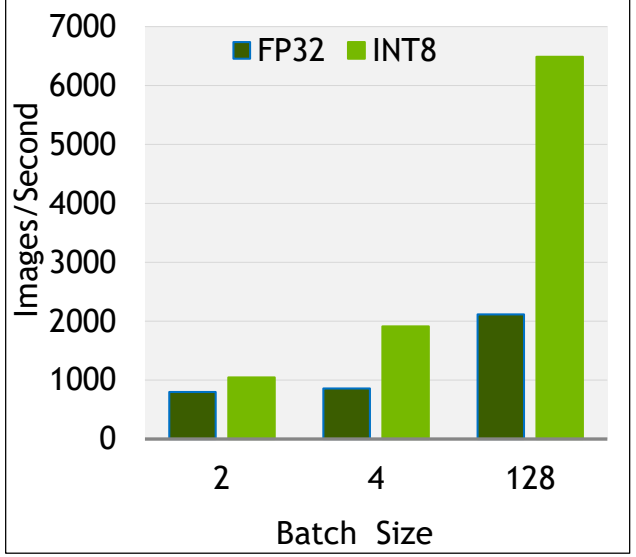

EFFICIENCY

Deploy 2x Larger Models with INT8 Precision

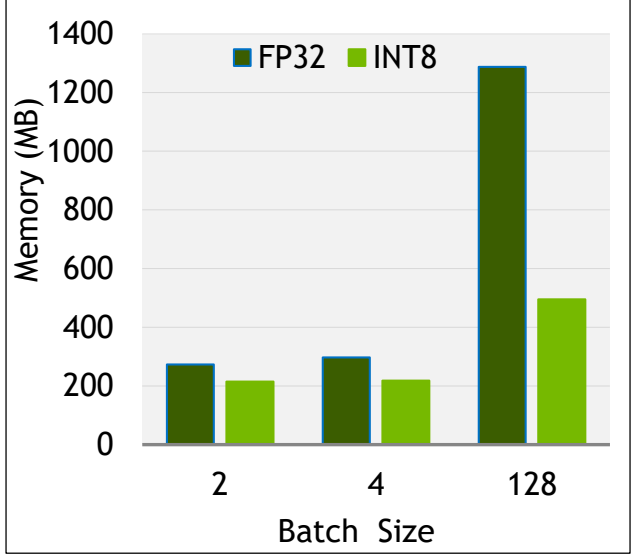

ACCURACY

Deliver full accuracy with INT8 precision

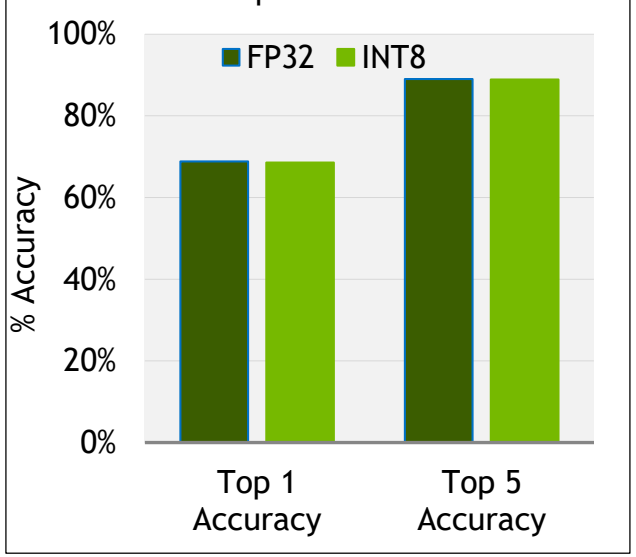

***GoogLenet**, FP32 vs INT8 precision + TensorRT on **Tesla P40 GPU**, 2 Socket Haswell E5-2698 v3@2.3GHz with HT off*

THANK YOU