


# HPC + AI


## THE CONVERGENCE OF SUPERCOMPUTING AND AI IN A POST-MOORE'S LAW WORLD

Piero Altoe'

IBD, Accelerated Computing




# GRAND CHALLENGES REQUIRE MASSIVE COMPUTING




## REINVENTING THE LI-ION BATTERY

3M Node Hours | 7 Days on Titan



## CLOUD RESOLVING CLIMATE SIMULATIONS

100M Node Hours | 840 Days on Piz Daint




## UNDERSTANDING HIV'S STRUCTURE

10M node Hours | 16 Days on BlueWaters


# BAD TIMING

## THE SLOW DEATH OF MOORE'S LAW



Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten  
New plot and data collected for 2010-2015 by K. Rupp

# TOP500 EFFECTS



**SOMETHING NEW:**  
**AI + HPC = *REVOLUTION***

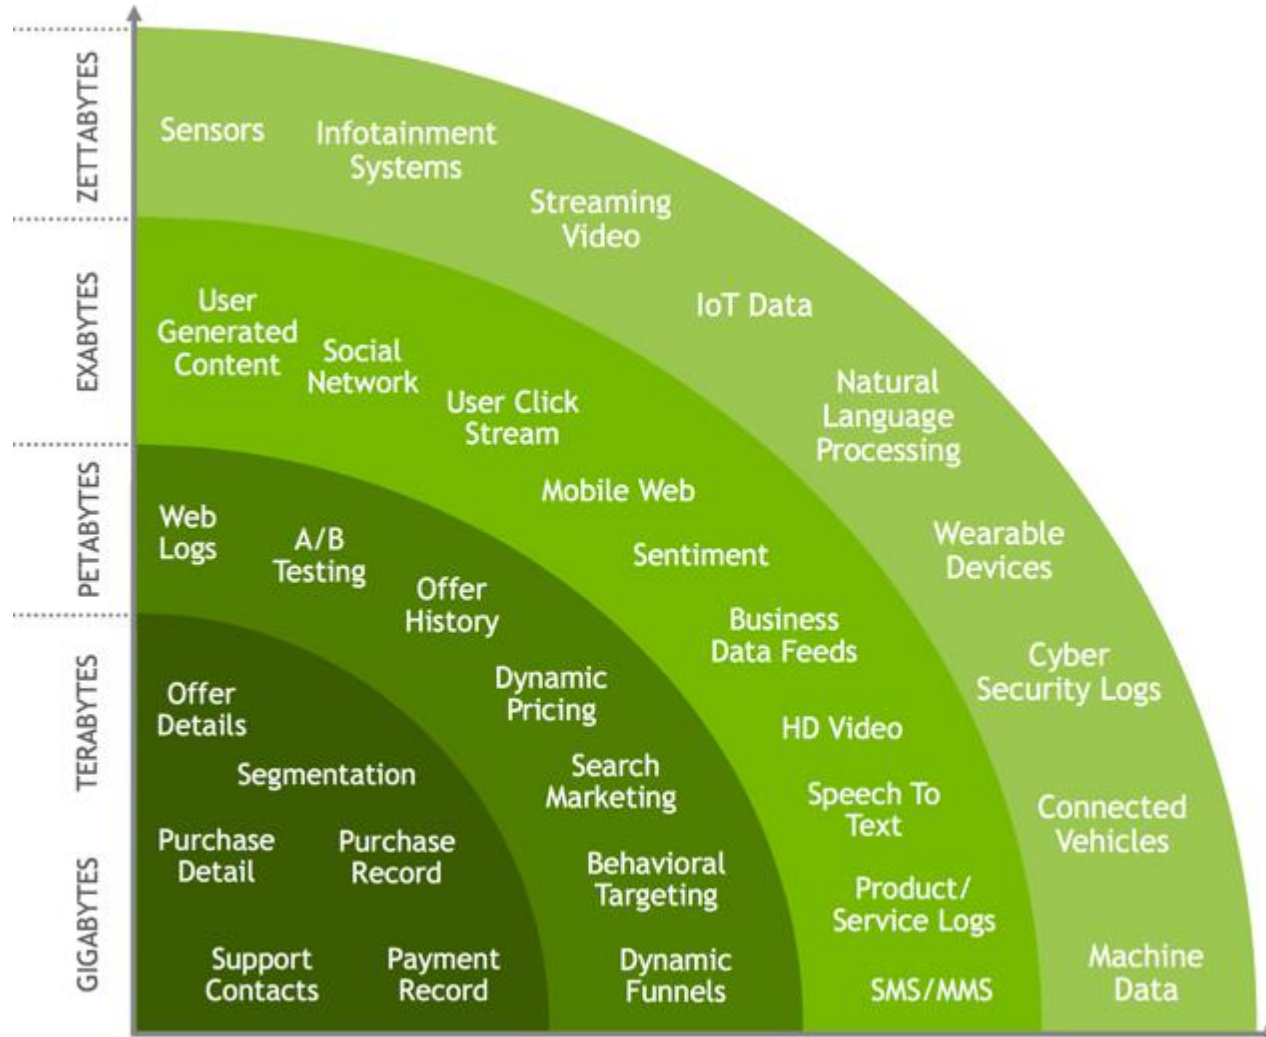


# INGREDIENTS: BIG DATA







(We are the sensors for our cloud service providers)

# INGREDIENTS: BIG DATA



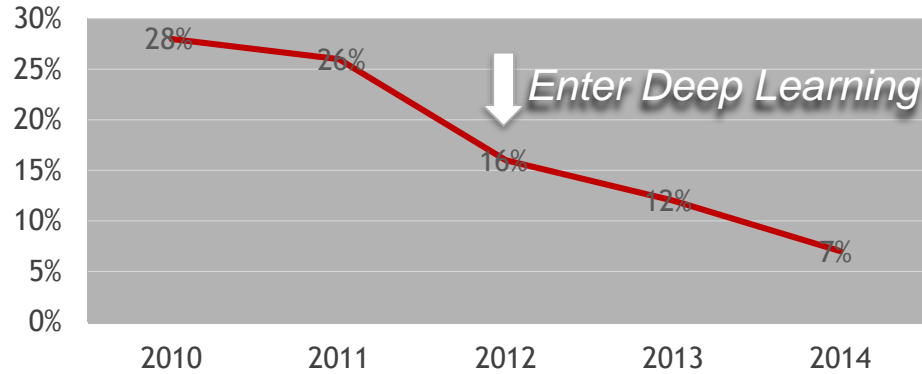
# INGREDIENTS: AI ALGORITHMS






# NOW, JUST ADD HPC AND STIR...

Classification Error Rates



GPU Entries



*Classification Task:*

*1.2M images • 1000 object categories*

Team	Date	Top-5 Test Error
GoogLeNet	2014	6.66%
Baidu Deep Image	01/12/2015	5.98%
Baidu Deep Image	02/05/2015	5.33%
Microsoft	02/05/2015	4.94%
Google	03/02/2015	<u>4.82%</u>
Baidu Deep Image	03/17/2015	4.83%

*Trained Human Performance: 5.1%*

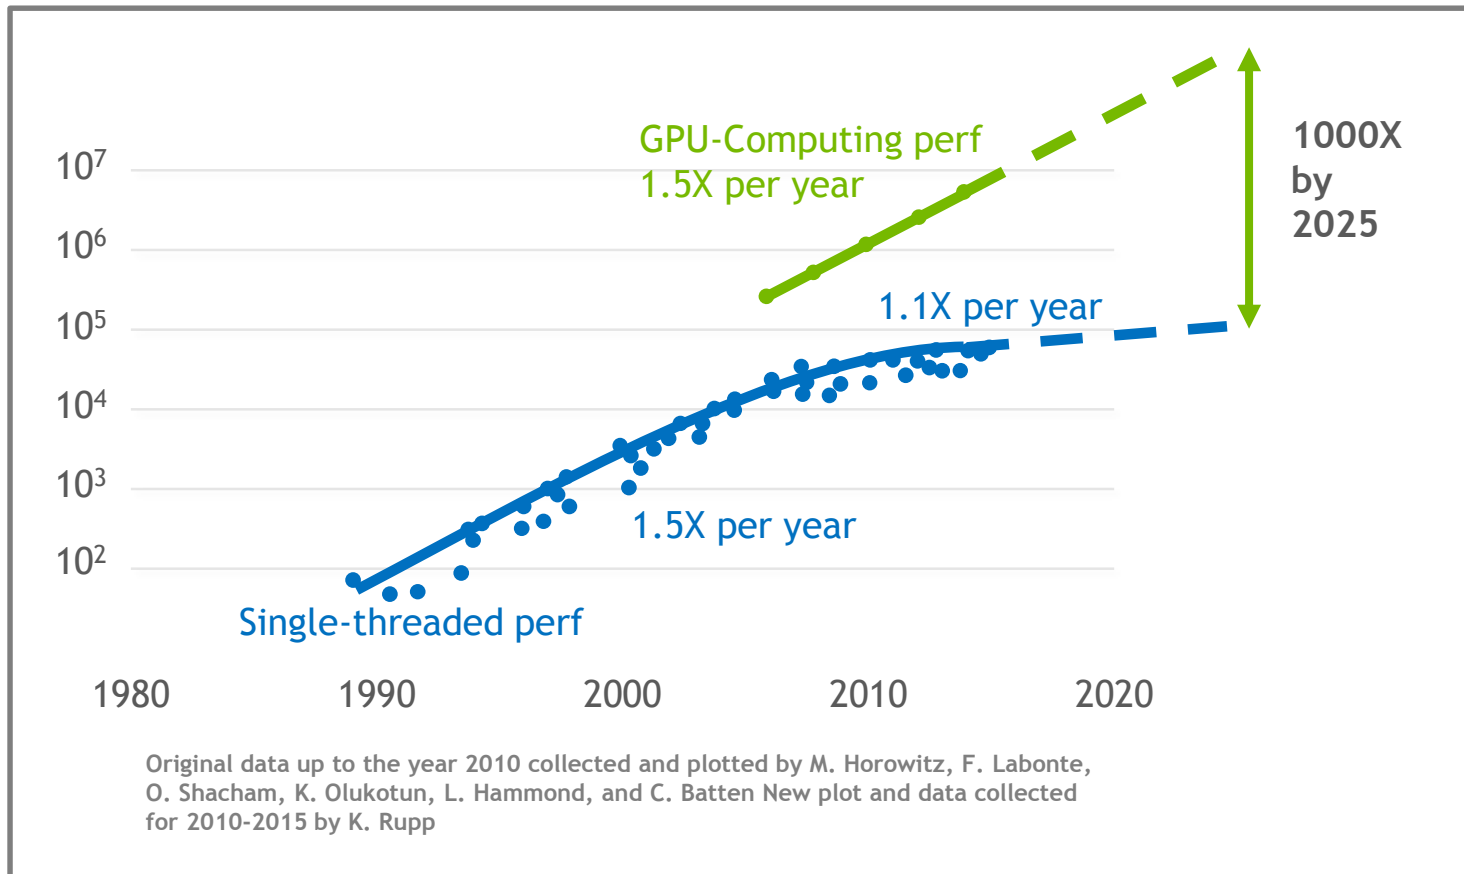
IMAGENET

# ALGORITHMS + BIG DATA + GPUS = THE BIG BANG OF MODERN AI




recognition/classification -> recursion/time series -> generative

# AI: A CHEAT IS COMING FROM THE GPU PARADIGM




# GPU COMPUTING PERFORMANCE OVER TIME

AI on a super-Moore's Law progression



4x in 3 years




12x in 3 years  
(65x in 5 years)

Amber dataset: Cellulose NVE; GoogLeNet dataset: Imagenet

# 2018: 10X AI GAIN IN ONE YEAR

DGX-1, SEP'17



DGX-2, Q3'18



PyTorch Stack: Time to Train FAIRSEQ

software improvements across the stack including NCCL, cuDNN, etc.


WHAT DOES THIS HAVE TO DO WITH  
*SCIENCE?*

(HPC + AI = ?)




# BIG DATA IN SCIENCE

Big Science ingests/outputs Big Data



Large Hadron Collider









Johns Hopkins  
Turbulence Database



Square Kilometer Array

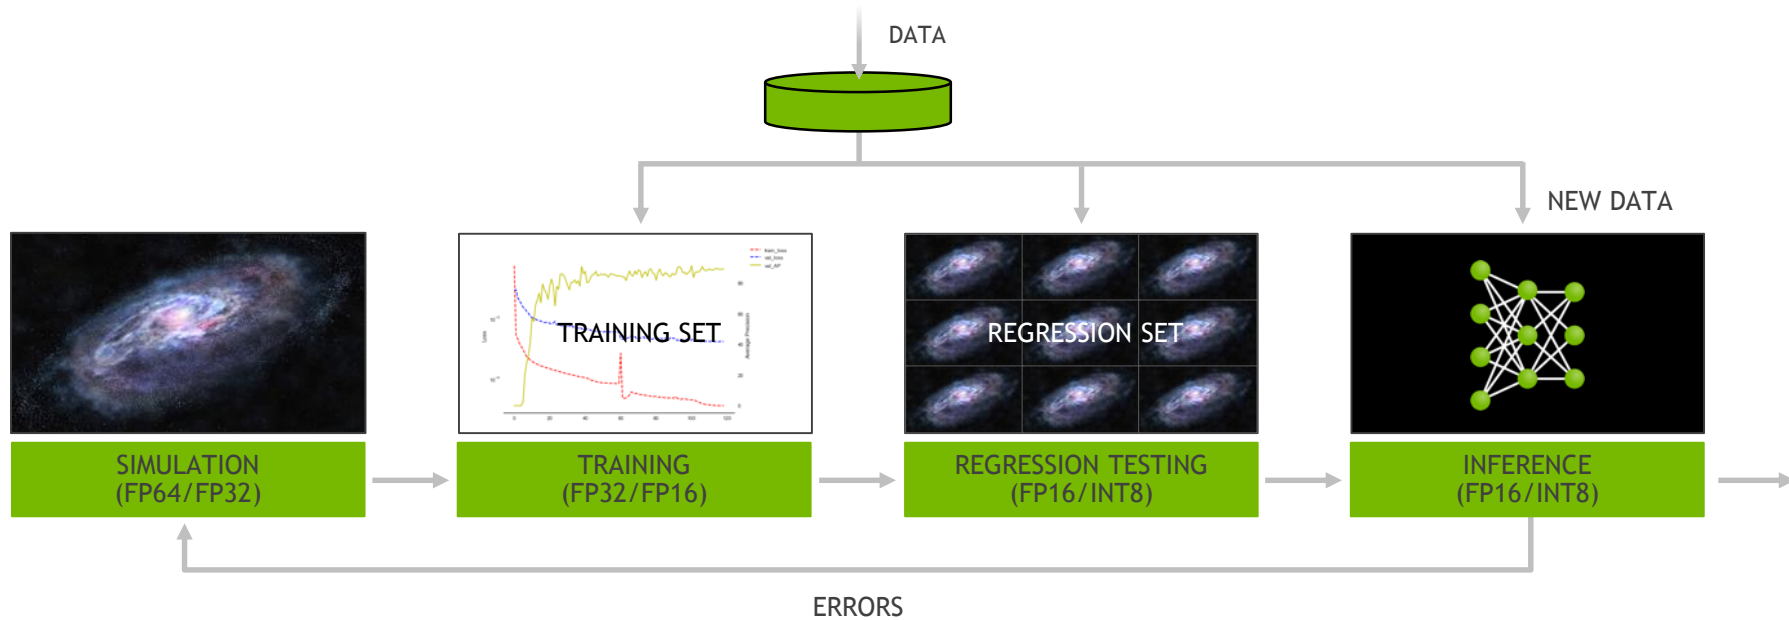
# 2015: USING NUMERIC SIMULATIONS TO TRAIN AI



# 2015: USING NUMERIC SIMULATIONS TO TRAIN AI







# AI WORKFLOW FOR HPC





# DEEP LEARNING COMES TO HPC





# APPROXIMATING QUANTUM CHEMISTRY

## Background

Developing a new drug costs \$2.5B and takes 10-15 years. Quantum chemistry (QC) simulations are important to accurately screen millions of potential drugs to a few most promising drug candidates.

## Challenge

QC simulation is computationally expensive so researchers use approximations, compromising on accuracy. To screen 10M drug candidates, it takes 5 years to compute on CPUs.

## Solution

Researchers at the University of Florida and the University of North Carolina leveraged GPU deep learning to develop ANAKIN-ME, to reproduce molecular energy surfaces with super speed (microseconds versus several minutes), extremely high (DFT) accuracy, and at 5-6 orders of magnitude lower cost.


## Impact

Faster, more accurate screening at far lower cost






# NEURAL NETWORK MODEL APPROACH

J. Smith, O.I., A. Roitberg. *Chem. Sci.*, 2017, **8**, 3192-3203




Training set: ~20M DFT data points.  
Molecules with 1 to 8 atoms from GDB database



Accessible molecular systems

# CORRELATION VALIDATION




All models are wrong, but some are  
useful.

— *George E. P. Box* —

AZ QUOTES

# CORRELATION VALIDATION



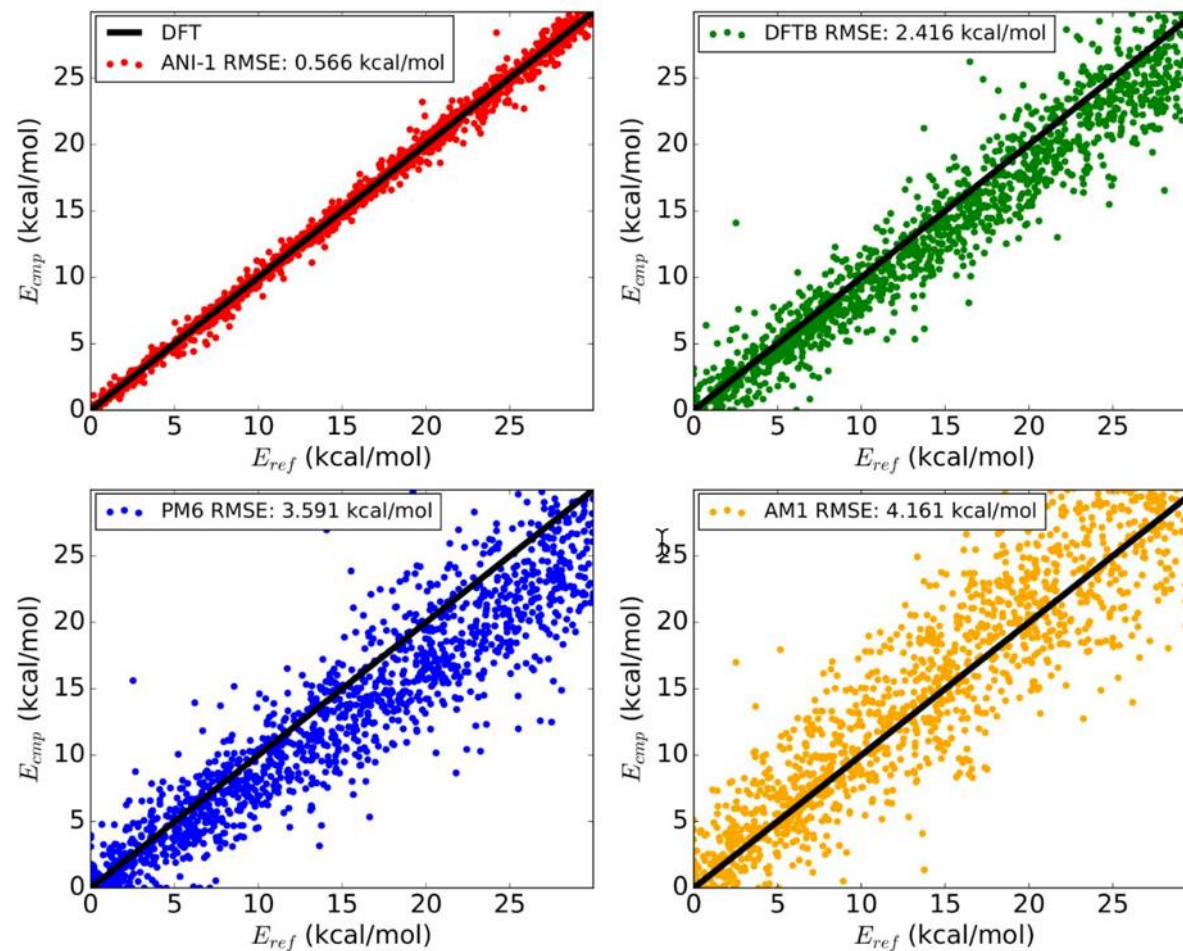
## Total energy correlation

ANI-1 vs. DFT

(131 molecules with 10 heavy atoms, 8200 total molecules + conformations) [units: kcal/mol]

J. Smith, O.I., A. Roitberg. **Chem. Sci.**, 2017, **8**, 3192-3203

# CORRELATION VALIDATION



# THE SHAPE OF AI SUPERCOMPUTING

# DESIGNED FOR AI SUPERCOMPUTING

Piz Daint

Europe's Fastest  
Supercomputer

19.6 PetaFLOPS

10.4 GF/W

Powered by Tesla P100

Largest simulation of Universe  
along with ORNL Titan

Set DL Scaling record with  
Microsoft Cognitive Toolkit






# DESIGNED FOR AI SUPERCOMPUTING

Tsubame 3

#1 Green500 System, June 2017



Powered by 2160 P100s



## Tokyo Tech's Tsubame 3 will be AI/HPC hybrid

20 February 2017 | By Peter Judge

**DatacenterDynamics**  
The Business of Data Centers

## JAPAN KEEPS ACCELERATING WITH TSUBAME 3.0 AI SUPERCOMPUTER

February 17, 2017 | Timothy Prickett Morgan

**THE NEXT PLATFORM**

## Next-Generation TSUBAME Will Be Petascale Supercomputer for AI


Michael Feldman | February 18, 2017 00:04 CET

**TOP 500**  
The List

“NVIDIA’s broad AI ecosystem will enable Tokyo Tech to begin training TSUBAME3.0 immediately to help us more quickly solve some of the world’s once unsolvable problems.”

- Satoshi Matsuoka, Prof Computer Science, TiTech & Project lead Tsubame 3

# DESIGNED FOR AI SUPERCOMPUTING



Most Powerful AI Supercomputer in Japan  
4,352 Tesla V100 GPUs  
37 PetaFLOPS FP64 HPC Performance  
0.55 ExaFLOPS AI Performance

**ABCI SUPERCOMPUTER** TO BE  
INSTALLED AT THE KASHIWA II CAMPUS OF  
THE UNIVERSITY OF TOKYO

# VOLTA TENSOR CORE GPU FUELS WORLD'S FASTEST SUPERCOMPUTER

Fused HPC and AI Computing In a Unified Platform



Summit Supercomputer  
Oakridge National Labs

AI: 3 Exaflops  
HPC: 122 Petaflops

## Genomics (CoMet)

World's First Exascale Run

Finding Genes-to-disease Connection

Same accuracy as FP64 w/ Tensor Core



## Quantum Chemistry (QMCPack)

Simulate New Materials


High-Temperature Semiconductors




Measured performance: Summit node vs Titan node

# AI SUPERCOMPUTING IS HERE


Extending The Reach of HPC By Combining Computational & Data Science



Turbulent Flow



Molecular Dynamics



Structural Analysis




N-body Simulation

**COMPUTATIONAL SCIENCE**



“What’s happening?”



“Is there cancer?”



“Next move?”



“What does she mean?”

**DATA SCIENCE**



Drug Discovery



Clean Energy



Understanding Universe



Monitoring Climate Change

**COMPUTATIONAL & DATA SCIENCE**



# GTC EUROPE

## 9-11 October MUNICH



**GPU** TECHNOLOGY  
CONFERENCE

AGENDA ▾ ATTEND ▾ PRESENT ▾ EXHIBIT ▾ MORE ▾


EUROPE ▾

OCTOBER 9-11, 2018

## MAKE THE MOST OF YOUR GTC EUROPE EXPERIENCE

Build your agenda with talks & more.

[VIEW SESSIONS](#)



JOIN OUR NVIDIA JETSON™ AGX XAVIER DEVELOPER  
MEETING.

[LEARN MORE](#)

# DEEP LEARNING INSTITUTE



DEEP LEARNING AI INDUSTRIES FOR DEVELOPERS PRODUCTS SOLUTIONS EDUCATION AI STARTUPS

Education ONLINE TRAINING WORKSHOPS UNIVERSITY PROGRAM MORE TRAINING PARTNERS RESOURCES CONTACT US

**NVIDIA  
DEEP LEARNING  
INSTITUTE**

Training You to Solve the World's Most Challenging Problems

DLI NOW OFFERS DEVELOPER CERTIFICATION IN AI AND ACCELERATED COMPUTING.

LEARN MORE

The NVIDIA Deep Learning Institute (DLI) offers hands-on training in AI and accelerated computing to solve real-world problems. Designed for developers, data scientists, and researchers, DLI content is available in three formats:

## ONLINE COURSES

DLI online courses teach you how to implement and deploy an end-to-end project in eight hours. Online courses can be taken anytime, anywhere, with access to a fully configured GPU-accelerated workstation in the cloud.

## ONLINE ELECTIVES

DLI electives explore how to apply a specific technology or development technique in two hours. Like full-length courses, electives can be taken anytime, anywhere, with access to GPUs in the cloud.

## INSTRUCTOR-LED WORKSHOPS

In-person workshops teach you how to implement and deploy an end-to-end project through hands-on training in eight hours. Offered at customer sites, conferences, and universities, full-day workshops include hands-on training and lectures delivered by DLI certified instructors.



# NVIDIA RESEARCH



RESEARCH PUBLICATIONS DEMOS NEWS RESEARCH AREAS TEAM ▾ COLLABORATIONS ▾ CAREERS ▾

## RESEARCH AT NVIDIA

Groundbreaking technology begins right here with the world's leading researchers.

## NEWS



AI Can Now Fix Your Grainy Photos by Only Looking at Grainy Photos.

[Read Article >](#)




NVIDIA Brings New Tensor Core GPU AI Tools, Super SloMo, Cutting-Edge Research

[Read Article >](#)



Improving Landmark Localization with a New Deep Learning Architecture

[Read Article >](#)



New AI Imaging Technique Reconstructs Photos with Realistic Results

[Read Article >](#)

**THIS IS ONLY THE BEGINNING...**

