DEEP LEARNING & ROBOTICS

Alison B Lowndes DL SA & Community Manager

July 2016

NVIDIA

THE BIG BANG IN MACHINE LEARNING

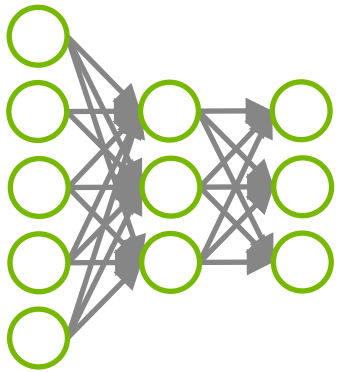

DNN

BIG DATA

GPU

“Google’s AI engine also reflects how the world of computer hardware is changing. (It) depends on machines equipped with GPUs… And it depends on these chips more than the larger tech universe realizes.”

WIRED

THE EXPANDING UNIVERSE OF MODERN AI

"THE BIG BANG"
Big Data
GPU
Algorithms

RESEARCH

Berkeley
Carnegie Mellon University
DEEPMIND
MIT
NYU
OpenAI
Université de Montréal
University of Oxford
University of Toronto

CORE TECHNOLOGY / FRAMEWORKS

facebook — torch
Preferred Networks — Chainer
Université de Montréal — theano
Google — TensorFlow
Berkeley — Caffe
Microsoft — CNTK
University of Oxford
NVIDIA — cuDNN

AI-as-a-PLATFORM

amazon web services
IBM Watson
Google
Microsoft Azure

START-UPS

api.ai
Personal Assistants
conversational interface

BLUE RIVER TECHNOLOGY
Agriculture
crop-yield optimization

clarifai
Tech
visual recognition platform

deep genomics
Genomics
genetic interpretation

drive.ai
Automotive
computer vision

MetaMind
eCommerce & Medical
recommendation engines

Morpho
Tech
computer vision

Orbital Insight
Geospatial
predictions from images

nervana
Tech
AI-as-a-service

SADAKO
Waste Management
sorting robots

SocialEyes
Medical
diabetic retinopathy

Education
teaching robots

1,000+ AI START-UPS
$5B IN FUNDING
Source: Venture Scanner

INDUSTRY LEADERS

Alibaba.com
AstraZeneca
Audi
Baidu 百度
Bloomberg
BMW
charles SCHWAB
cisco
ebay
FANUC ROBOTICS
Ford
GE
gsk
The Home Depot
Massachusetts General Hospital
Mercedes-Benz
MERCK
Pinterest
Schlumberger
Shell
SIEME
TARG
TESL
TOYO
UBE
VOLV
Walma
YAHO
Yand
yelp

DEEP LEARNING EVERYWHERE

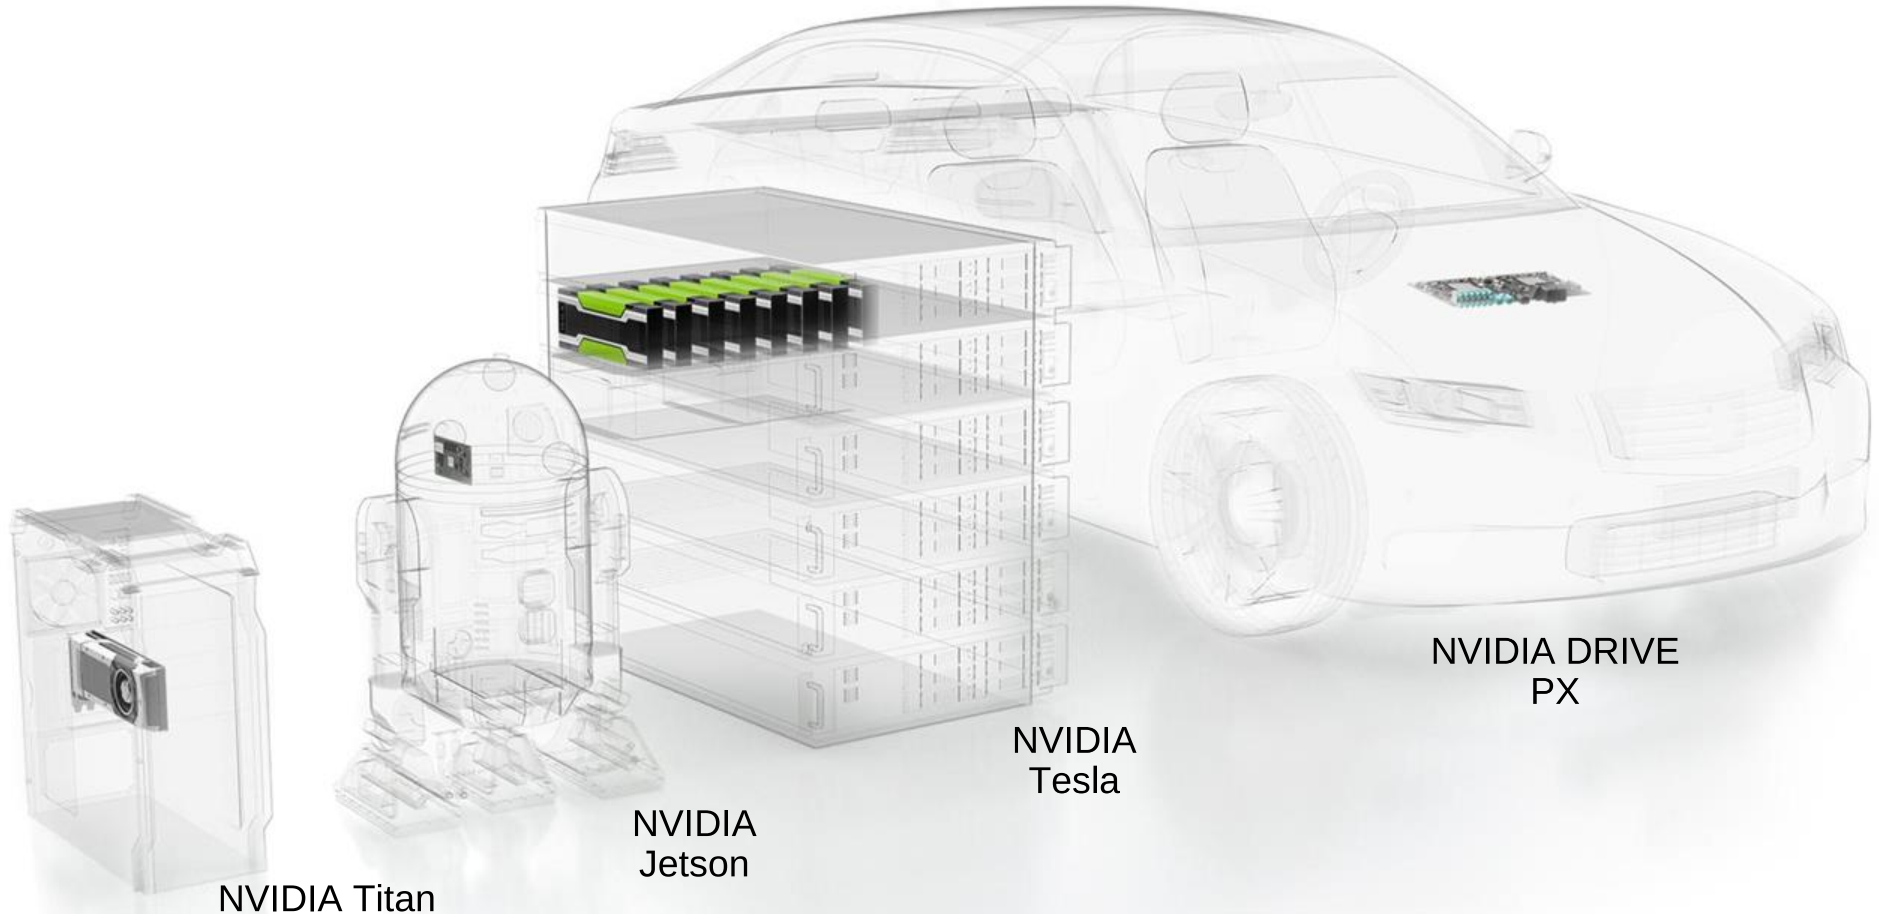

NVIDIA DRIVE PX

NVIDIA Tesla

NVIDIA Jetson

NVIDIA Titan X

NN playground

Visualising a neural network

http://playground.tensorflow.org

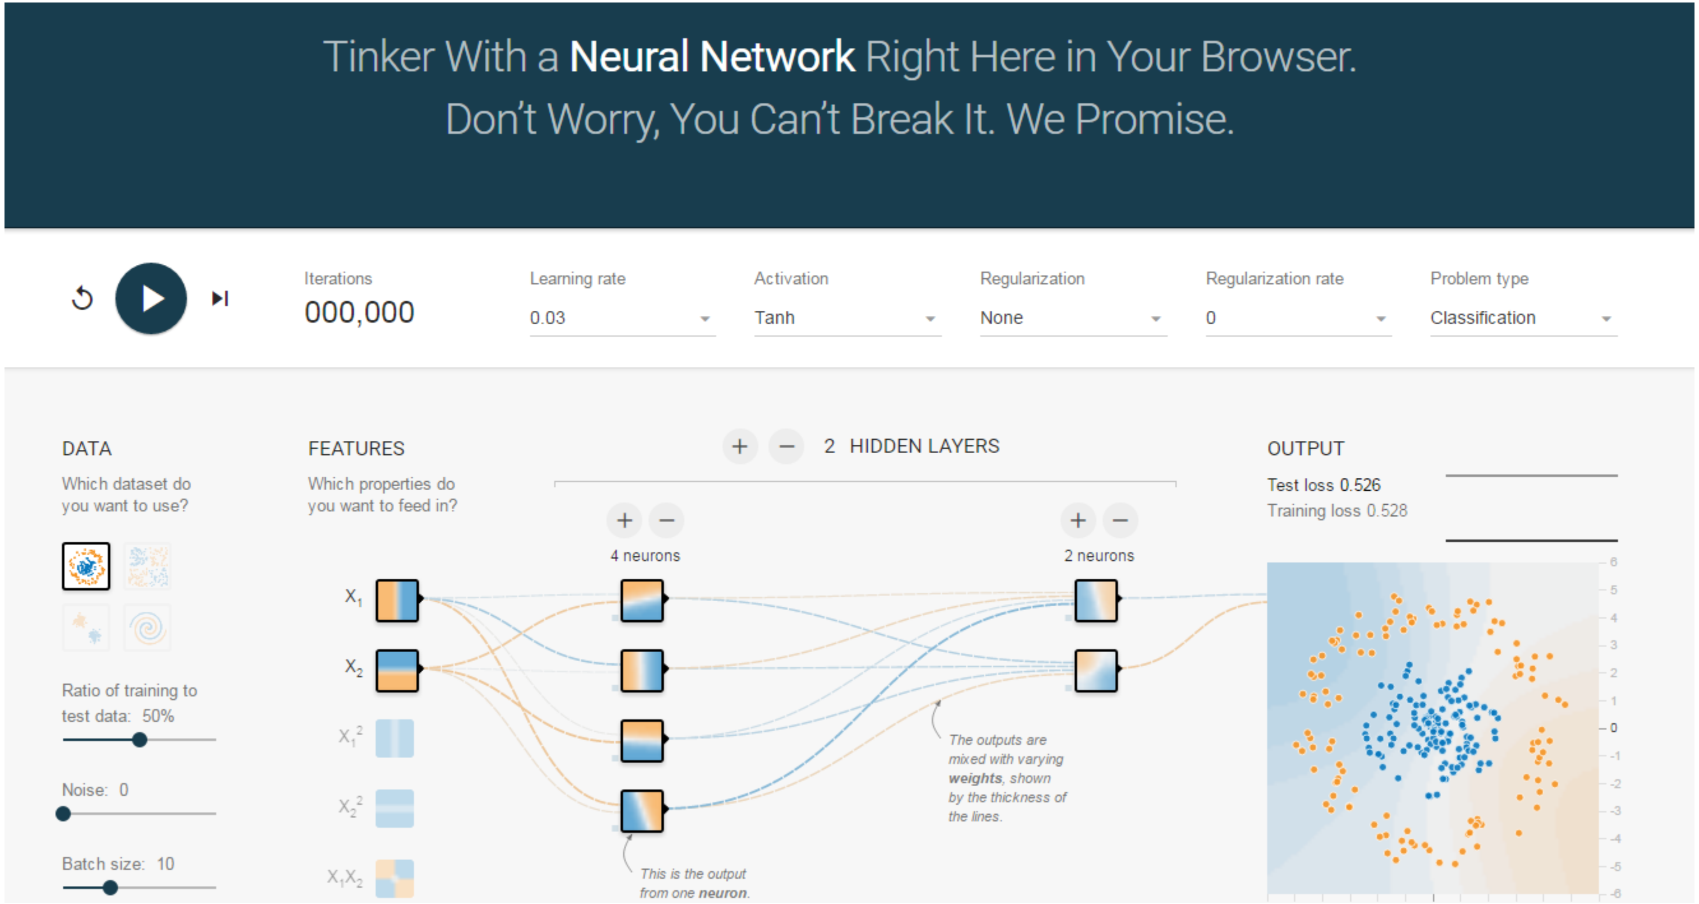

Recurrent neural networks

x – input; **h** – hidden state vector; **y** – output

f – maps input and previous hidden state into new hidden state; *g* – maps hidden state into **y**

f – can be a huge feed-forward network

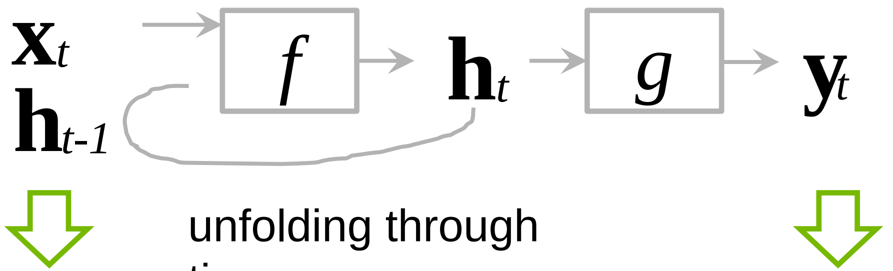

unfolding through time

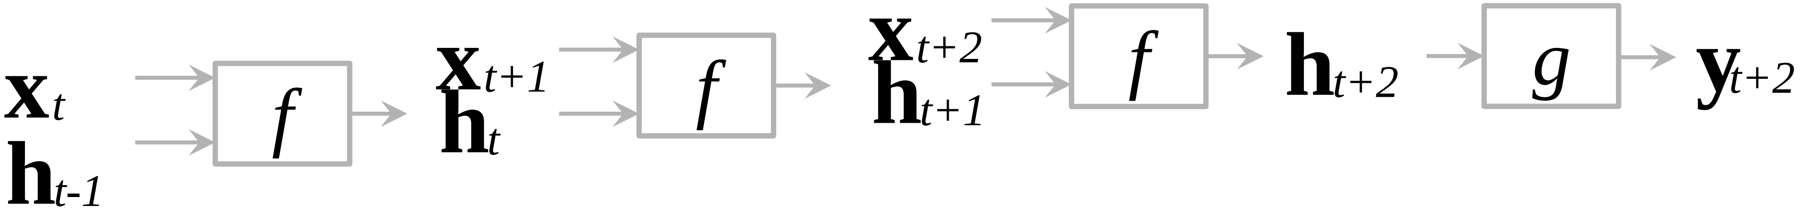

"classical" feed-forward network with shared weights *f*

Most commonly trained by back-propagation

To put *f*'s weight update together:

- sum with **much smaller learning rate**
- averaging

The problem with RNN

Exploding gradient – large increase of the gradient’s norm due to long term dependencies. Results in large increase of the cost function during training.

Address with Rectified-linear (ReLu) activation function

Vanishing gradient – opposite behaviour, long term components go exponentially fast to 0. Results in bad prediction of long term dependencies.

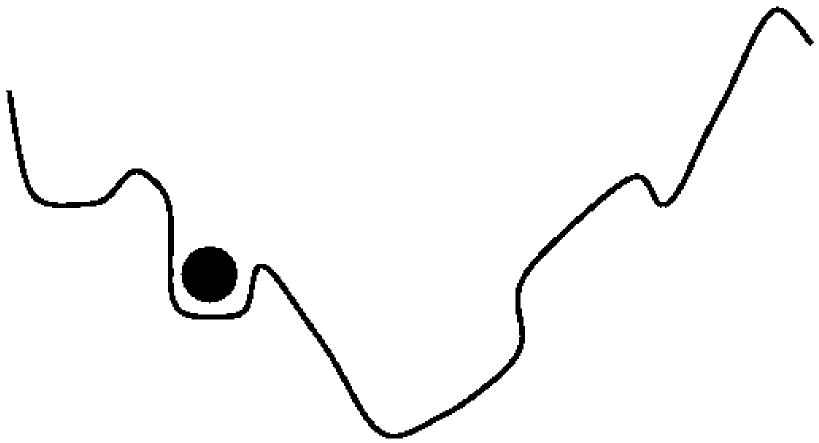

http://www.willamette.edu/

Long short-term memory (LSTM)

Hochreiter (1991) analysed vanishing gradient “*LSTM falls out of this almost naturally*”

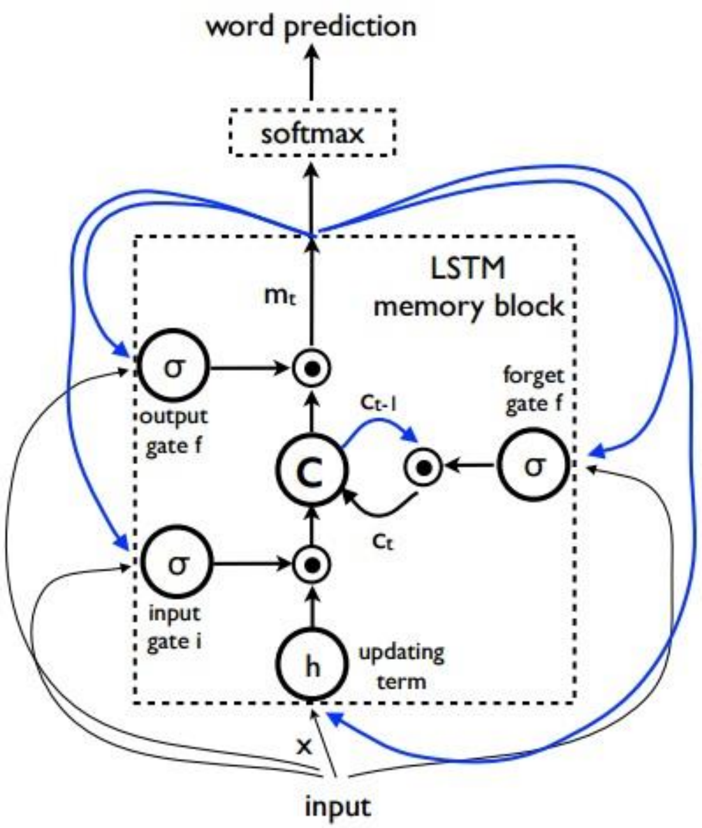

Fig from Vinyals et al, Google April 2015 NIC Generator

Training via backprop unfolded in time

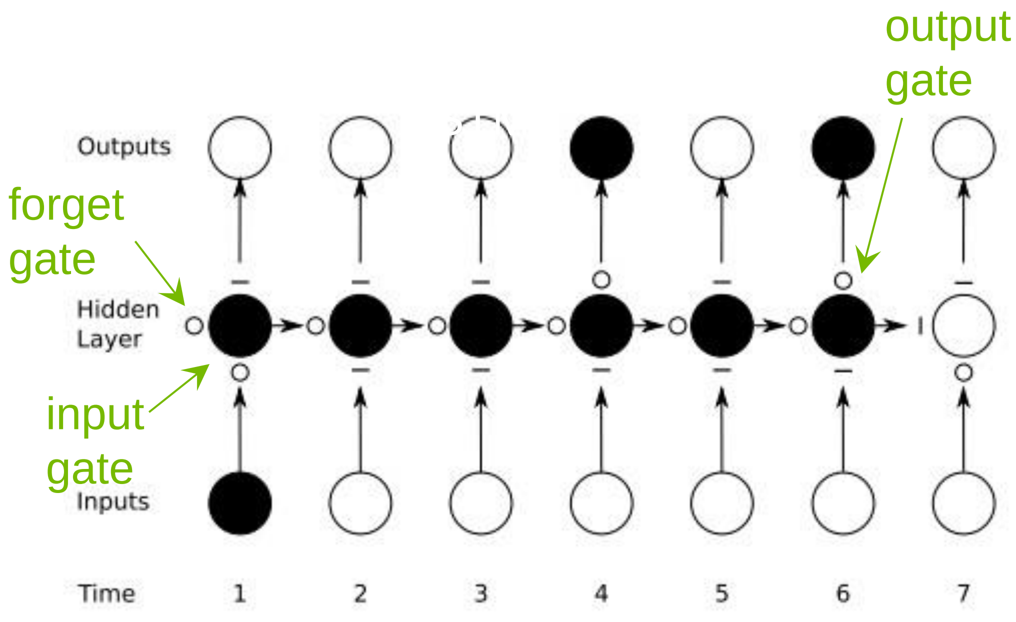

Long time dependencies are preserved until input gate is closed (-) and forget gate is open (O)

Fig from Graves, Schmidhuber et al, Supervised Sequence Labelling with RNNs

Gates control importance of the corresponding activations

Optimising RNNs with cuDNN v5

ParallelForAll

devblogs.nvidia.com/parallelforall/optimizing-recurrent-neural-networks-cudnn-5/

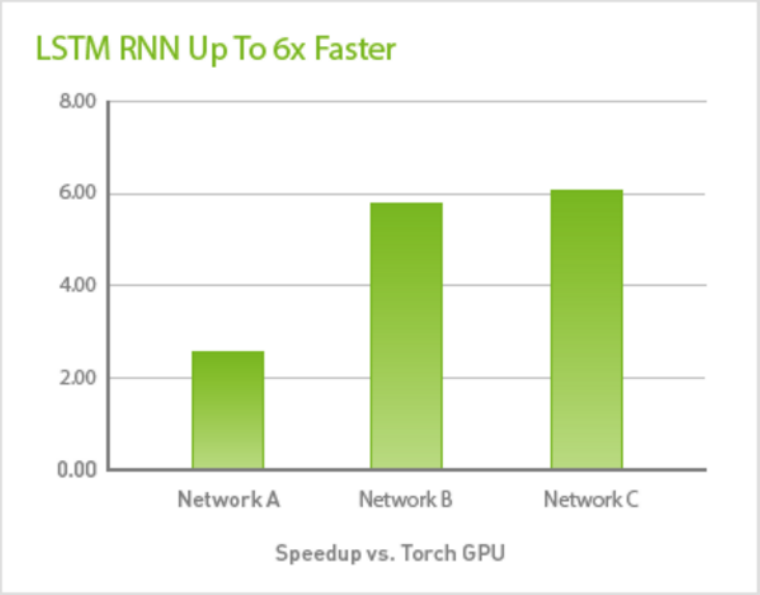

Figure 1: cuDNN 5 + Torch speedup vs. Torch-rnn implementation, M40, Intel® Xeon® Processor E5-2698

Supports:

- ReLU & tanh activation functions
- Gated Recurrent Units (GRU)
- Long Short-Term Memory (LSTM)

REINFORCEMENT LEARNING

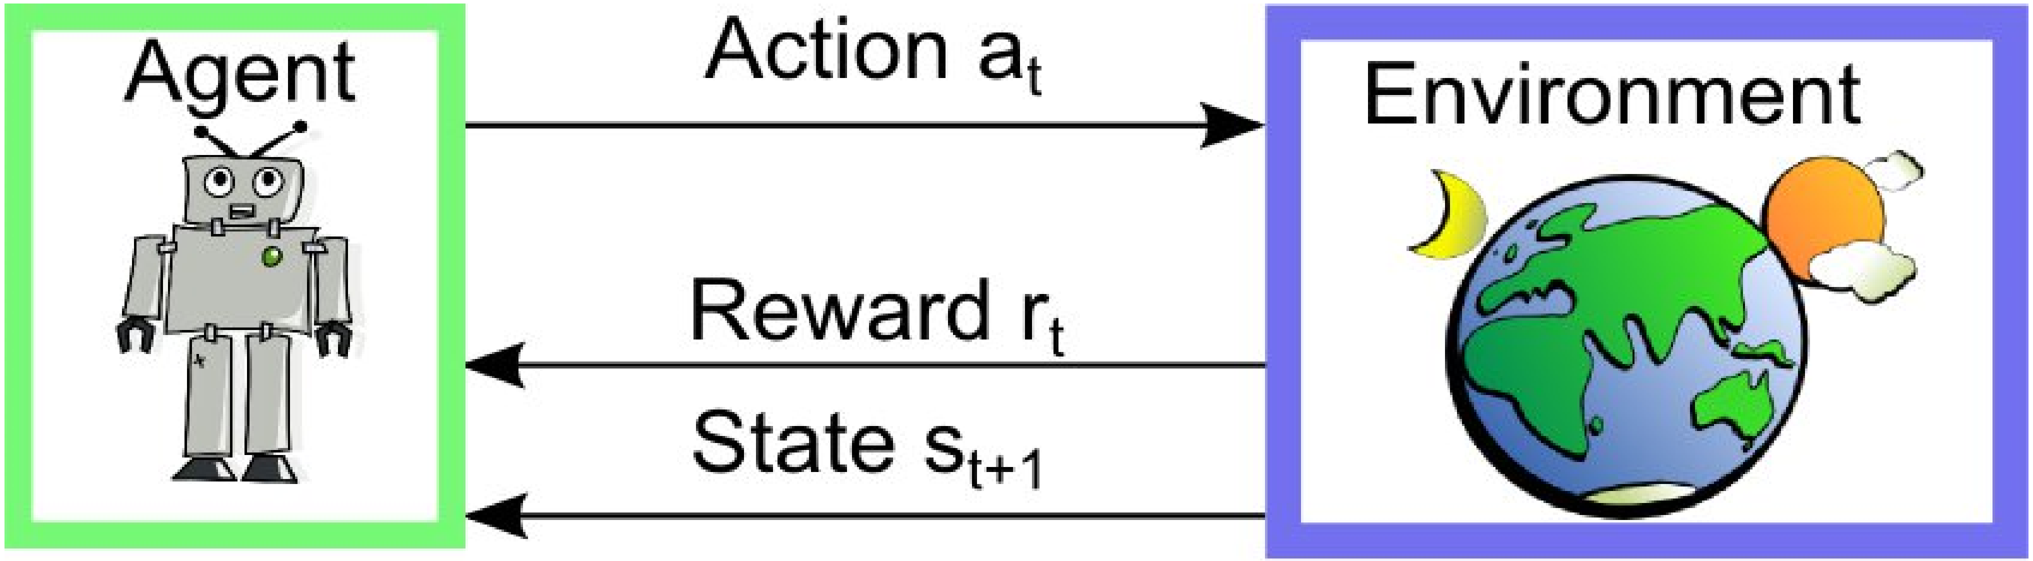

Reinforcement Learning Setup

http://www.ausy.tu-darmstadt.de/

Bellman Principle and Q-function

Q(s; a): the maximum expected return achievable by following any strategy after seeing sequence *s* and taking action *a*

Bellman principle (**dynamic programming**)

if the optimal value **Q(s(t+1); a(t+1))** of the sequence **s** at the next time-step **(t+1)** was known for all possible actions **a'**, then the optimal strategy is to select the action **a'** which maximizes the expected value of

r(t+1) + Q(s(t+1); a(t+1)) => max

DQN: deep Q-learning network

Mastering Breakout

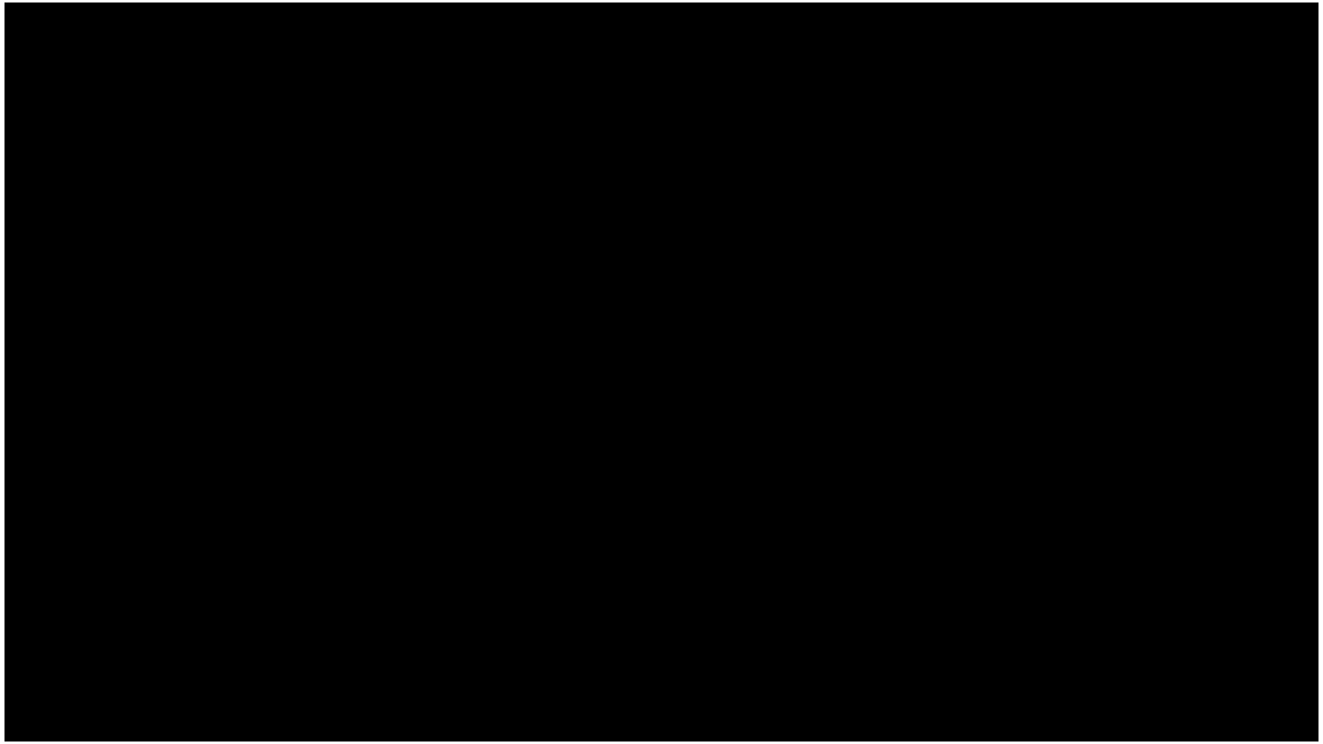

Video: https://www.youtube.com/watch?v=TmPfTpjtdgg

Further resources

- RLPy Framework - Value-Function-Based Reinforcement

- TeachingBox - Java based RL framework
  - http://servicerobotik.hs-weingarten.de/en/teachingbox.php

- BeliefBox - Bayesian reinforcement learning library and toolkit
  - https://code.google.com/p/beliefbox/

- Deep Q-Learning with Tensor Flow
  - https://github.com/nivwusquorum/tensorflow-deepq

GOOGLE DEEPMIND ALPHAGO CHALLENGE

Sergey Levine

http://videolectures.net/iclr2016_levine_deep_learning/
http://arxiv.org/pdf/1603.02199v3.pdf

Pieter Abbeel

gym.openai.com

Video: https://www.youtube.com/watch?v=ATvp0Hp7RUI

END-TO-END PRODUCT FAMILY

HYPERSCALE HPC

Tesla M4, M40

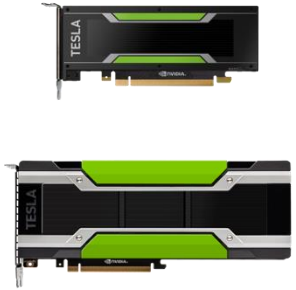

Hyperscale deployment for DL training, inference, video & image processing

MIXED-APPS HPC

Tesla K80

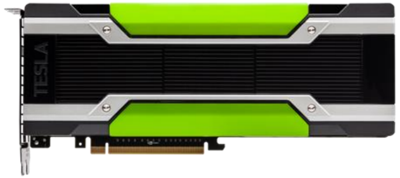

HPC data centers running mix of CPU and GPU workloads

STRONG-SCALING HPC

Tesla P100

Hyperscale & HPC data centers running apps that scale to multiple GPUs

FULLY INTEGRATED DL SUPERCOMPUTER

DGX-1

For customers who need to get going now with fully integrated solution

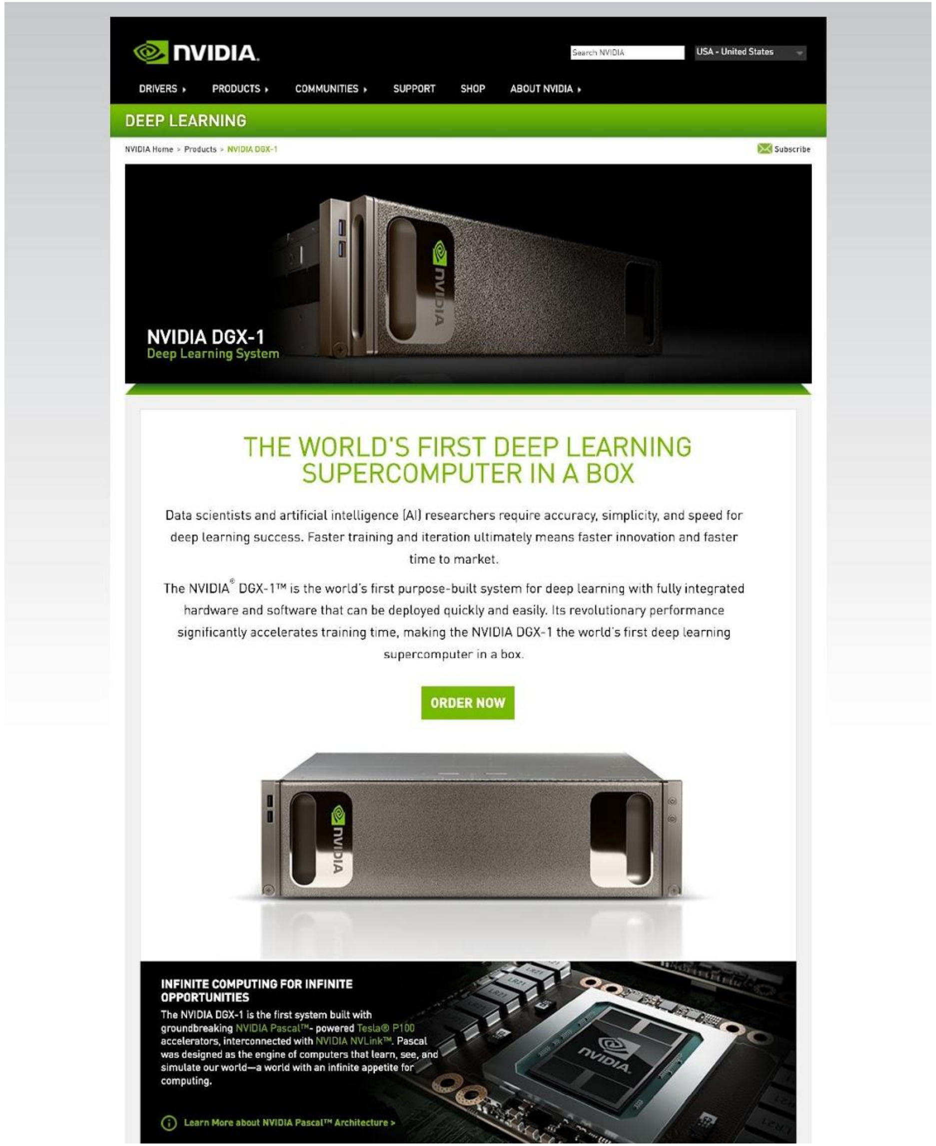

LEARN MORE

nvidia.com/dgx1

nvidia.com/deeplearning

ONLINE CLASSES

Getting started with deep learning

developer.nvidia.com/deep-learning

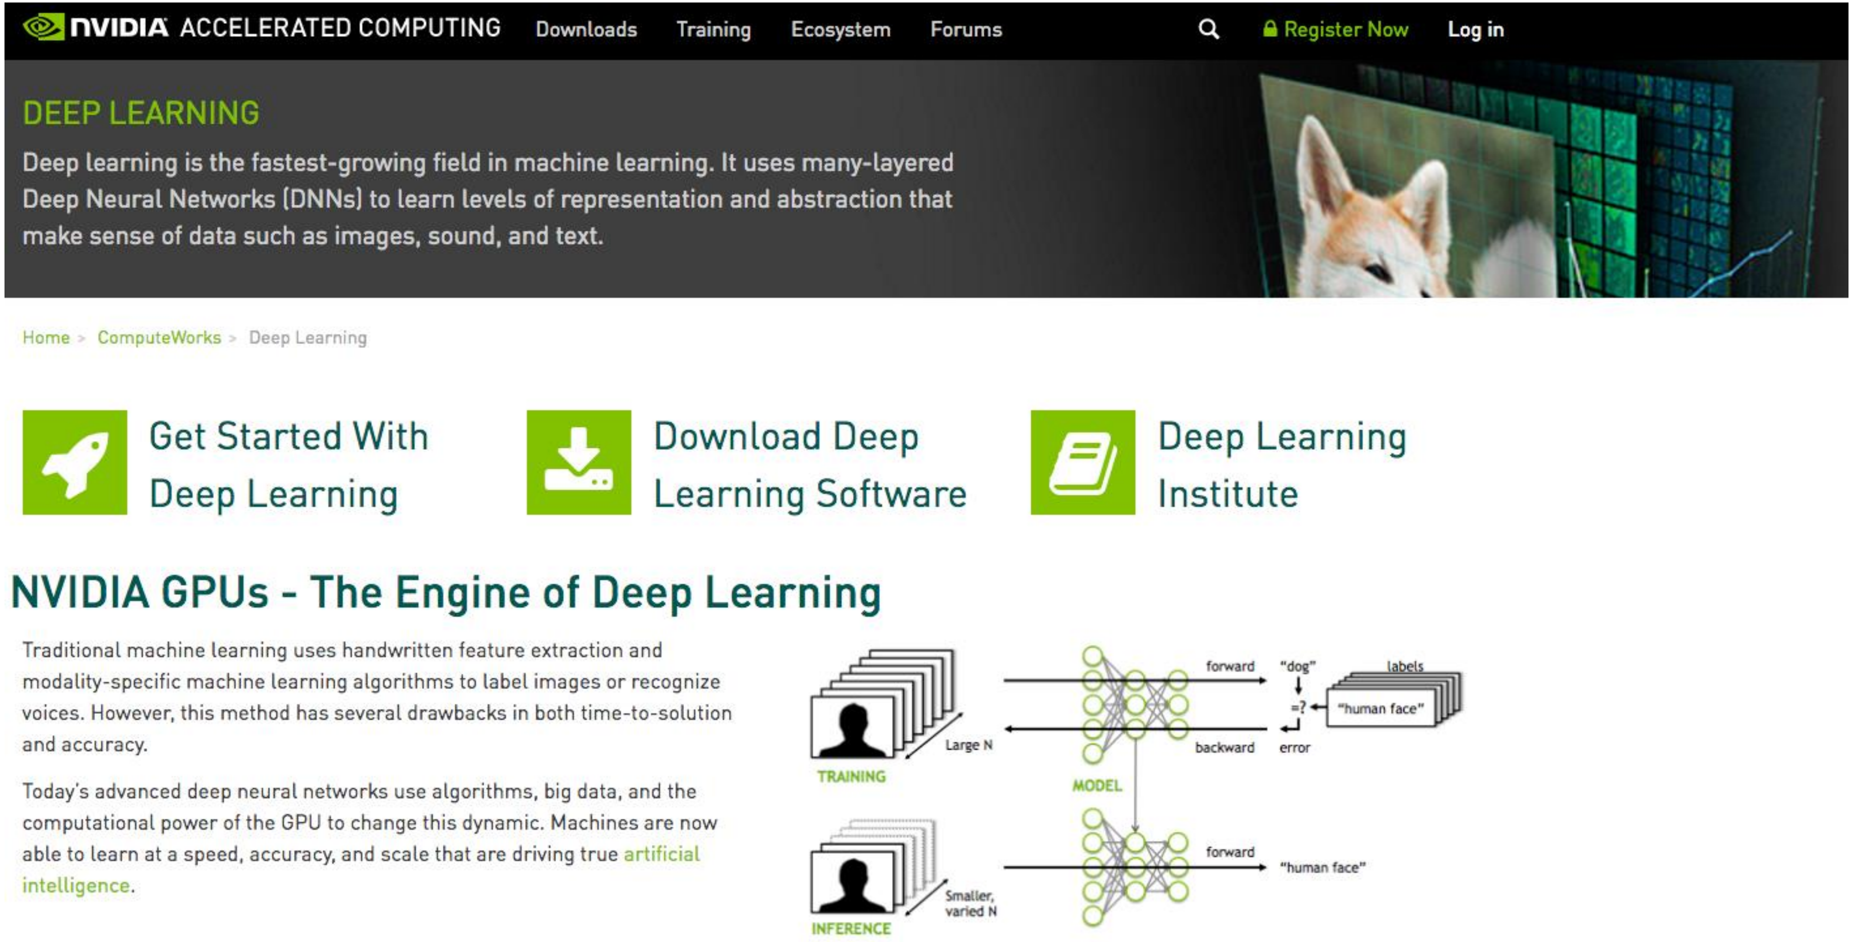

Questions?

alowndes@nvidia.com

NVIDIA