



DEEP
LEARNING
INSTITUTE

DEEP LEARNING AND BEYOND

Will Ramey, Sr. Director, Global Head of Developer Programs, NVIDIA Corporation

ACCELERATED DATA SCIENCE

DATA ANALYTICS

Extracting insights from big data




MACHINE LEARNING

Learning from examples in the data



DEEP LEARNING

Automating feature engineering




GPU-ACCELERATED DATA SCIENCE

Use Cases in Every Industry



CONSUMER INTERNET

- Ad Personalization
- Click Through Rate Optimization
- Churn Reduction



OIL & GAS

- Sensor Data Tag Mapping
- Anomaly Detection
- Robust Fault Prediction




FINANCIAL SERVICES

- Claim Fraud
- Customer Service Chatbots/Routing
- Risk Evaluation



MANUFACTURING

- Remaining Useful Life Estimation
- Failure Prediction
- Demand Forecasting



HEALTHCARE

- Improve Clinical Care
- Drive Operational Efficiency
- Speed Up Drug Discovery



TELECOM

- Detect Network/Security Anomalies
- Forecasting Network Performance
- Network Resource Optimization (SON)



RETAIL

- Supply Chain & Inventory Management
- Price Management / Markdown Optimization
- Promotion Prioritization And Ad Targeting



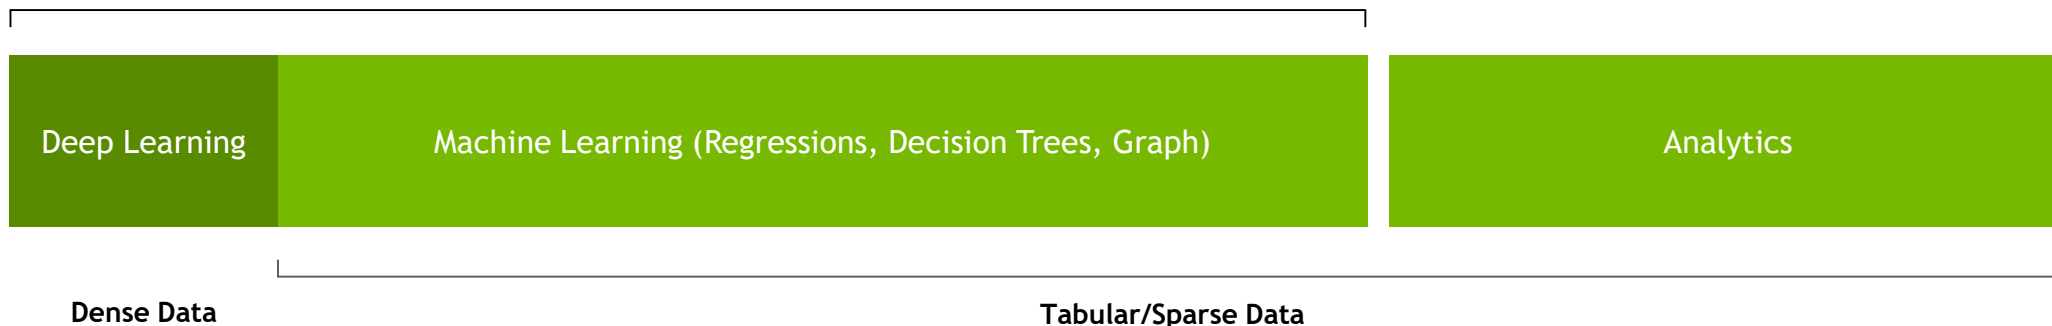
AUTOMOTIVE

- Personalization & Intelligent Customer Interactions
- Connected Vehicle Predictive Maintenance
- Forecasting, Demand, & Capacity Planning

BEYOND DEEP LEARNING

Opportunities to Accelerate Data Science

ARTIFICIAL INTELLIGENCE

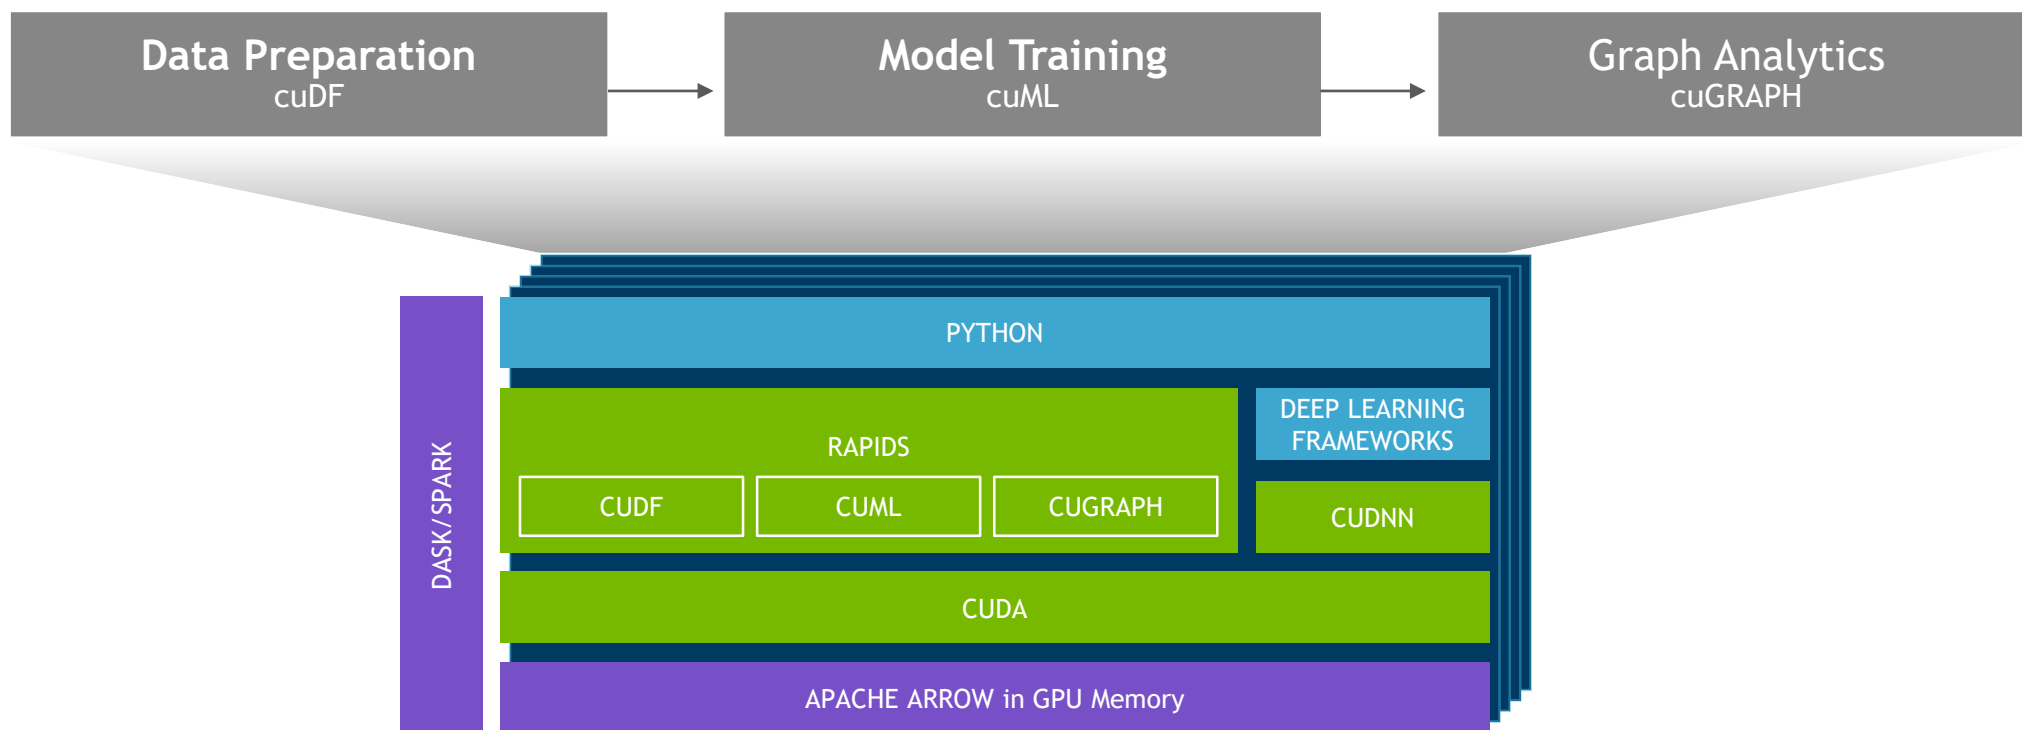


MACHINE LEARNING / DATA ANALYTICS

2.2 exabytes (2.2B GB) of data created daily - McKinsey
\$166B in 2018 revenues for big data and business analytics - IDC


RAPIDS DATA SCIENCE PLATFORM

GPU-accelerated Open Source Libraries



DL IS CRITICAL FOR INTERNET APPLICATIONS

Users Expect Intelligence in Services




Source: Jeff Dean, leads the Google Brain team, making machines intelligent. www.wsdm-conference.org/2016/slides/WSDM2016-Jeff-Dean.pdf

DEEP LEARNING IS SWEEPING ACROSS INDUSTRIES



Internet Services

Image/Video Classification
Speech Recognition
Natural Language Processing



Medicine

Cancer Cell Detection,
Diabetic Grading,
Drug Discovery



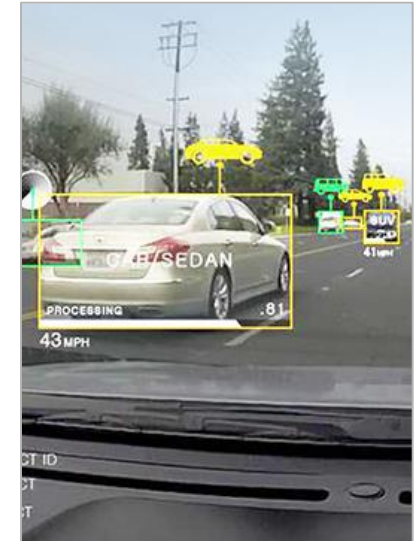
Media & Entertainment

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

THE EXPANDING UNIVERSE OF MODERN AI

"THE BIG BANG"

Big Data
Algorithms
GPU

RESEARCH



FRAMEWORKS



AI-as-a-PLATFORM



START-UPS



Personal Assistants
conversational interface



Agriculture
crop-yield optimization



Tech
visual recognition platform



Genomics
genetic interpretation



Automotive
computer vision



eCommerce & Medical
recommendation engines



Tech
computer vision



Geospatial
predictions from images



Tech
AI-as-a-service



Waste Management
sorting robots



Medical
diabetic retinopathy




Education
teaching robots

4,000+ AI START-UPS
\$33B IN FUNDING





Source: Crunchbase & Pitchbook

INDUSTRY LEADERS




WHAT PROBLEM ARE YOU SOLVING?

Defining the AI/DL task

INPUTS	BUSINESS QUESTIONS	AI / DL TASK	EXAMPLE OUTPUTS		
			HEALTHCARE	RETAIL	FINANCE
 Text Data  Images  Video  Audio	Is “it” present or not?	Detection	Cancer Detection	Targeted Ads	Cybersecurity
	What type of thing is “it”?	Classification	Image Classification	Basket Analysis	Credit Scoring
	To what extent is “it” present?	Segmentation	Tumor Size / Shape Analysis	Build 360° Customer View	Credit Risk Analysis
	What is the likely outcome ?	Prediction	Survivability Prediction	Sentiment & Behavior Recognition	Fraud Detection
	What will likely satisfy the objective?	Recommendations	Therapy Recommendation	Recommendation Engine	Algorithmic Trading

CAMBRIAN EXPLOSION

Convolutional Networks



Encoder/Decoder



ReLU



BatchNorm



Concat




Dropout



Pooling

Recurrent Networks



LSTM



GRU



Beam Search



WaveNet



CTC



Attention

Generative Adversarial Networks



3D-GAN



MedGAN



Conditional GAN




Coupled GAN




Speech Enhancement GAN

Reinforcement Learning



DQN




Simulation



DDPG

New Species



Mixture of Experts



Neural Collaborative Filtering




Block Sparse LSTM

DEEP LEARNING APPLICATION DEVELOPMENT



DEEP LEARNING APPLICATION DEVELOPMENT

Untrained
Neural Network
Model



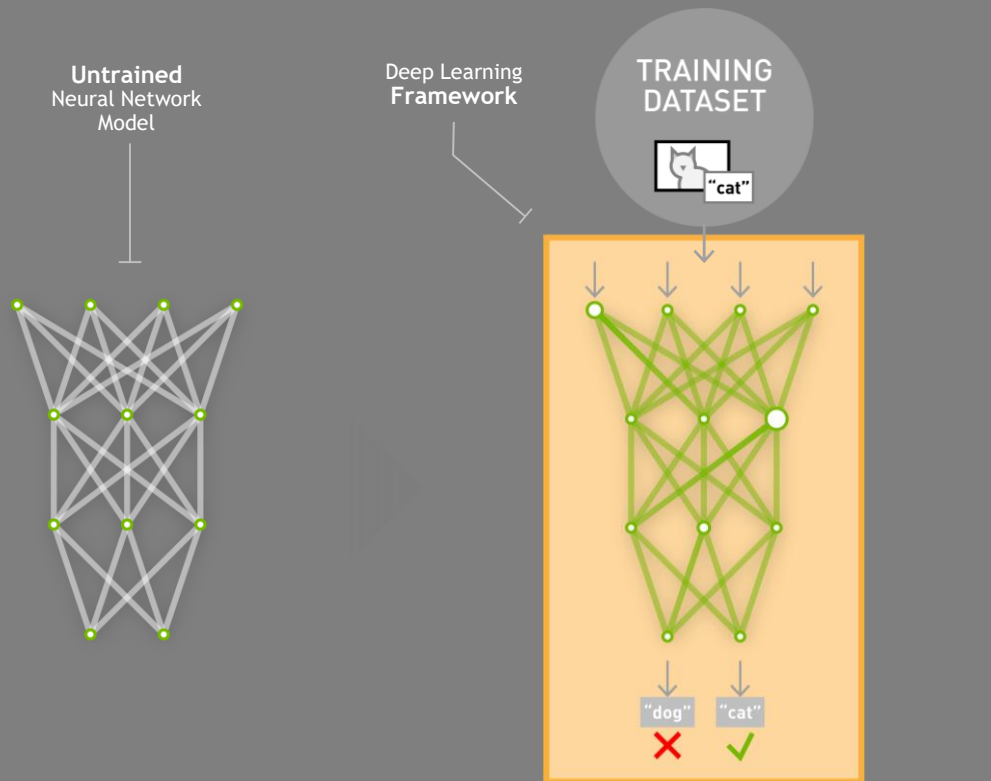
TRAINING
DATASET



DEEP LEARNING APPLICATION DEVELOPMENT

TRAINING


Learning a new capability
from existing data



DEEP LEARNING APPLICATION DEVELOPMENT

TRAINING

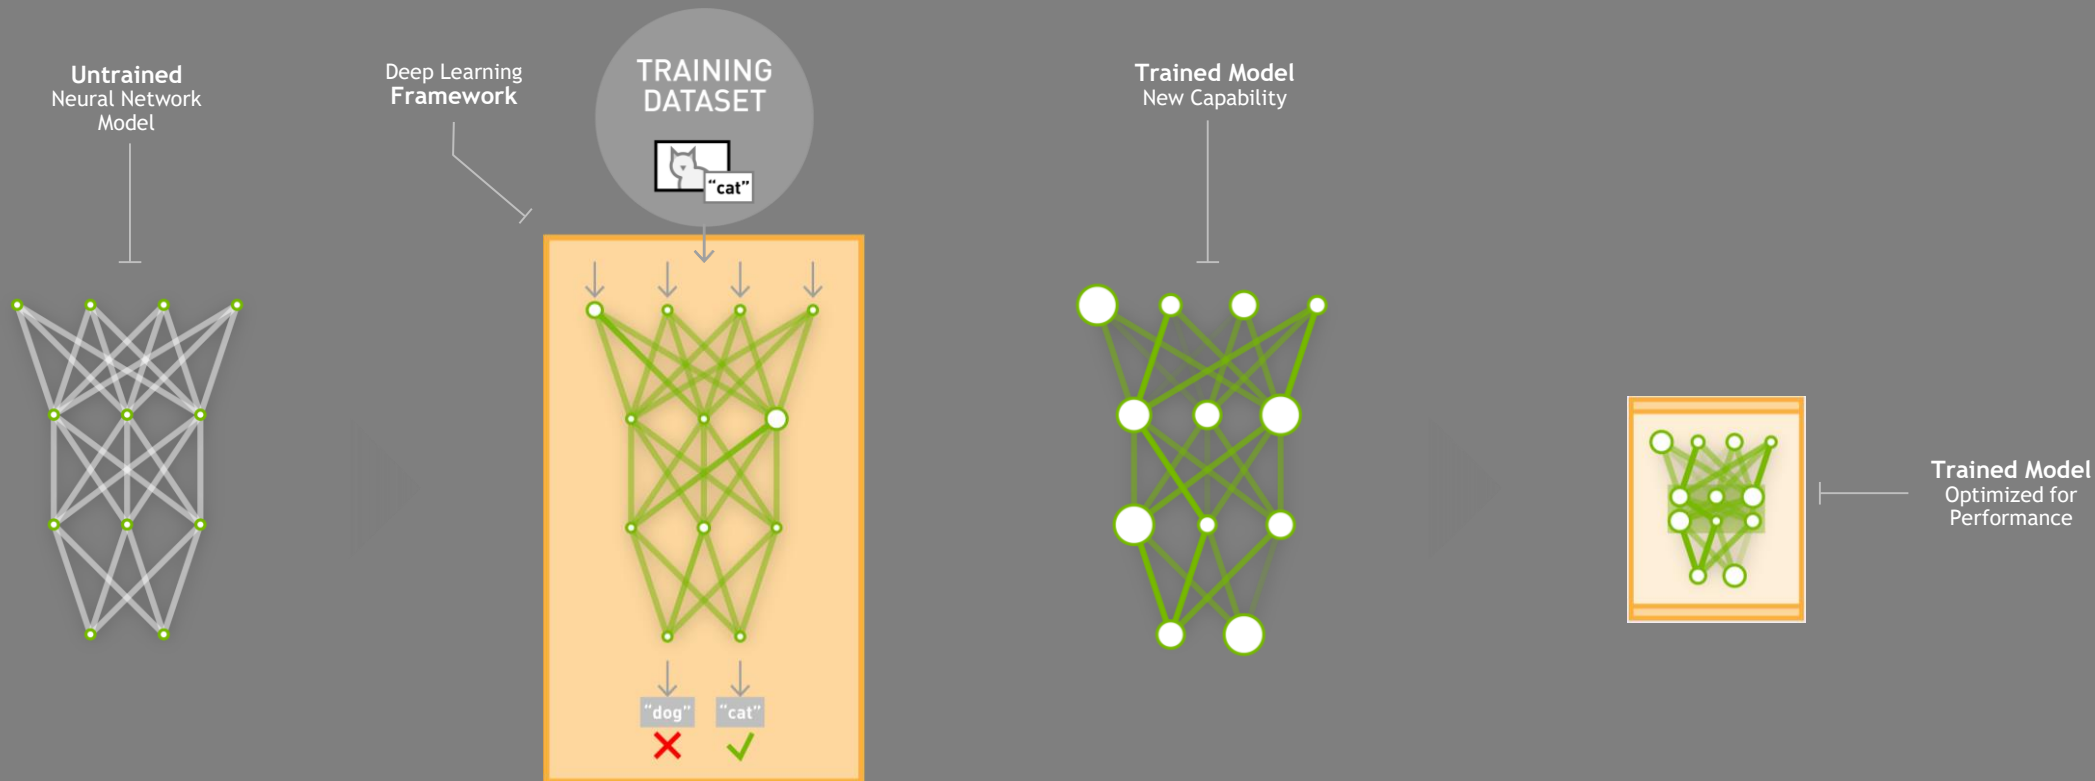
Learning a new capability
from existing data



DEEP LEARNING APPLICATION DEVELOPMENT

TRAINING

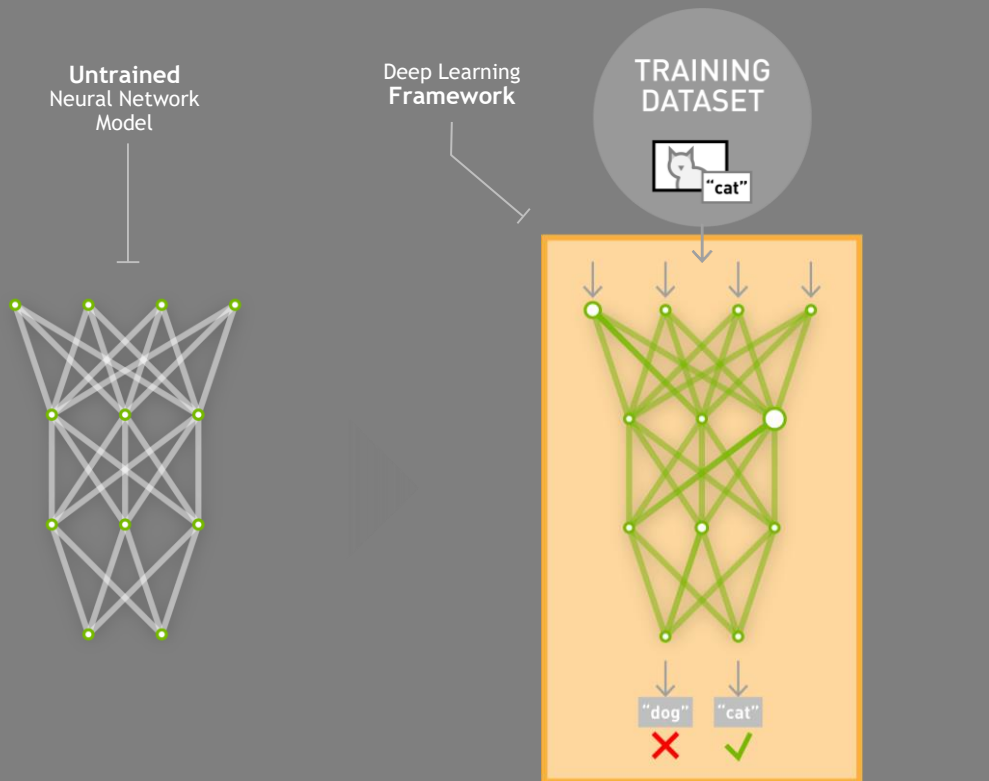
Learning a new capability
from existing data



DEEP LEARNING APPLICATION DEVELOPMENT

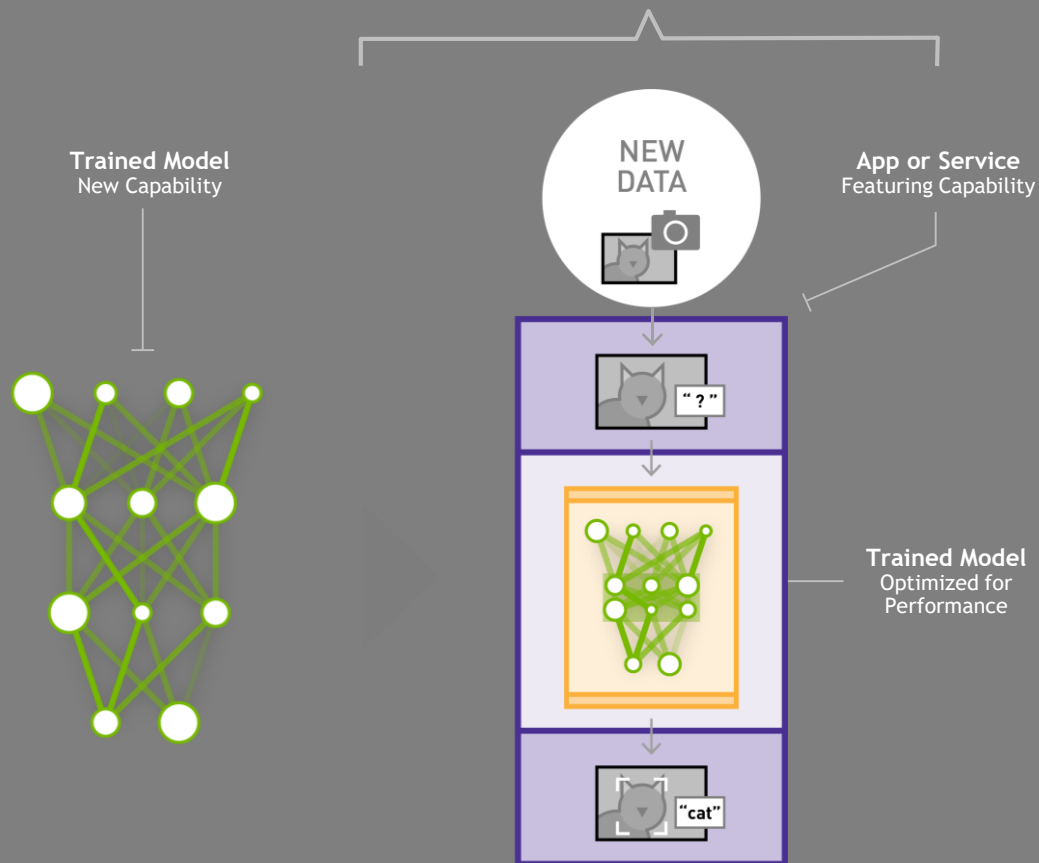
TRAINING

Learning a new capability
from existing data




INFERENCE

Applying this capability
to new data



CHALLENGES

DEEP LEARNING NEEDS	WHY
Data Scientists	New computing model
Latest Algorithms	Rapid evolving
Fast Training	Impossible -> Practical
Deployment Platforms	Must be available everywhere



ADVANCE YOUR DEEP LEARNING KNOWLEDGE AT GTC

Don't miss the world's most important event for GPU developers

WWW.GPUTECHCONF.COM

NVIDIA DEEP LEARNING INSTITUTE

Hands-on Training for Data
Scientists and Software Engineers


Helping the world to solve challenging
problems using AI and deep learning

On-site workshops and online courses
presented by certified instructors

Covering complete workflows for proven
application use cases

Fundamentals, Autonomous Vehicles,
Finance, Healthcare, Video Analytics,
IoT/Robotics, ...

www.nvidia.com/dli



Inception: Virtual Accelerator for AI Startups


BENEFITS

AI Expertise



DEEP
LEARNING
INSTITUTE


Go-To-Market Support



Technology Access



MEMBERSHIP



www.nvidia.com/inception

IMPACT

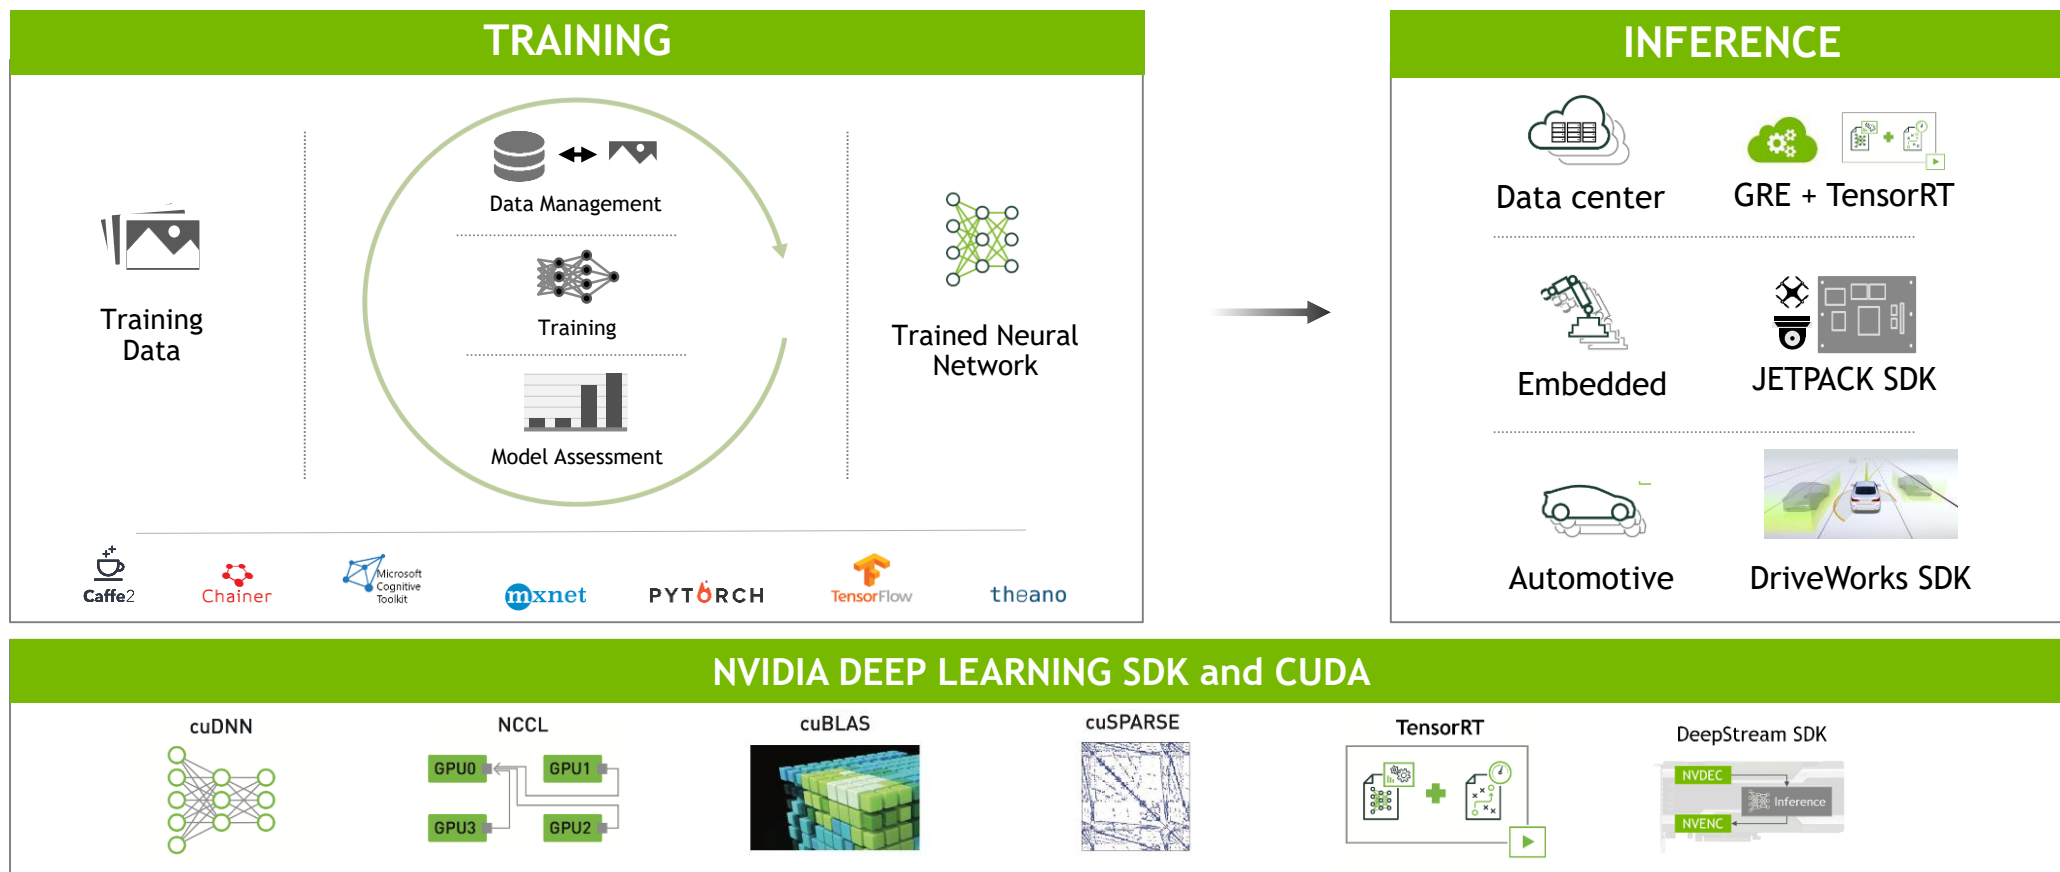
Showcasing Innovation



Creating a Global Community



NVIDIA DEEP LEARNING SOFTWARE PLATFORM



NVIDIA GPU CLOUD

Simple Access to a Comprehensive Catalog of GPU-accelerated Software

Discover 30 GPU-Accelerated Containers

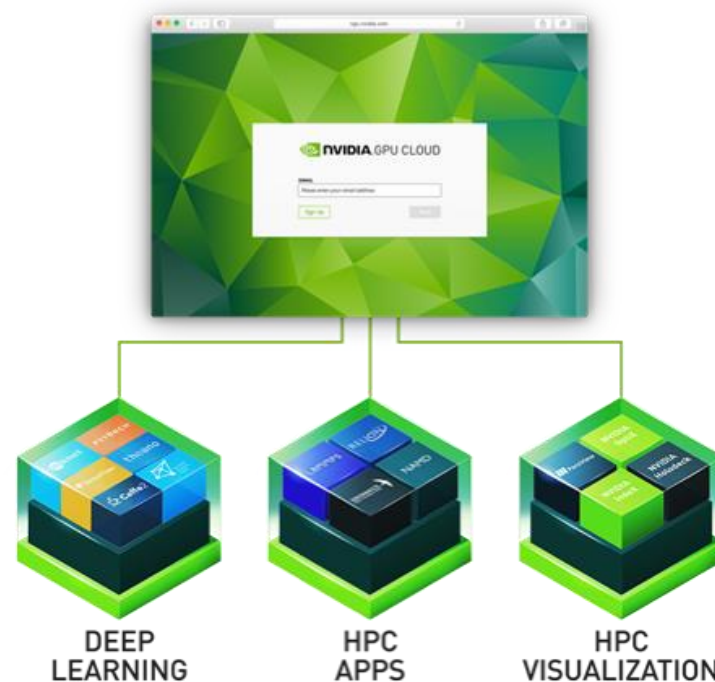
Deep learning, third-party managed HPC applications, NVIDIA HPC visualization tools, and partner applications

Innovate in Minutes, Not Weeks

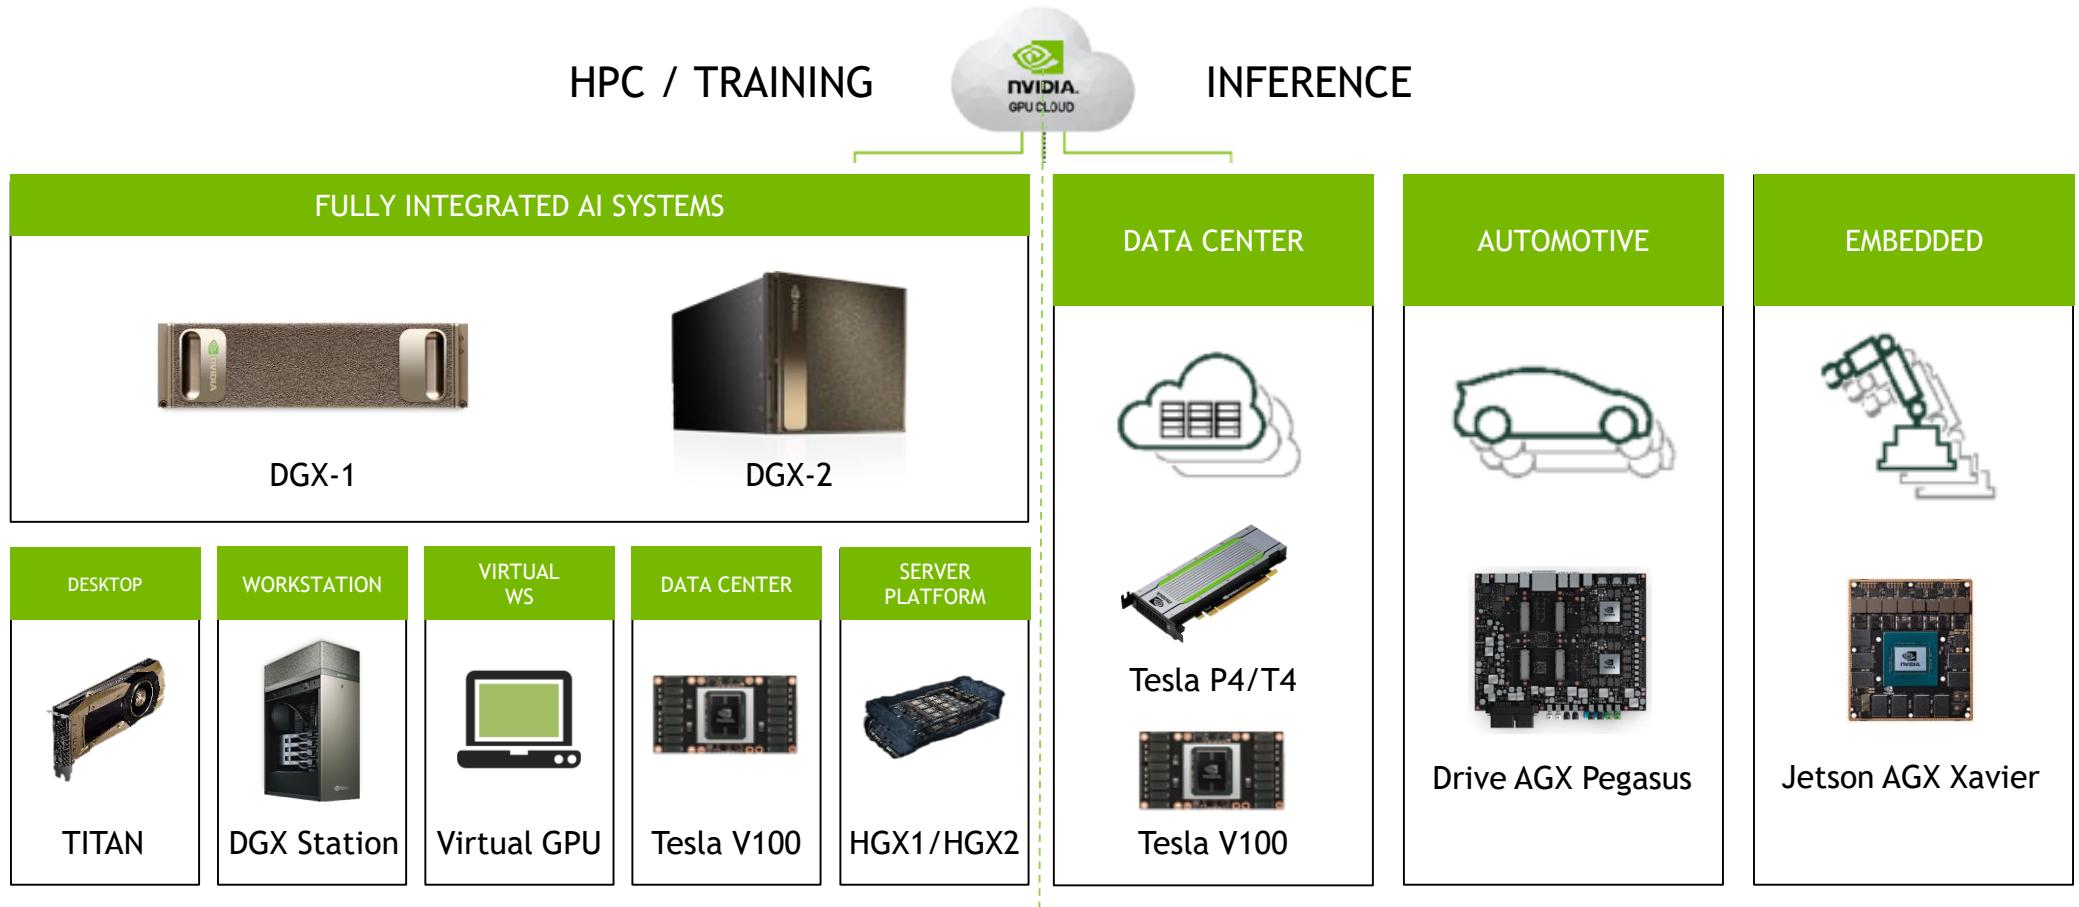
Get up and running quickly and reduce complexity

Access from Anywhere

Use on PCs with NVIDIA Volta or Pascal™ architecture GPUs, NVIDIA DGX Systems, and supported cloud providers



END-TO-END PRODUCT FAMILY



SOLUTIONS

DEEP LEARNING NEEDS	SOLUTIONS
Data Scientists	Deep Learning Institute, GTC
Latest Algorithms	NGC, GPU Accelerated Frameworks, DL SDK
Fast Training	DGX, V100/T100, TITAN V
Deployment Platforms	TensorRT, V100/T4, Drive AGX, Jetson AGX

READY TO GET STARTED?



Project Checklist

What problem are you solving, what are the AI/DL tasks?

What data do you have/need, how is it labeled?

Which tools & environment will you use?

On what platform(s) will you train and deploy?



WHAT'S NEXT?

Work through DLI online courses

www.nvidia.com/dli

Review examples of AI in action

news.developer.nvidia.com


Listen to the NVIDIA AI Podcast

blogs.nvidia.com/ai-podcast

Register for GTC near you

www.nvidia.com/gtc

THANK YOU



DEEP
LEARNING
INSTITUTE

www.nvidia.com/dli



TRAINING YOU TO SOLVE
THE WORLD'S MOST
CHALLENGING PROBLEMS



NVIDIA DEEP LEARNING INSTITUTE

The NVIDIA Deep Learning Institute (DLI) offers hands-on training in artificial intelligence (AI) and accelerated computing to solve real-world problems. Start a self-paced online course at your convenience or take the course in an instructor-led workshop.

No matter how you learn best, DLI gives you access to a fully configured, GPU-accelerated workstation in the cloud, complete with software tools, neural networks, and datasets. All you need is a web browser and Internet connection to get started.

www.nvidia.com/dli



DEEP
LEARNING
INSTITUTE