

ARTIFICIAL INTELLIGENCE & DEEP LEARNING



ALISON B LOWNDES




AI DevRel | EMEA

August 2016

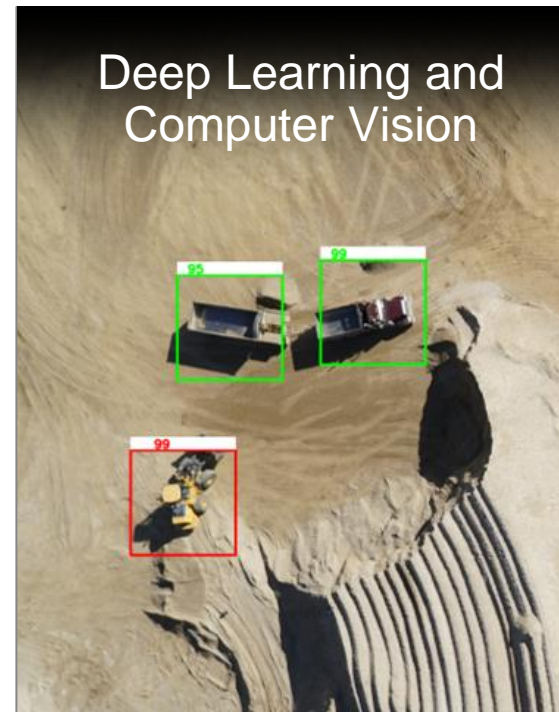




NVIDIA

Pioneered GPU Computing | Founded 1993 | \$7B | 9,500 Employees



NVIDIA GPU: MORE THAN GRAPHICS




NVIDIA


Computing for the Most Demanding Users



GPU Computing



Computing Human Imagination



Computing Human Intelligence

ALL OVER

GTC CHINA

BEIJING

Sept. 13

GTC TAIWAN

TAIPEI

Sept. 21

GTC EUROPE

AMSTERDAM

Sept. 28 & 29

GTCx Australia

MELBOURNE

Oct. 4

GTC JAPAN

TOKYO

Oct. 5

GTCx KOREA

SEOUL

Oct. 7

GTC D.C.

WASHINGTON

Oct. 26 & 27

GTCx INDIA

MUMBAI

Dec. 6

THE EXPANDING UNIVERSE OF MODERN AI

"THE BIG BANG"

Big Data
GPU
Algorithms

RESEARCH



CORE TECHNOLOGY / FRAMEWORKS



AI-as-a-PLATFORM



START-UPS




1,000+ AI START-UPS
\$5B IN FUNDING

Source: Venture Scanner

INDUSTRY LEADERS



Exponential Growth



SOURCE: RAY KURZWEIL, "THE SINGULARITY IS NEAR: WHEN HUMANS TRANSCEND BIOLOGY", P.67, THE VIKING PRESS, 2006. DATAPOINTS BETWEEN 2000 AND 2012 REPRESENT BCA ESTIMATES.

HOW

GPU Computing



CUDA

Framework to Program NVIDIA GPUs

A simple sum of two vectors (arrays) in C

```
void vector_add(int n, const float *a, const float *b, float *c)
{
    for( int idx = 0 ; idx < n ; ++idx )
        c[idx] = a[idx] + b[idx];
}
```

GPU friendly version in CUDA

```
__global__ void vector_add(int n, const float *a, const float *b, float *c)
{
    int idx = blockIdx.x*blockDim.x + threadIdx.x;
    if( idx < n )
        c[idx] = a[idx] + b[idx];
}
```


THE ENGINE OF MODERN AI

EDUCATION

TORCH



CAFFE



THEANO



MATCONVNET



MOCHA.JL



PURINE



MINERVA



MXNET*



BIG SUR



TENSORFLOW



WATSON



CNTK



START-UPS

CHAINER



DL4J



KERAS





OPENDEEP




NVIDIA DEEP LEARNING PLATFORM

*U. Washington, CMU, Stanford, TuSimple, NYU, Microsoft, U. Alberta, MIT, NYU Shanghai

CONVOLUTIONAL NEURAL NETWORKS



Biological vs artificial




Deep neural network (DNN)

Raw data



Typically millions of images


Low-level features



Mid-level features



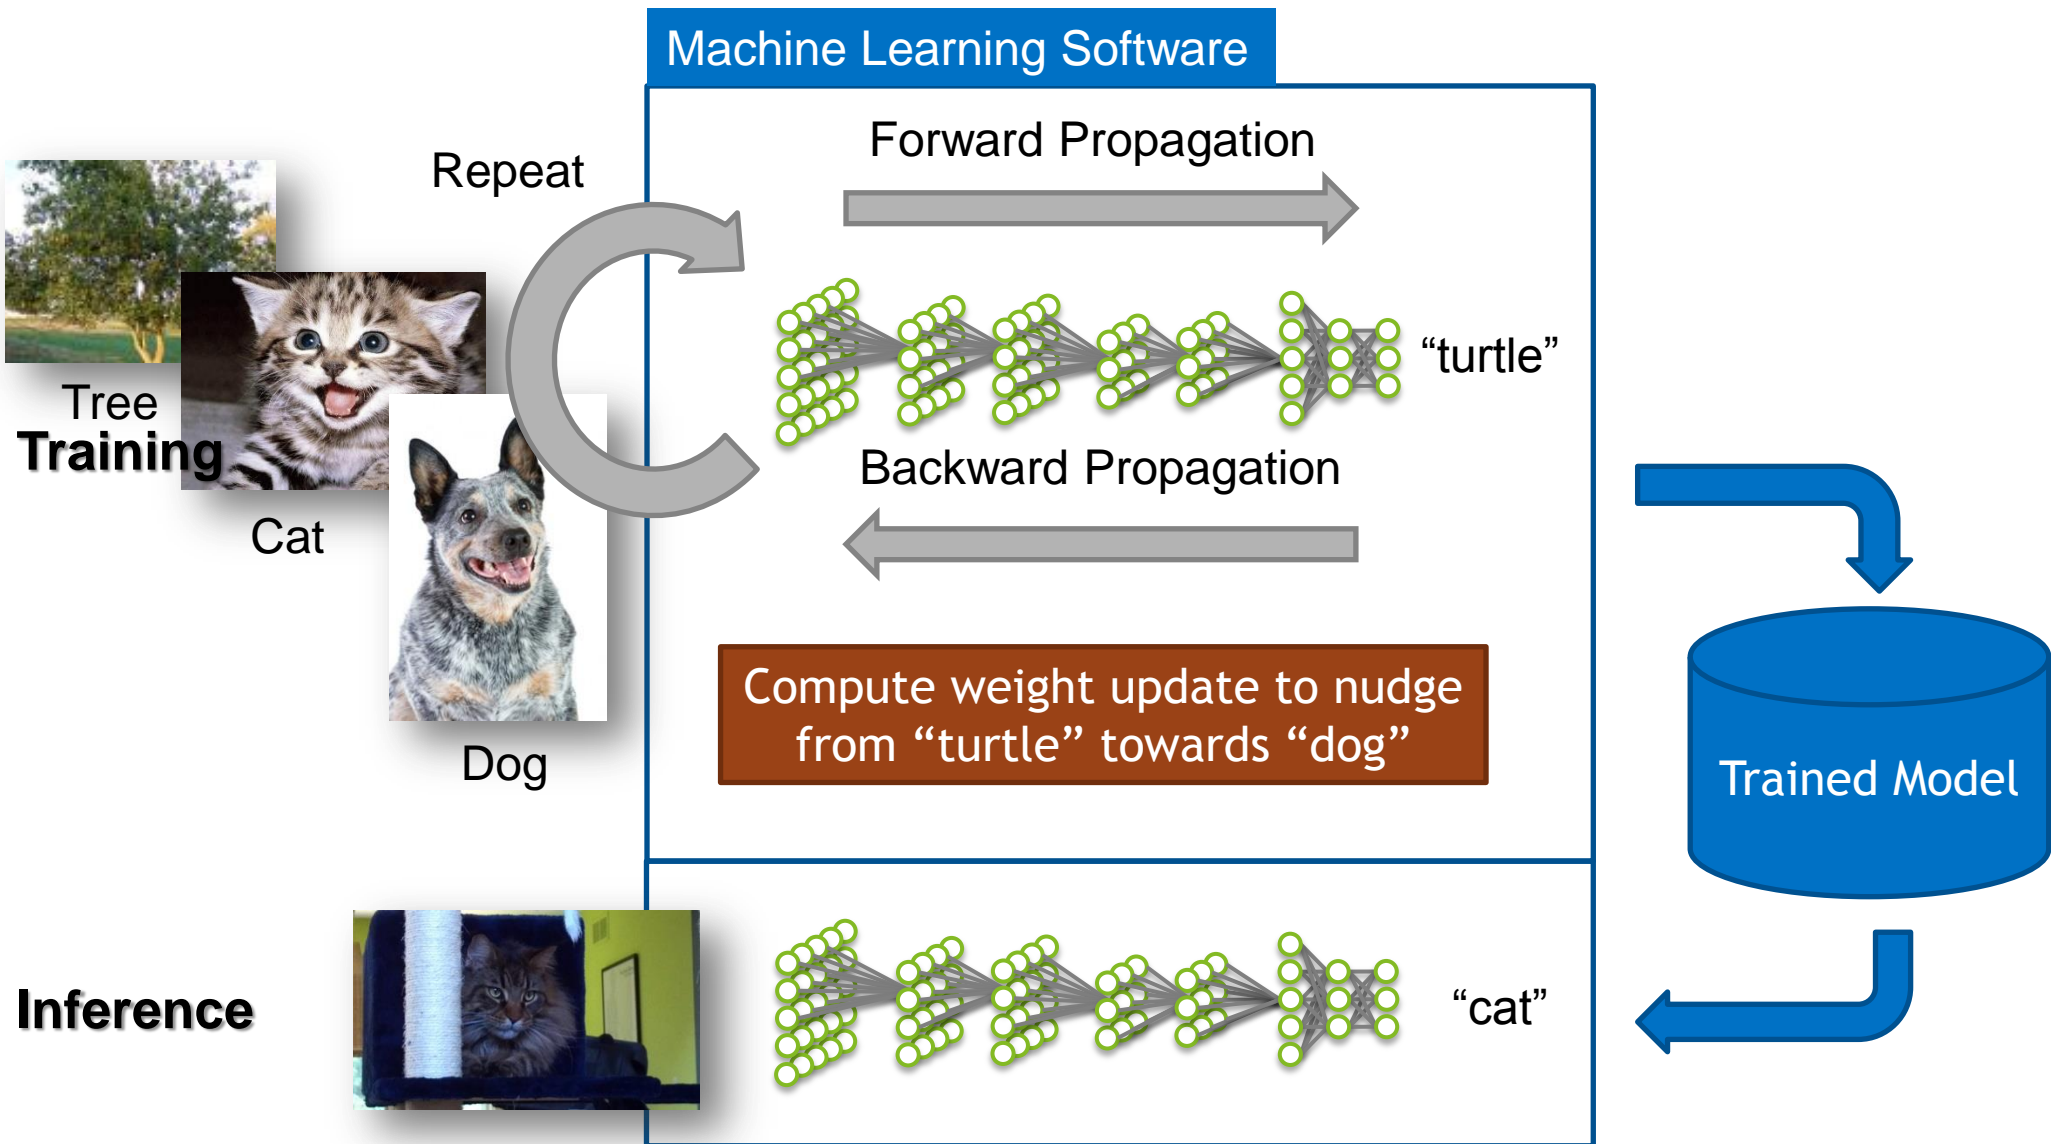
High-level features




Robust

Generalizable

Scalable



The Theory of Kevin (and Jamie aged 18 mths)

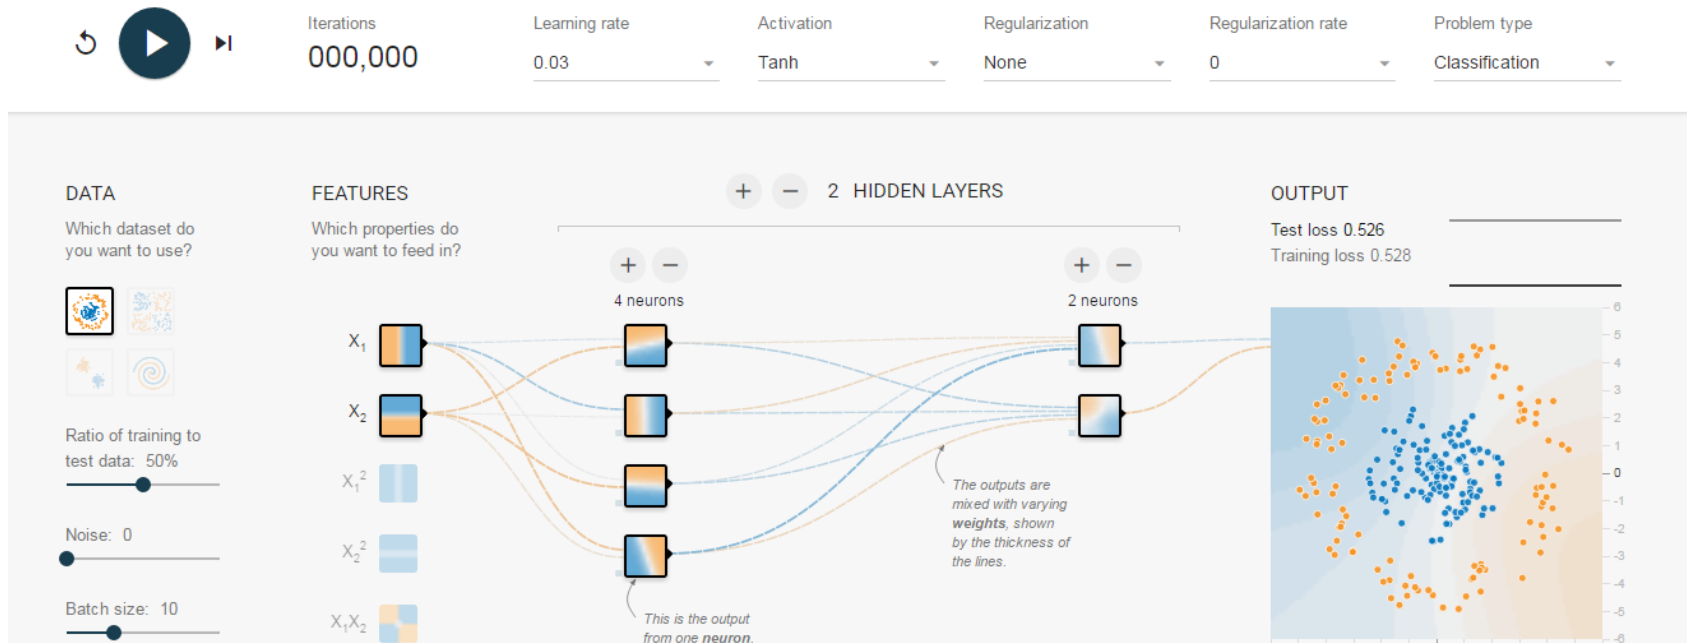


NN playground

Visualising a neural network

<http://playground.tensorflow.org>

Tinker With a **Neural Network** Right Here in Your Browser.
Don't Worry, You Can't Break It. We Promise.

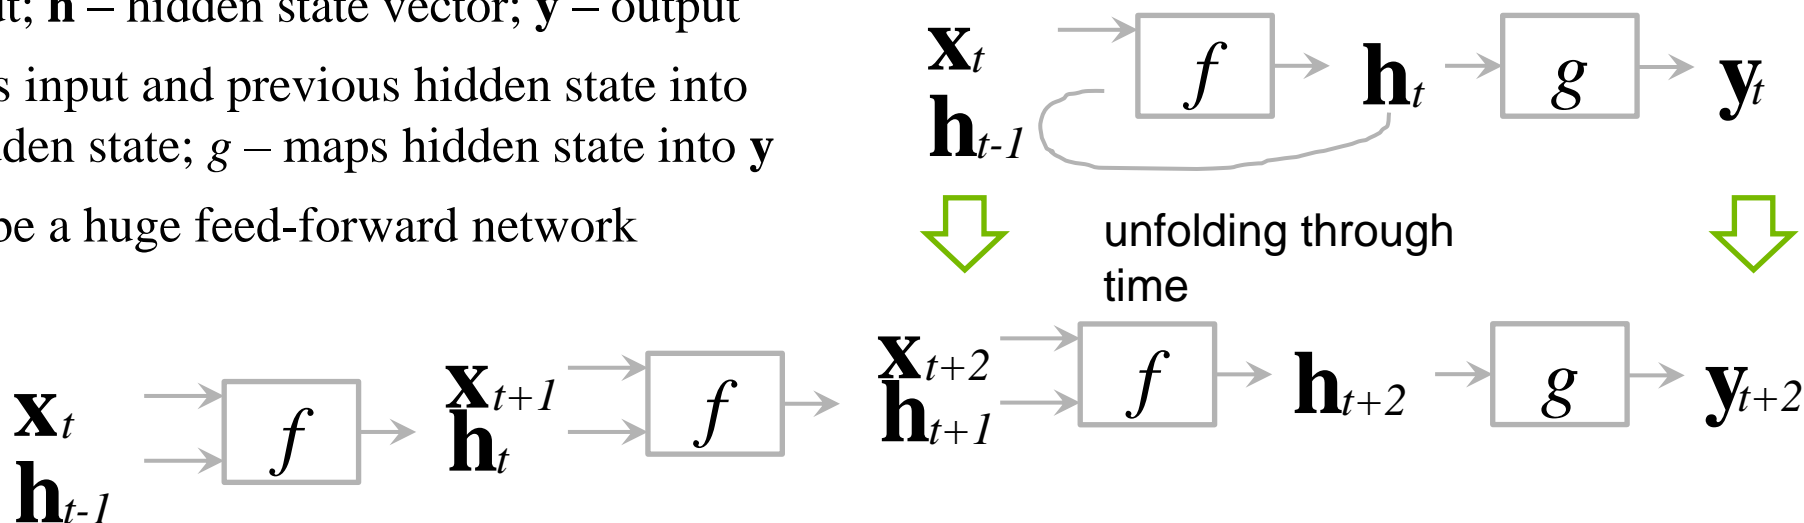


Recurrent neural networks

\mathbf{x} – input; \mathbf{h} – hidden state vector; \mathbf{y} – output

f – maps input and previous hidden state into new hidden state; g – maps hidden state into \mathbf{y}

f – can be a huge feed-forward network




“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

Long short-term memory (LSTM)

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



**Training
via
backprop
unfolded
in time**



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

**Gates control importance of
the corresponding
activations**

Fig from Vinyals et al, Google April 2015 NIC Generator

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

THE NEXT STEP – NATURAL LANGUAGE PROCESSING

CNN + RNN




Figure 4. Diagram of our multimodal Recurrent Neural Network generative model. The RNN takes a word, the context from previ-

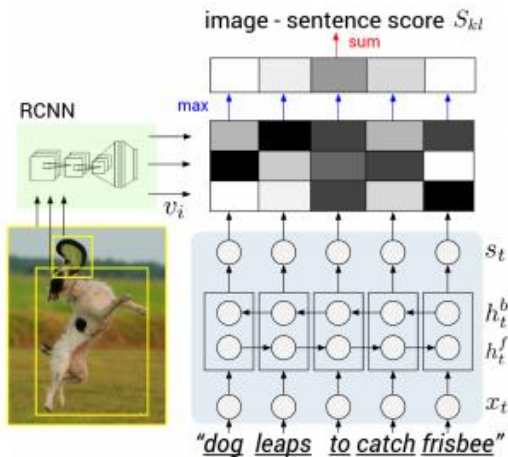


Figure 3. Diagram for evaluating the image-sentence score S_{kl} . Object regions are embedded with a CNN (left). Words (enriched by their context) are embedded in the same multimodal space with a BRNN (right). Pairwise similarities are computed with inner products (magnitudes shown in grayscale) and finally reduced to image-sentence score with Equation 8.



"girl in pink dress is jumping in air."



"black and white dog jumps over bar."



"young girl in pink shirt is swinging on swing."



"man in blue wetsuit is surfing on wave."



"little girl is eating piece of cake."



"baseball player is throwing ball in game."



"woman is holding bunch of bananas."



"black cat is sitting on top of suitcase."



"a young boy is holding a baseball bat."



"a cat is sitting on a couch with a remote control."



"a woman holding a teddy bear in front of a mirror."



"a horse is standing in the middle of a road."

Dynamic Memory Networks

MetaMind now SalesforceIQ <https://arxiv.org/pdf/1603.01417v1.pdf>



What is this sculpture made out of ?



Answer: **metal**



What color are the bananas ?



Answer: **green**



Who is on both photos ?



Answer: **girl**



What is the pattern on the cat 's fur on its tail ?



Answer: **stripes**



Did the player hit the ball ?



Answer: **yes**



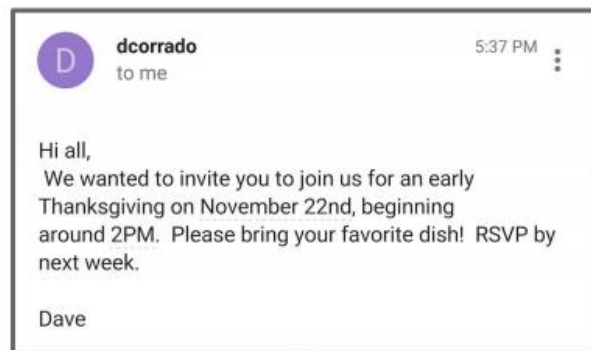
What is the boy holding ?



Answer: **surfboard**

Smart Reply

Incoming Email



Small
Feed-Forward
Neural Network

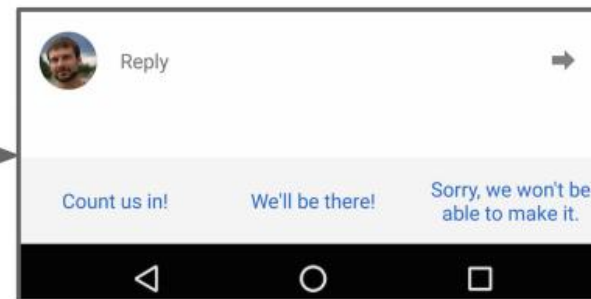
Activate
Smart Reply?

yes/no



Deep
Recurrent
Neural Network

Generated Replies



Example Data:



CHEST 2V FRONTAL/LATERAL XXXX, XXXX XXXX PM

Comparison: None

Indication: Burmese male
has complete TB treatmentFindings: Both lungs are
clear and expanded with
no infiltrates. Basilar focal
atelectasis is present in the
lingula. Heart size normal.
Calcified right hilar XXXX
are present

Impression: No active disease.

MeSH

Major


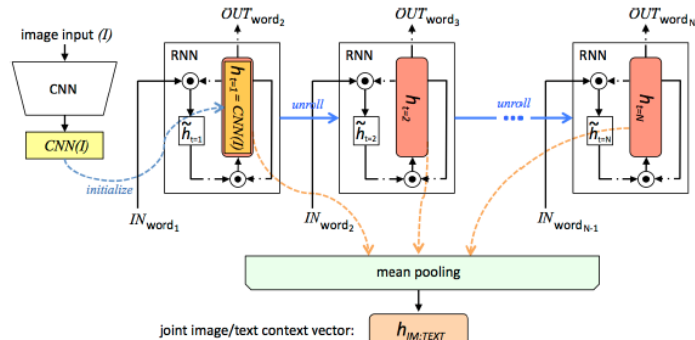
Pulmonary Atelectasis
/ lingula / focal
Calcinosis
/ lung / hilum / right

Generating image annotation:


Trained model

Joint image/text context vector:

- Averaging CNN/RNN encoding vectors for composite image labeling, considering both image and text contexts



Diagnosing with LSTM



Left: Diabetic Ketoacidosis, Right: Brain Neoplasm

<http://arxiv.org/pdf/1511.03677v6.pdf>

Bellman Principle and Q-function

$Q(\mathbf{s}; \mathbf{a})$: the maximum expected return achievable by following any strategy after seeing sequence \mathbf{s} and taking action \mathbf{a}

Bellman principle (dynamic programming)


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

$$r(\mathbf{t}+1) + Q(\mathbf{s}(\mathbf{t}+1); \mathbf{a}(\mathbf{t}+1)) \Rightarrow \max$$

DQN: deep Q-learning network

REINFORCEMENT LEARNING

Mastering Breakout



WHY

Deep learning teaches robots

engadget

Amazon robot challenge winner counts on deep learning AI



MIT
Technology
Review


China Is Building a Robot Army of Model Workers

Forbes

Japan Must Refocus From US-dominated AI to Integrating Deep Learning into Manufacturing

Pieter Abbeel

gym.openai.com



DEEP-LEARNING RESEARCH ROVER

TURBO 2.0




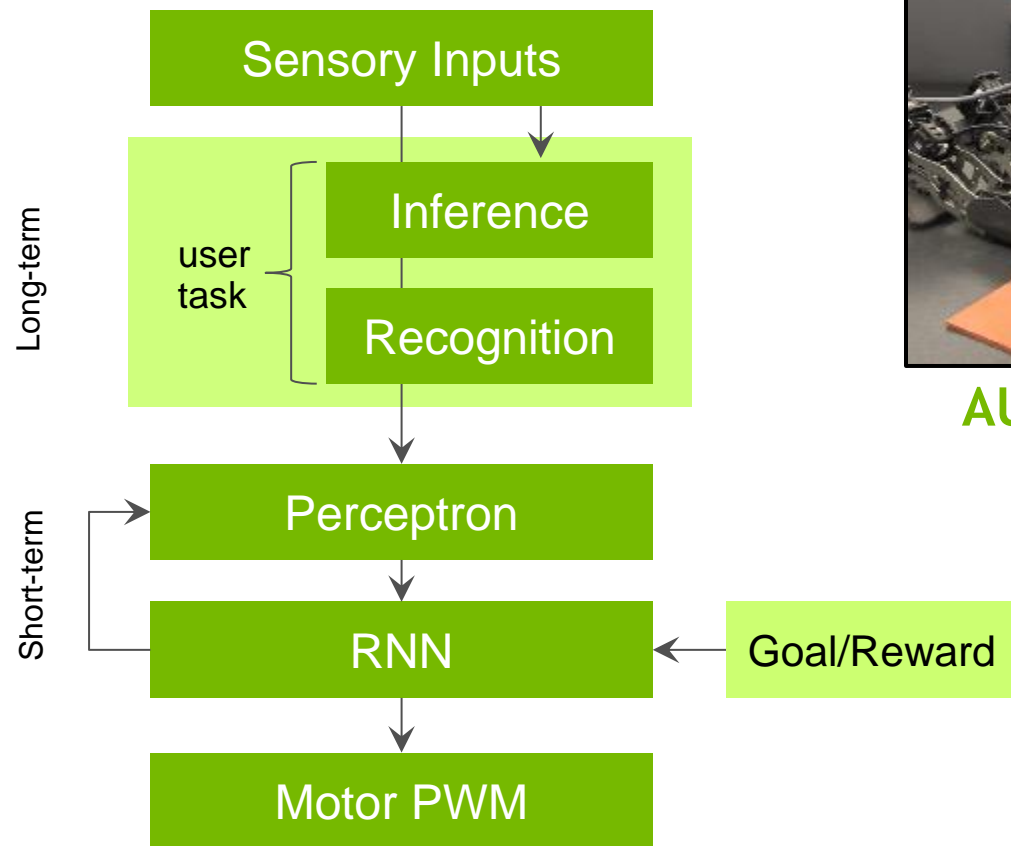
github.org/dusty-nv




GitHub

DEEP REINFORCEMENT LEARNING

MOTION CONTROL



AUTONOMOUS NAVIGATION




Da Vinci medical robotics




Automating Cardiac MRI analysis


DL performance matches expert cardiologist at computing ejection fraction - a key indicator of heart disease



imaging






Manual
annotation



C.M.S. Nambakhsh et al./Medical Image Analysis 17 (2013) 1010–1024

1019



Software
volume
estimate

GOOGLE DEEPMIND ALPHAGO CHALLENGE



DEEP LEARNING REINVENTS CARS

THE VERGE

Meet Uber's Self-Driving Car



TECH
INSIDER

Google, Apple, Tesla race to develop self-driving cars by 2020

TRUCKS.COM
FUELING YOUR MIND FOR THE ROAD AHEAD

Daimler test self-driving truck platoon in live traffic

theguardian

Volvo to test self-driving cars on London's roads next year

FORTUNE

GM buying self-driving tech startup for more than \$1 billion

FORTUNE

Baidu wants to bring self-driving cars to the masses by 2021

 **REUTERS**

Daimler aims to launch self-driving car by 2020

TE **TechCrunch**

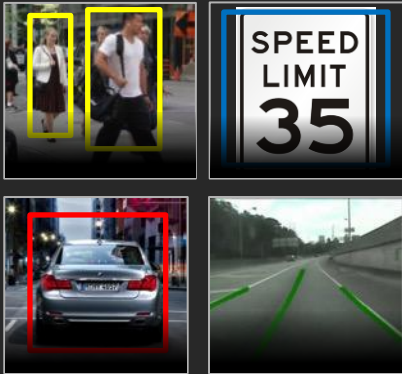
Otto, founded by ex-Googlers, is bringing self-driving technology to trucks

CHALLENGE for self-driving cars






REVOLUTIONIZED by DEEP LEARNING

CONVENTIONAL




DEEP NEURAL NETWORK






POWERED BY TEGRA X1




Visual computing and AI will make future cars safer and delightful to drive.


NVIDIA DRIVE PX™ 2 can perform 24 trillion deep learning operations per second. The size of a lunchbox, it can fit neatly into a trunk.

The DRIVE PX platform includes DriveWorks, for developing applications across the entire self-driving pipeline; DIGITS, for training and visualizing deep neural networks; and DriveNet, our reference deep neural network.



NVIDIA DRIVE™ PX 2
Selected by Volvo on
Journey Towards a
Crash-Free Future






NVIDIA DRIVENET VIDEO

Paper on
<http://arxiv.org/abs/1604.07316>

MIT & DARPA put LIDAR on a chip using PHOTONICS



WORLD'S FIRST AUTONOMOUS CAR RACE

10 teams, 20 identical cars | DRIVE PX 2: The “brain” of every car | 2016/17 Formula E season



NVIDIA Deep Learning Platform

NVIDIA DEEP LEARNING PLATFORM


DL FRAMEWORK (CAFFE, CNTK, TENSORFLOW, THEANO, TORCH...)

DEEP LEARNING SDK

DEVELOPERS



DEPLOYMENT



AUTOMOTIVE - DRIVEPX



EMBEDDED - JETSON




CUDA FOR DEEP LEARNING DEVELOPMENT


DEEP LEARNING SDK



DIGITS




cuDNN



cuSPARSE



cuBLAS



NCCL

TITAN X



DEVBOX




GPU CLOUD



NVIDIA cuDNN

Building blocks for accelerating deep neural networks on GPUs

- ▶ High performance deep neural network training
- ▶ Accelerates Deep Learning: Caffe, CNTK, Tensorflow, Theano, Torch
- ▶ Performance continues to improve over time



*AlexNet training throughput based on 20 iterations,
CPU: 1x E5-2680v3 12 Core 2.5GHz.*




“NVIDIA has improved the speed of cuDNN with each release while extending the interface to more operations and devices at the same time.”

— Evan Shelhamer, Lead Caffe Developer, UC Berkeley

developer.nvidia.com




NVIDIA SDK



cuDNN

Deep Learning Primitives


IGNITING ARTIFICIAL
INTELLIGENCE




developer.nvidia.com/cudnn

- GPU-accelerated Deep Learning subroutines
- High performance neural network training
- Accelerates Major Deep Learning frameworks: Caffe, Theano, Torch
- Up to 3.5x faster AlexNet training in Caffe than baseline GPU

Millions of Images Trained Per Day



Tiled FFT up to 2x faster than FFT



WHAT'S NEW IN CUDNN 5?

Pascal GPU, RNNs, Improved Performance

LSTM recurrent neural networks deliver up to 6x speedup in Torch

Improved performance:

- Deep Neural Networks with 3x3 convolutions, like VGG, GoogleNet and ResNets
- 3D Convolutions
- FP16 routines on Pascal GPUs

5.9x

Speedup for char-rnn
RNN Layers

2.8x

Speedup for DeepSpeech 2
RNN Layers

Performance relative to torch-rnn
(<https://github.com/jcjohnson/torch-rnn>)
DeepSpeech2: <http://arxiv.org/abs/1512.02595>
Char-rnn: <https://github.com/karpathy/char-rnn>

DIGITS™

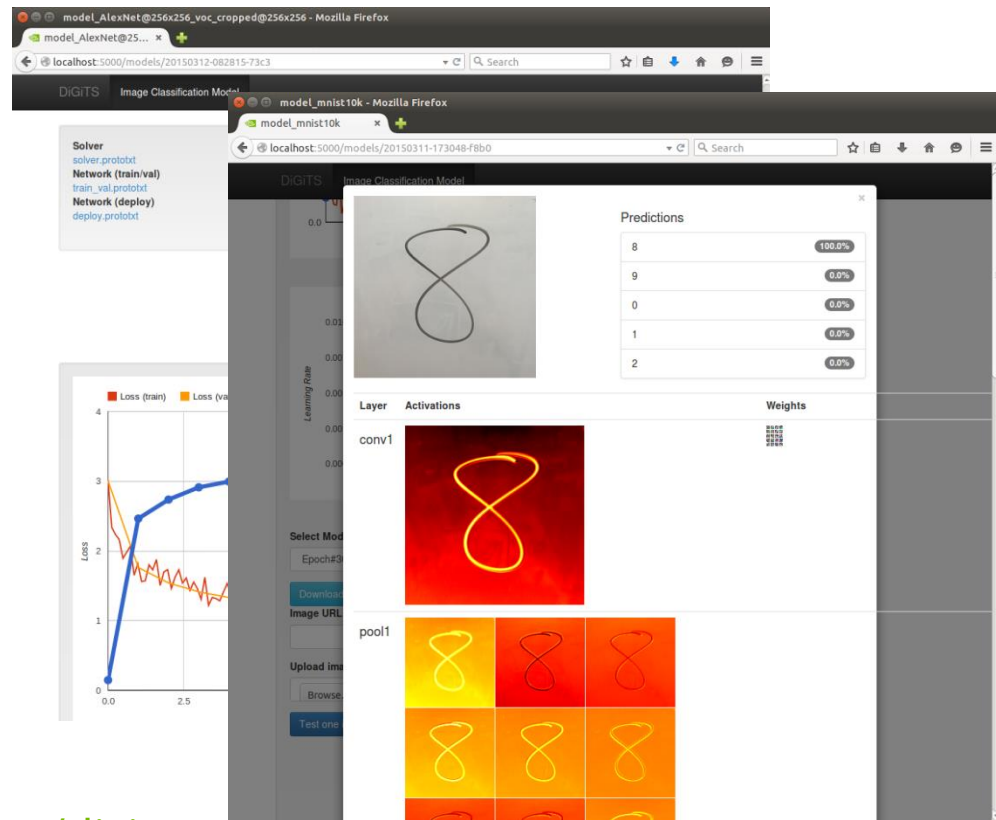
Interactive Deep Learning GPU Training System

Quickly design the best deep neural network (DNN) for your data

Train on multi-GPU (automatic)

Visually monitor DNN training quality in real-time


Manage training of many DNNs in parallel on multi-GPU systems



developer.nvidia.com/digits

DEEP VISUALIZATION TOOLBOX

IMAGE RECOGNITION




DIGITS 4

Object Detection Workflow

- Object Detection Workflows for Automotive and Defense
- Targeted at Autonomous Vehicles, Remote Sensing

<https://devblogs.nvidia.com/parallelforall/>



DetectNet Generic Image Model

Job Status Done

- Initialized at 10:26:49 PM (1 second)
- Running at 10:26:51 PM (16 seconds)
- Done at 10:27:08 PM (Total - 18 seconds)

Object Detection Task Done ▾

Infer Model Done ▾

Notes

None

Detections are shown as red boxes. If no boxes are shown, then there were no detections for this image.

Description	Statistics	Visualization
"data" Activation	Data shape: [3 384 1248] Mean: 92.5834 Std deviation: 72.3209 	

OBJECT DETECTION


REMOTE SENSING




PICKING AND GRASPING



ADVANCED DRIVER ASSISTANCE SYSTEMS (ADAS)



MEDICAL DIAGNOSTICS




NCCL ‘nickel’

Accelerating Multi-GPU Communications

A topology-aware library of accelerated collectives to improve the scalability of multi-GPU applications

- Patterned after MPI’s collectives: includes all-reduce, all-gather, reduce-scatter, reduce, broadcast
- Optimized intra-node communication
- Supports multi-threaded and multi-process applications



github.com/NVIDIA/nccl

GRAPH ANALYTICS with NVGRAPH

developer.nvidia.com/nvgraph



GPU Optimized Algorithms






Reduced cost & Increased performance



Standard formats and primitives
Semi-rings, load-balancing



Performance Constantly Improving





GPU INFERENCE ENGINE (GIE)

High-performance deep learning inference for production deployment

DATA CENTER




AUTOMOTIVE



EMBEDDED



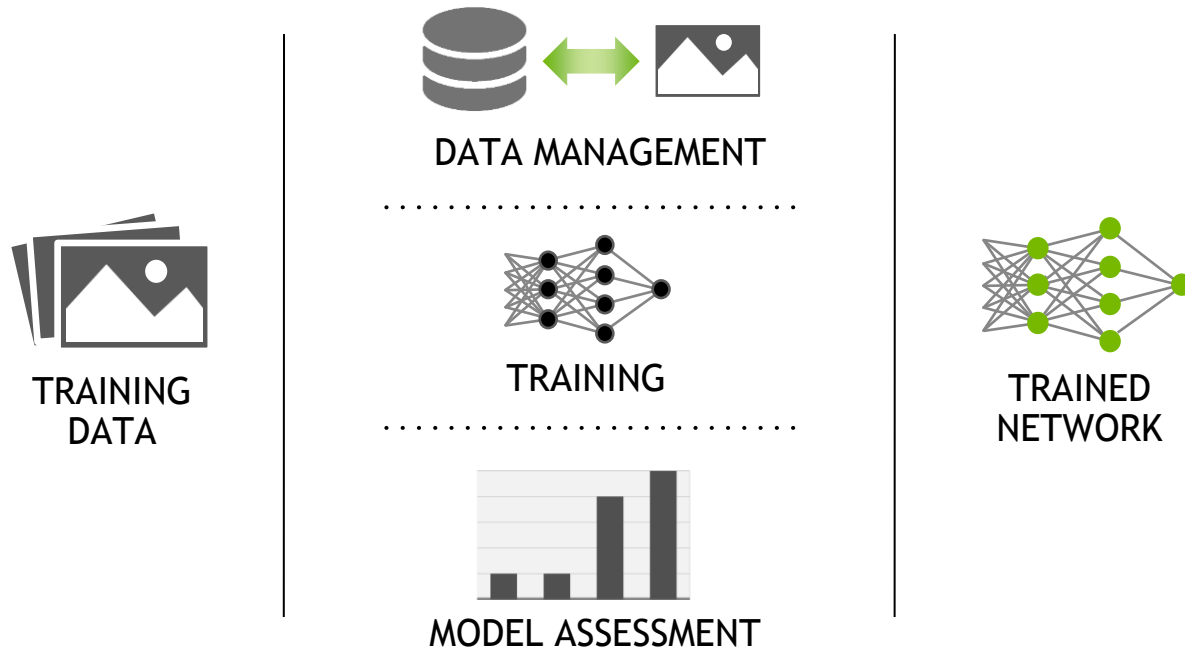
developer.nvidia.com/gie




GoogLeNet, CPU-only vs Tesla M4 + GIE on
Single-socket Haswell E5-2698 v3@2.3GHz with HT

NVIDIA DEEP LEARNING SOFTWARE PLATFORM

DEVELOP WITH DIGITS



DEPLOY WITH GIE



EMBEDDED



DATA CENTER



AUTOMOTIVE

NVIDIA DEEP LEARNING SDK


WHAT'S NEW IN DEEP LEARNING SOFTWARE

DIGITS 4



Object Detection

GIE



High performance deep
learning inference

cuDNN 5.1



Improved performance for
VGG, ResNet style networks

Deep Learning Hardware

TITAN X

The Ultimate


Pascal Architecture | 11TFLOPS | 12GB G5X



TITAN X PERFORMANCE

Slash DL Training Time by 40%


Caffe Googlenet Training Time




	TITAN X PASCAL	TITAN X MAXWELL
CUDA cores	3584	3072
Boost Clock	1.53 GHZ	1.08GHZ
Memory	12GB G5X	12GB G5
Memory Bandwidth (GB/s)	480	336
GFLOPS (INT8)	44	-
GFLOPS (FP32)	11	7

INTRODUCING TESLA P100


Five Technology Breakthroughs Made it Possible




Pascal Architecture




16nm
FinFET



COWOS with
HBM2 Stacked Memory




NVLink



New AI
Algorithms

FP16 vs FP32


	FP16	FP32
Max. Value	~61,000	~1e38
Grid points between each power of 2	2048	~16,700,000
Smallest number you can add to one and get a different number (ULP relative to 1)	~.000489	~.00000006



LSDalton

Quantum
Chemistry


12X speedup
in 1 week



Numeca

CFD


10X faster kernels
2X faster app



PowerGrid

Medical
Imaging


40 days to
2 hours



INCOMP3D

CFD


3X speedup



NekCEM

Computational
Electromagnetics


2.5X speedup
60% less energy



COSMO

Climate
Weather


40X speedup
3X energy efficiency



CloverLeaf

CFD


4X speedup
Single CPU/GPU code



**MAESTRO
CASTRO**

Astrophysics


4.4X speedup
4 weeks effort



Today, Tesla GPUs power the fastest supercomputers in the US, Europe and Japan, and 20 of the 25 most energy-efficient in the world.

NVIDIA and IBM have been selected to power two next-gen USA supercomputers, which are expected to be the worlds fastest when they come online.

Powered by Tesla GPUs and our breakthrough NVLink interconnect technology.





NVIDIA DGX-1


WORLD'S FIRST DEEP LEARNING SUPERCOMPUTER

Engineered for deep learning | 170TF FP16 | 8x Tesla P100
NVLink hybrid cube mesh | Accelerates major AI frameworks

CUDA 8 - WHAT'S NEW


P100 Support

Stacked Memory
NVLINK
FP16 math




Unified Memory

Larger Datasets
Demand Paging
New Tuning APIs
Standard C/C++ Allocators
CPU/GPU Data Coherence & Atomics




Libraries

New nvGRAPH library
cuBLAS improvements for Deep Learning



Developer Tools

Critical Path Analysis
2x Faster Compile Time
OpenACC Profiling
Debug CUDA Apps on Display GPU




NVIDIA® DGX-1™



DGX STACK

Fully integrated Deep Learning platform



Instant productivity — plug-and-play, supports every AI framework

Performance optimized across the entire stack

Always up-to-date via the cloud



Mixed framework environments — virtualized and containerized

Direct access to NVIDIA experts


NVIDIA DGX-1 SOFTWARE STACK

Optimized for Deep Learning Performance


Accelerated Deep Learning



Container Based Applications



NVIDIA Cloud Management





A SUPERCOMPUTER FOR AUTONOMOUS MACHINES


Bringing AI and machine learning to
a world of robots and drones

Jetson TX1 is the first embedded
computer designed to process deep
neural networks

1 TeraFLOPS in a credit-card sized
module



System on Module




	JETSON TX1
GPU	1 TFLOP/s 256-core Maxwell
CPU	4x 64-bit ARM A57 CPUs 1.6 GHz
Memory	4 GB LPDDR4 25.6 GB/s
Video decode	4K 60Hz H.264
Video encode	4K 30Hz H.264
CSI	Up to 6 cameras 1400 Mpix/s
Display	2x DSI, 1x eDP 1.4, 1x DP 1.2/HDMI
Wi-Fi	802.11 2x2 ac
Networking	1 Gigabit Ethernet
PCI-E	Gen 2 1x1 + 1x4
Storage	16 GB eMMC, SDIO, SATA
Other	3x UART, 3x SPI, 4x I2C, 4x I2S, GPIOs
Power	10-15W, 6.6V-19.5VDC
Size	50mm x 87mm

NVIDIA JETPACK




Deep Learning SDK



DIGITS Workflow




VisionWorks



Jetson Media SDK

and other technologies:
CUDA, Linux4Tegra, NSIGHT EE, OpenCV4Tegra, OpenGL, Vulkan, System Trace, Visual Profiler



Jetson TX1 Developer Kit

Jetson TX1
Developer Board
5MP Camera



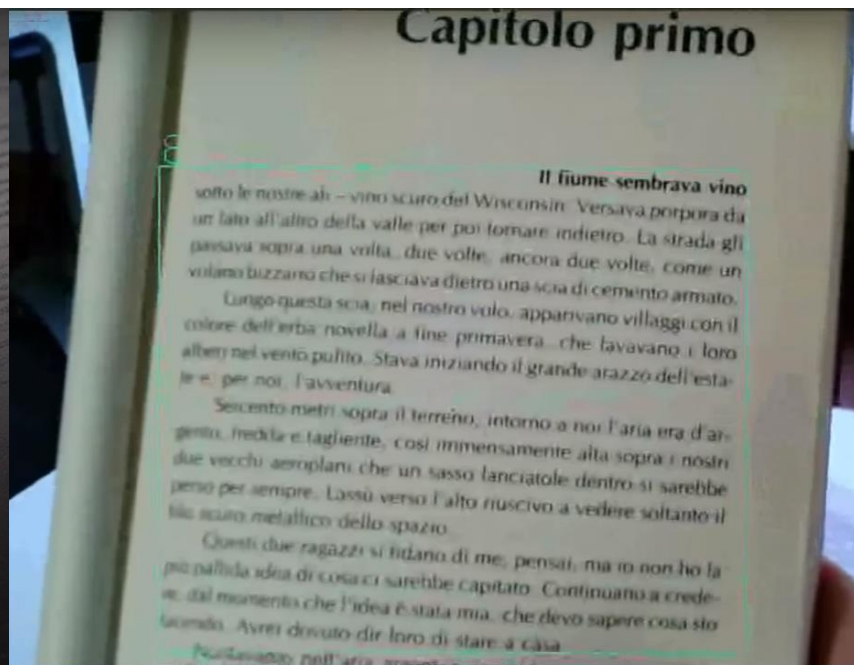
Horus
Technology

HORUS

BLIND ASSIST DEVICE

EMERGING COMPANIES
SUMMIT

SOCIAL
INNOVATION
AWARD WINNER



Il fiume sembrava vino sotto le nostre ali — vino scuro del Wisconsin. Versava porpora da un lato all'altro della valle per poi tornare indietro. La strada gli passava sopra una volta, due volte, ancora due volte, come un volano bizzarro che si lasciava dietro una scia di cemento armato.


Lungo questa scia, nel nostro volo, apparivano villaggi con il colore dell'erba novella a fine primavera, che lavavano i loro alberi nel vento pulito. Stava iniziando il grande arazzo dell'estate e, per noi, l'avventura.

Seicento metri sopra il terreno, intorno a noi l'aria era d'argento, fredda e tagliente, così immensamente alta sopra i nostri due vecchi aeroplani che un sasso lanciato dentro si sarebbe perso per sempre. Lassù verso l'alto riuscivo a vedere soltanto il blu scuro metallico dello spazio.

Questi due ragazzi si fidano di me, pensai, ma io non ho la più pallida idea di cosa ci sarebbe capitato. Continuano a credere, dal momento che l'idea è stata mia, che devo sapere cosa sto facendo. Avrei dovuto dir loro di stare a casa.

Horus.tech

```
cv::Mat rect = rectFromRect(cv::Rect(0,0,1000,1000));
auto c = std::string(0x readBlock(rect));
const cv::Mat img = cv::imread(c);
wavPlayer->playWAV(c);
ptSrc.stop();
rectToTrack = nullptr;
imshow("Rectified", rect);
```





**Potentially Hazardous Asteroid
detection/classification**

Wrap Up

NVIDIA DEEP LEARNING EVERYWHERE, EVERY PLATFORM



TITAN X

Available via retail in
200+ countries



DGX-1

The HPC appliance for
instant productivity



TESLA

Servers in every shape and size



CLOUD

Everywhere



Getting started with deep learning

developer.nvidia.com/deep-learning

 **NVIDIA** ACCELERATED COMPUTING

Downloads Training Ecosystem Forums

Q [Register Now](#) [Log in](#)

DEEP LEARNING

Deep learning is the fastest-growing field in machine learning. It uses many-layered Deep Neural Networks (DNNs) to learn levels of representation and abstraction that make sense of data such as images, sound, and text.



[Home](#) > [ComputeWorks](#) > [Deep Learning](#)



Get Started With
Deep Learning



Download Deep
Learning Software




Deep Learning
Institute

NVIDIA GPUs - The Engine of Deep Learning

Traditional machine learning uses handwritten feature extraction and modality-specific machine learning algorithms to label images or recognize voices. However, this method has several drawbacks in both time-to-solution and accuracy.

Today's advanced deep neural networks use algorithms, big data, and the computational power of the GPU to change this dynamic. Machines are now able to learn at a speed, accuracy, and scale that are driving true **artificial intelligence**.



NVIDIA Deep Learning Institute

Hands-on Training for Data Scientists and Software Engineers



Training organizations and individuals to solve challenging problems using Deep Learning

On-site workshops and online courses presented by certified experts

Covering complete workflows for proven application use cases

Image classification, object detection, natural language processing, recommendation systems, and more


www.nvidia.com/dli

GPU TECHNOLOGY CONFERENCE

Sep 28-29, 2016 | Amsterdam
www.gputechconf.eu #GTC16EU

EUROPE'S BRIGHTEST MINDS & BEST IDEAS

GET A 20% DISCOUNT WITH CODE *ALLOGTCEU2016*




DEEP LEARNING &
ARTIFICIAL INTELLIGENCE



AUTONOMOUS VEHICLES



VIRTUAL REALITY &
AUGMENTED REALITY



SUPERCOMPUTING & HPC

GTC Europe is a two-day conference designed to expose the innovative ways developers, businesses and academics are using parallel computing to transform our world.

2 Days | 1,000 Attendees | 50+ Exhibitors | 50+ Speakers | 10+ Tracks | 15+ Hands-on Labs | 1-to-1 Meetings

COME DO YOUR LIFE'S WORK

JOIN NVIDIA

We are looking for great people at all levels to help us accelerate the next wave of AI-driven computing in Research, Engineering, and Sales and Marketing.

Our work opens up new universes to explore, enables amazing creativity and discovery, and powers what were once science fiction inventions like artificial intelligence and autonomous cars.

Check out our career opportunities:

- www.nvidia.com/careers
- Reach out to your NVIDIA social network or NVIDIA recruiter at DeepLearningRecruiting@nvidia.com



NVIDIA®

Questions?
alowndes@nvidia.com

