

5th HVM New Materials Conference Summit
Cambridge, UK 6-7 November 2019

www.cir-strategy.com/events



WHY EVERY INDUSTRY SHOULD BE THINKING ABOUT AI TODAY

Jonny Hancox
jhancox@nvidia.com



Introduction to NVIDIA



Why the current AI hype?

Industry examples

Getting started with AI

NVIDIA

“THE AI COMPUTING COMPANY”


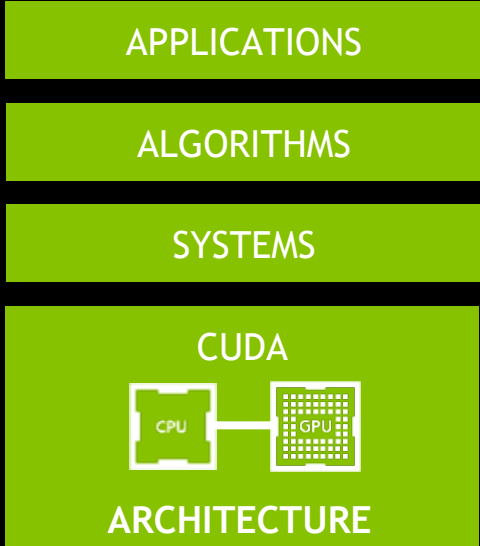


GPU Computing


Computer Graphics

Artificial Intelligence

THE FORCES SHAPING COMPUTING



40 YEARS OF CPU TREND DATA




ALEXNET: THE SPARK OF THE MODERN AI ERA

The background features a complex network of thin, glowing green lines connecting various nodes. Some nodes are bright green, while others are a soft blue. The overall effect is a sense of interconnectedness and digital activity against a dark, almost black background.

WHY THE CURRENT AI HYPE?

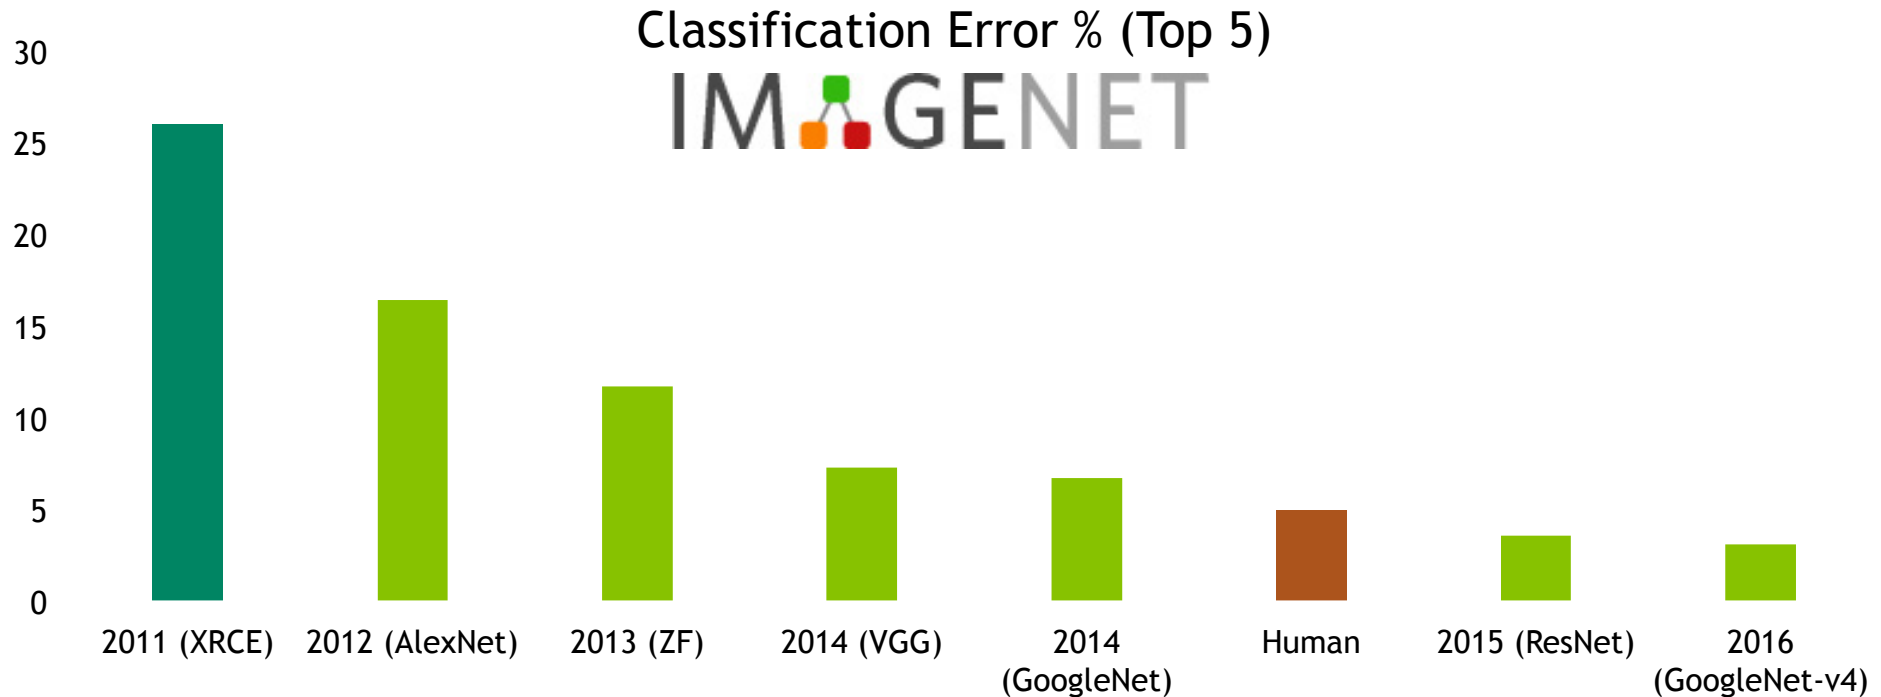
THE PERCEPTRON

The fundamental unit of artificial neural networks

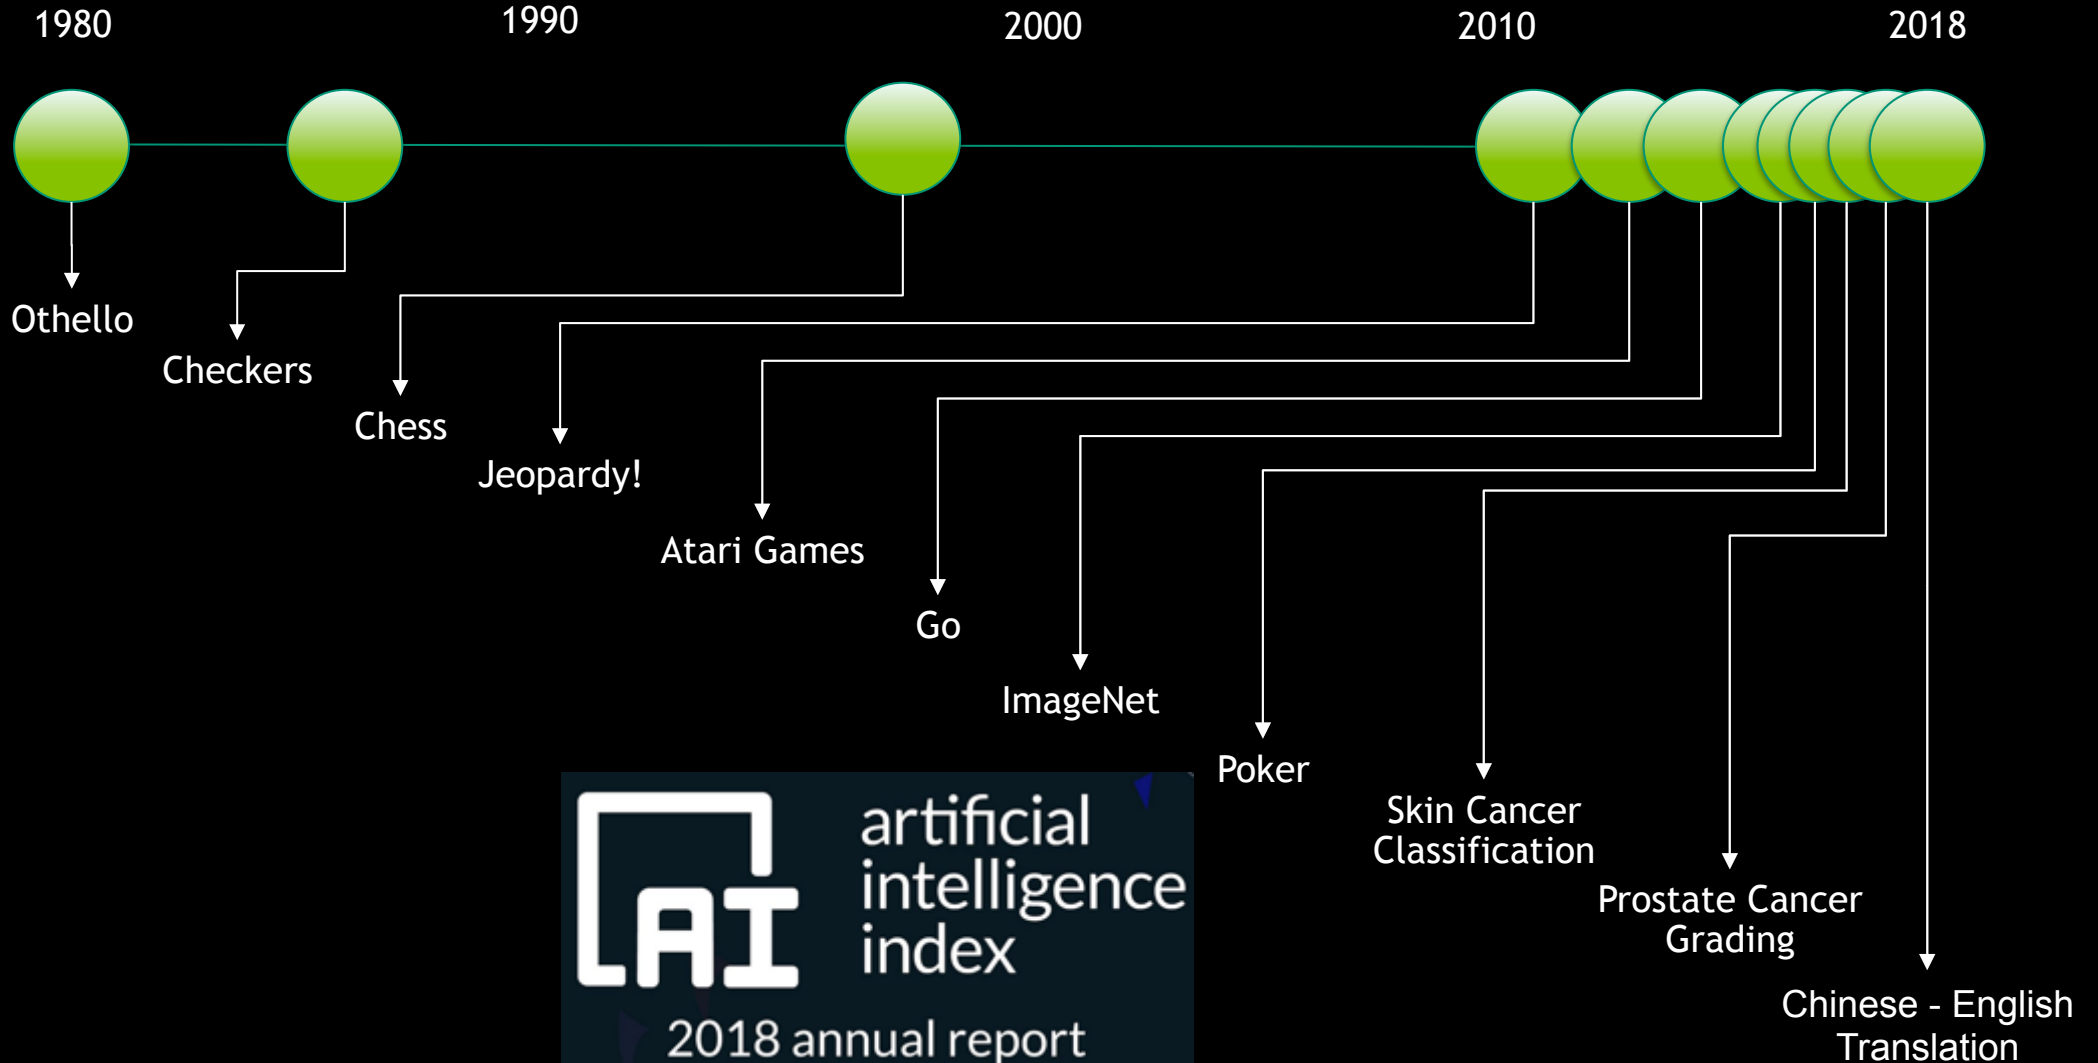


CLASSIFYING INTERNET IMAGES

Alex Krizhevsky's 'AlexNet' spawned a new generation of DNNs



SUPERHUMAN PERFORMANCE




DEEP LEARNING IS SWEEPING ACROSS INDUSTRIES

INTERNET SERVICES



Image/Video classification
Speech recognition
Natural language processing

MEDICINE



Cancer cell detection
Diabetic grading
Drug discovery

MEDIA & ENTERTAINMENT



Video captioning
Content based search
Real time translation

SECURITY & DEFENSE



Face recognition
Video surveillance
Cyber security

AUTONOMOUS MACHINES



Pedestrian detection
Lane tracking
Recognize traffic signs

*“You need a chief AI officer...
If you have a lot of data and you want to
create value from that data, one of the
things you might consider is building up an
AI team.”*

Andrew Ng, Baidu chief scientist

EVOLUTION OF COMPUTING



PC Internet
WinTel, Yahoo!
1 billion PC users

1995



Mobile-Cloud
iPhone, Amazon AWS
2.5 billion mobile users

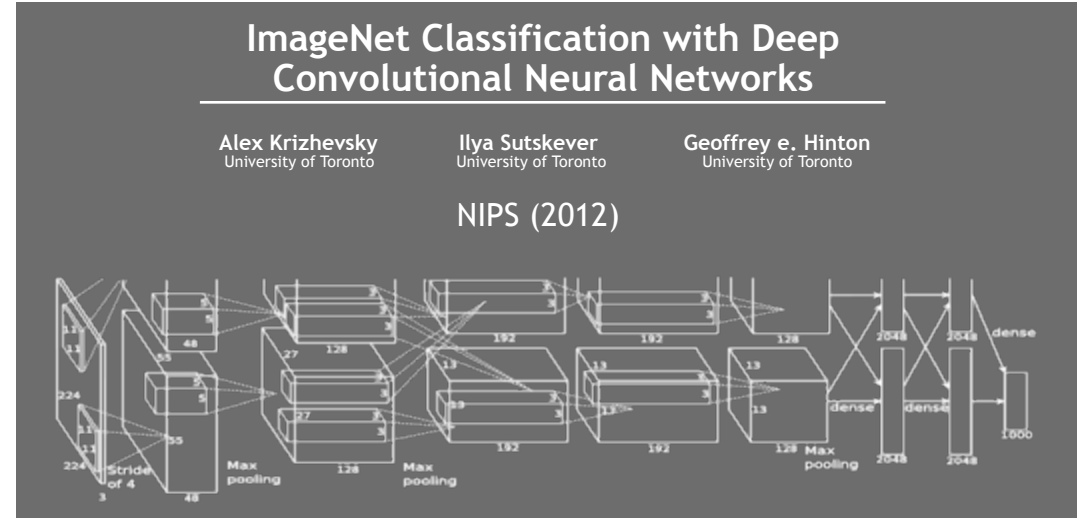
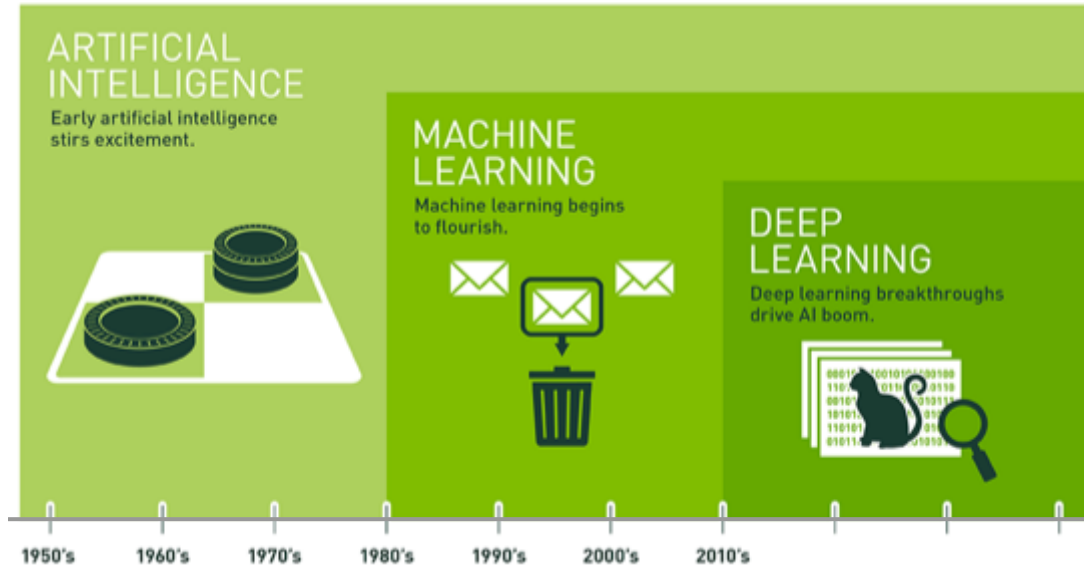
2005



AI & IOT
Deep Learning, GPU
100s of billions of devices

2015

AI AND DEEP LEARNING



Deep Learning

A NEW COMPUTING MODEL

Algorithms that learn from examples



Expert Written
Computer Program

Car

Vehicle

Coupe


TRADITIONAL APPROACH

Requires domain experts

Time consuming

Error prone

Not scalable to new problems



Deep Neural Network

Car

Vehicle

Coupe

DEEP LEARNING APPROACH

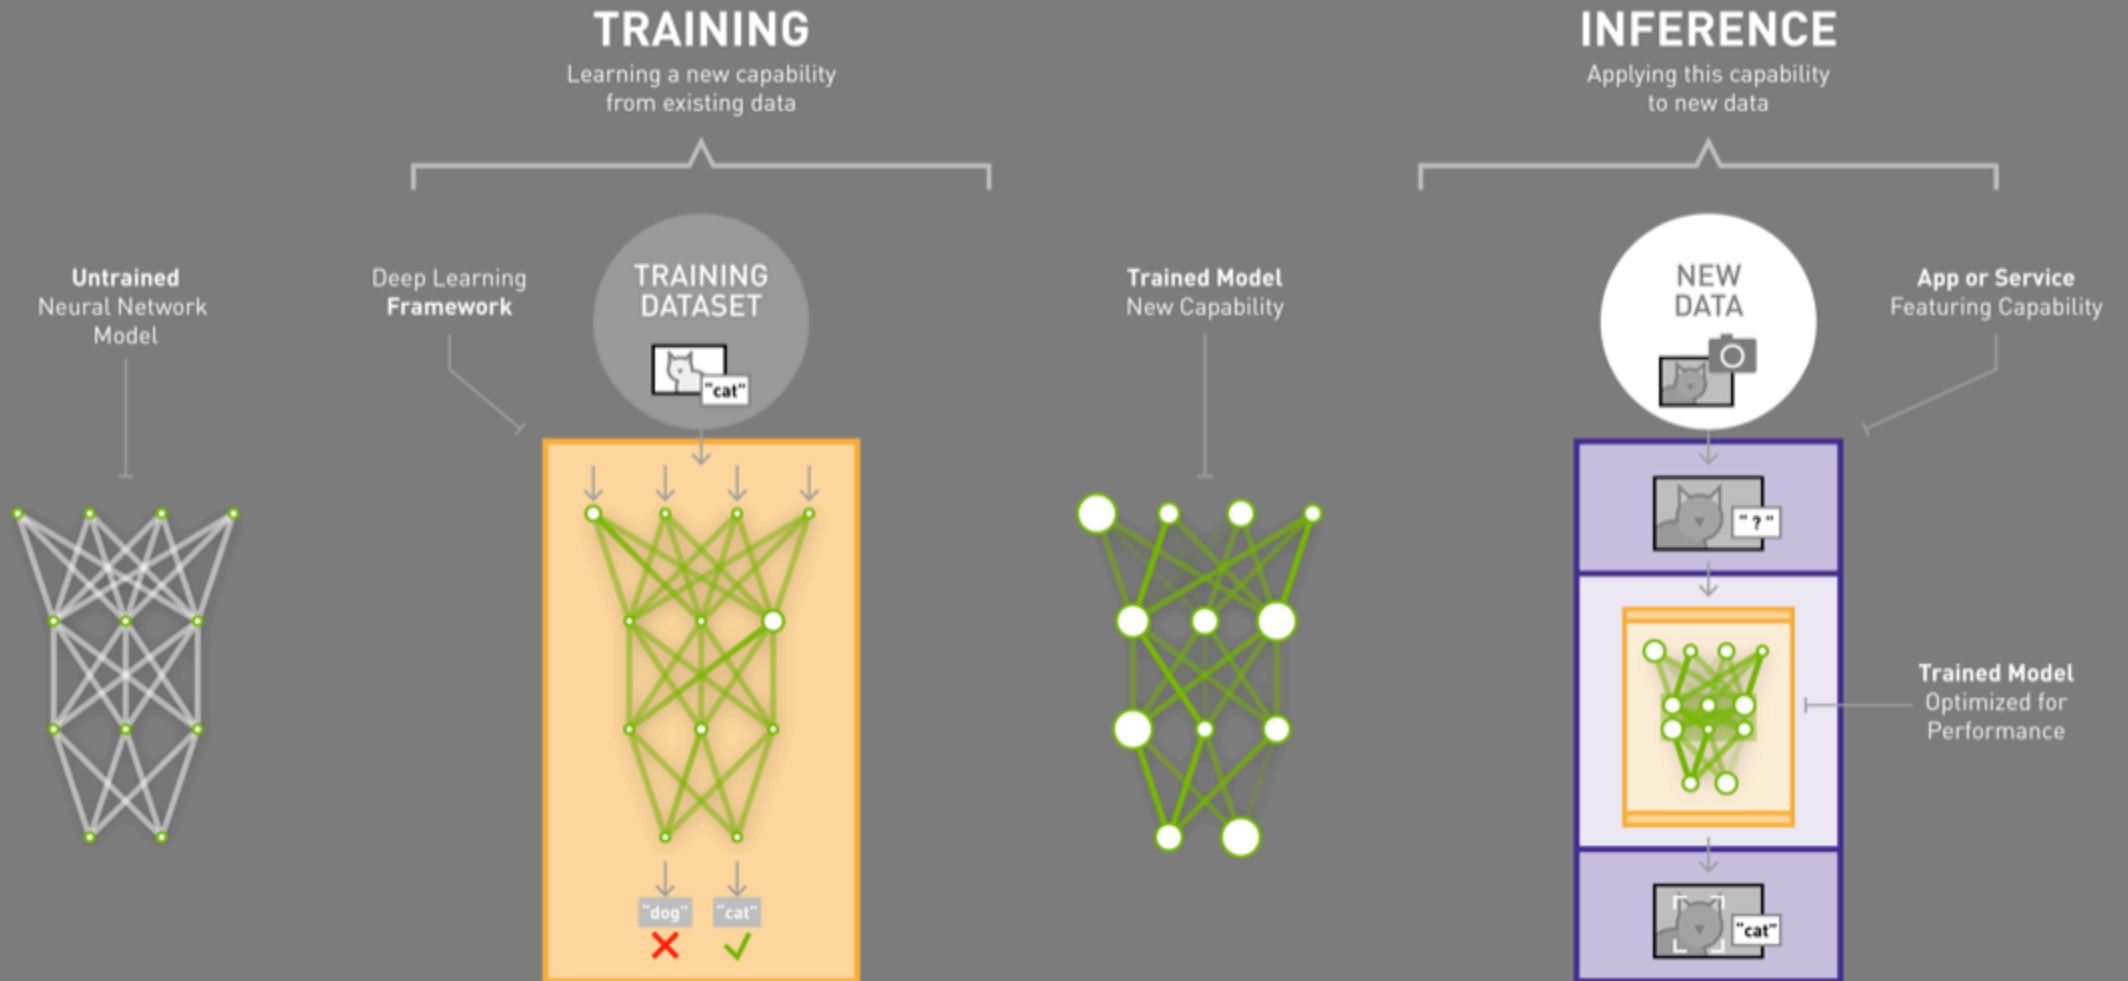
Learn from data

Easily to extend

Speedup with GPUs


DEEP LEARNING

How it works



THE BIG BANG OF MODERN AI





Evidence of an exploding AI ecosystem is everywhere. From AI conference attendees to the number of students studying AI to a massive growth in investments.





INDUSTRY EXAMPLES

ODD ONE OUT?




NVIDIA DRIVE FOR AUTONOMOUS VEHICLES

Autonomous vehicles will revolutionize the \$10 trillion transportation industry.

NVIDIA DRIVE is an open platform and enables researchers and programmers to develop new algorithms or adapt them for specific vehicles.

To train the network, data from all over the world needs to be collected and fed into an NVIDIA DGX supercomputer.

Simulation expands the training set and covers dangerous scenarios that can't be captured on the road. The trained model is deployed on an in-car supercomputer, for capabilities like pedestrian detection and driver monitoring.




NVIDIA CLARA AI — A MEDICAL IMAGING SUPERCOMPUTER

The latest breakthroughs of AI and computational imaging are giving radiologists incredibly powerful tools for early detection and treatment.

The NVIDIA Clara AI Toolkit supercharges existing instruments with state-of-the-art image reconstruction, object detection and segmentation, and visualization capabilities.

The American College of Radiology chose Clara AI to empower its 38,000 members to build, share, and validate AI algorithms in thousands of hospitals in the U.S.

TRADITIONAL MEDICAL IMAGING



ENHANCED WITH NVIDIA CLARA AI


IMAGING AND VISUALIZATION APPS

CUDA | CUDNN | TENSORRT | OGL |

RTX

GPU CONTAINERS | VGPU

NVIDIA GPU SERVER



SUPERCHARGING GENOMIC ANALYTICS

China's healthcare industry is turning to AI to address the needs of its elderly population. Genetics giant BGI—which has over 1PB of data—is classifying targetable peptides for personalized immunotherapy for cancer patients.

By running the open source RAPIDS data processing and machine learning libraries built on CUDA-X AI on an NVIDIA DGX-1 AI supercomputer, BGI sped up analysis 18x using cuDF, and 10x using XGBoost. The company is now expanding analysis to millions of peptide candidates.



IMPROVING DEMAND FORECASTS

With >100,000 different products in its 4,700 U.S. stores, the Walmart Labs data science team predicts demand for 500 million item-by-store combinations every week.

By performing forecasting with the open-source RAPIDS data processing and machine learning libraries built on CUDA-X AI on NVIDIA GPUs, Walmart speeds up feature engineering 100x and trains machine learning algorithms 20x faster, resulting in faster delivery of products, real-time reaction to shopper trends, and inventory cost savings at scale.






HOW TO GET STARTED

NGC: GPU-OPTIMIZED SOFTWARE HUB

Ready-to-run GPU Optimized Software, Anywhere

50+ Containers

DL, ML, HPC




15+ Model Training Scripts

NLP, Image Classification, Object Detection & more



NGC




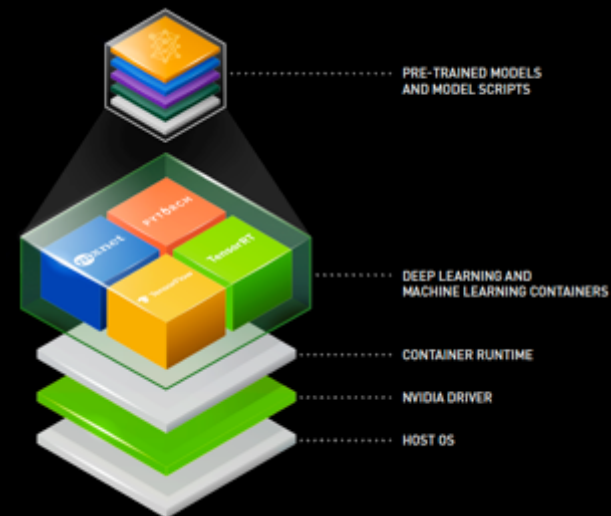
60 Pre-trained Models

NLP, Image Classification, Object Detection & more



Industry Workflows

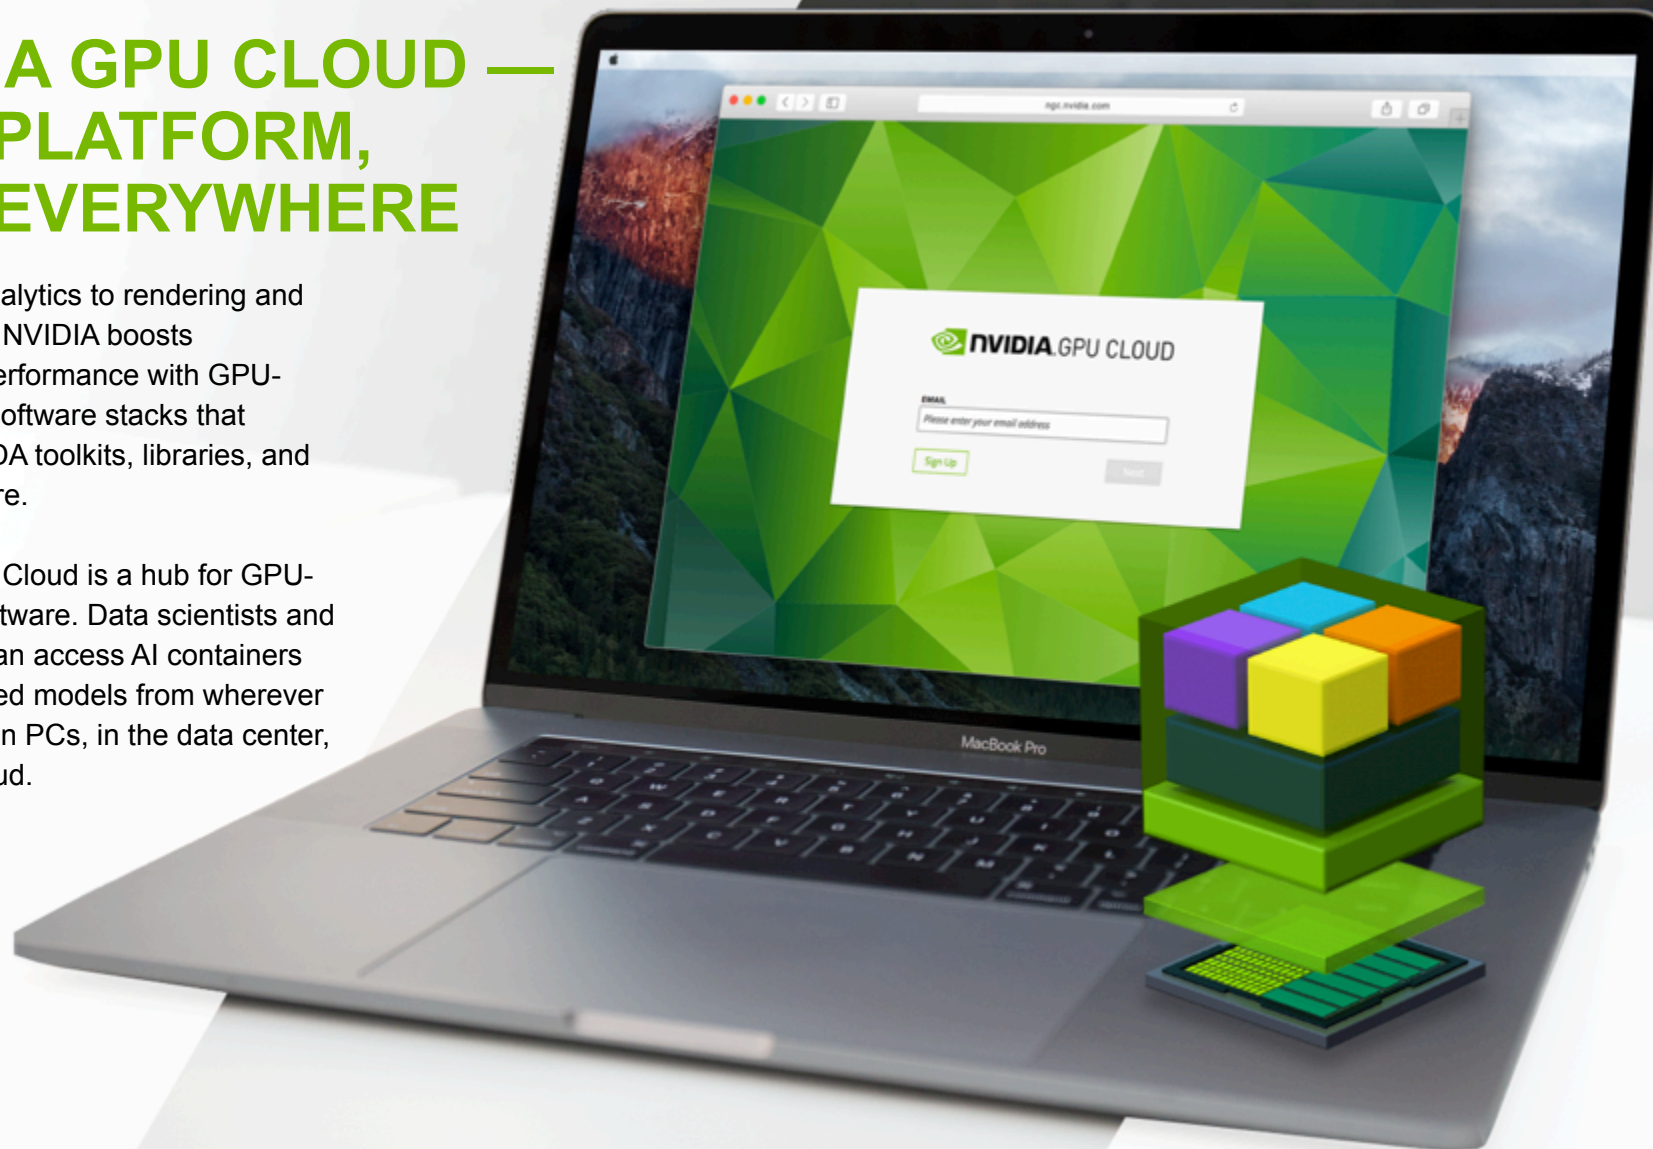
Medical Imaging, Intelligent Video Analytics



NVIDIA GPU CLOUD — ONE PLATFORM, RUN EVERYWHERE

From data analytics to rendering and visualization, NVIDIA boosts application performance with GPU-accelerated software stacks that combine CUDA toolkits, libraries, and top AI software.

NVIDIA GPU Cloud is a hub for GPU-optimized software. Data scientists and developers can access AI containers and pre-trained models from wherever they want – on PCs, in the data center, or via the cloud.



Science



Medical Imaging



Deep Learning



Machine Learning



Rendering and
Visualization



Hyperscale Inference

NVIDIA ENTERPRISE GPU PRODUCT FAMILY

Computing For Modern Enterprise Workloads

DATACENTER SERVERS



SPECIALIZED

(Max Performance)

MAINSTREAM

(Max Utility)

AI & HPC
Model Development

Rendering

Enterprise Application
Deployment



V100
Tensor Core, NVLink
32GB HBM2
250W/300W

RTX 8000/6000
RT Core, Tensor Core
48/24 GB GDDR6
250W/300W

T4
Tensor Core, RT Core
16GB GDDR6, 70W

RTX 8000/6000
Tensor Core
48/24 GB GDDR6

RTX 8000/6000/5000*/4000*
RT Core, Tensor Core

WORKSTATIONS



DATA SCIENCE

VISUALIZATION

AI Development

Design & Graphics

*Not designed & cannot be qualified for Servers

WORLD'S MOST POWERFUL AI TRAINING TOOL

Building amazing AI applications begins with training neural networks. NVIDIA DGX-2™ is the world's most powerful tool for AI training, uniting 16 GPUs to deliver 2 petaflops of training performance.

In December 2018, DGX-2 set six world records in the debut of MLPerf, a new set of industry benchmarks designed to test deep learning performance.

Image Classification

70 mins

Object Detection

176 mins

Object Instance Segmentation

14 mins

Translation (Recurrent)


10 mins

Translation (Non-Recurrent)

19 mins

Recommendation


0.4 mins



JETSON AGX AND ISAAC DELIVER AI TO ROBOTICS AND IOT INDUSTRY

Jetson™ AGX Xavier delivers the energy-efficient computational power needed for embedded systems like robots, drones, and smart cities. And the new Jetson Nano™ will enable millions more small, low-power AI systems for embedded IoT apps.

From Xavier to Nano, all of NVIDIA's AI computers run on the same CUDA-X AI software stack.



NVIDIA DEEP LEARNING INSTITUTE

Hands-on, self-paced and instructor-led training in deep learning and accelerated computing for developers

Request onsite instructor-led workshops at your organization: www.nvidia.com/requestdli


Take self-paced courses and electives online, view upcoming workshops, and learn about the University Ambassador Program: www.nvidia.com/dli




Deep Learning Fundamentals




Autonomous Vehicles



Medical Image Analysis



Genomics




Finance



Digital Content Creation




Game Development



Accel. Computing Fundamentals

More industry-specific training coming soon...



THANK YOU!

