SAND2017-12551PE

Deep Learning: Neural-Inspired AI

Tim Draelos

November 15, 2017

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International Inc. for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.


A B C D E F G H I J K L M N O P Q R S
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
ALPHAGO
01:27:15
LEE SEDOL
00:45:18

INTRODUCING

amazon echo

Always ready, connected,
and fast. **Just ask.**


Tesla Motors / MGN


MediaMarkt


UN Deputy Chief Interviews Social Robot Sophia
THE FUTURE OF
EVERYTHING


FUTURE
INVESTMENT
INITIATIVE
FUTURE
INVESTMENT


METTLE
& METAL
THE DARPA ROBOTICS CHALLENGE


Tara


TU/e

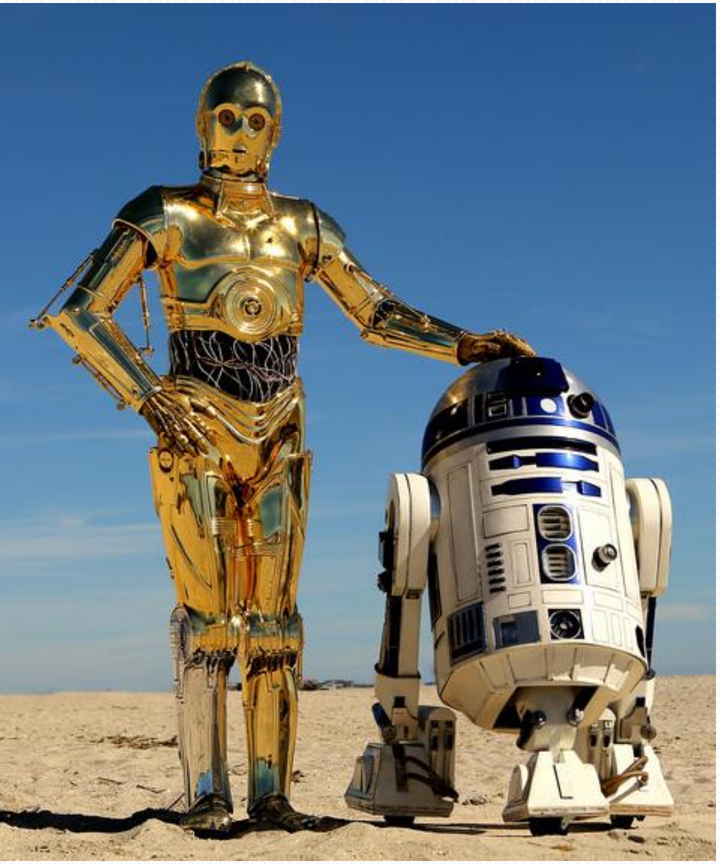

Deep Learning

- Representation Learning to Discover/Detect Features
  - Hierarchical Feature Selection / Feature Extraction
    - Features = Explanatory factors, attributes of data samples
  - Good representation (feature set) captures the *a posteriori* distribution of underlying explanatory factors (e.g., Classes) of an observed environment.
    - P(Class | Data) ➔ P(Class | Features)
    - Each data sample is a biased representation of its class

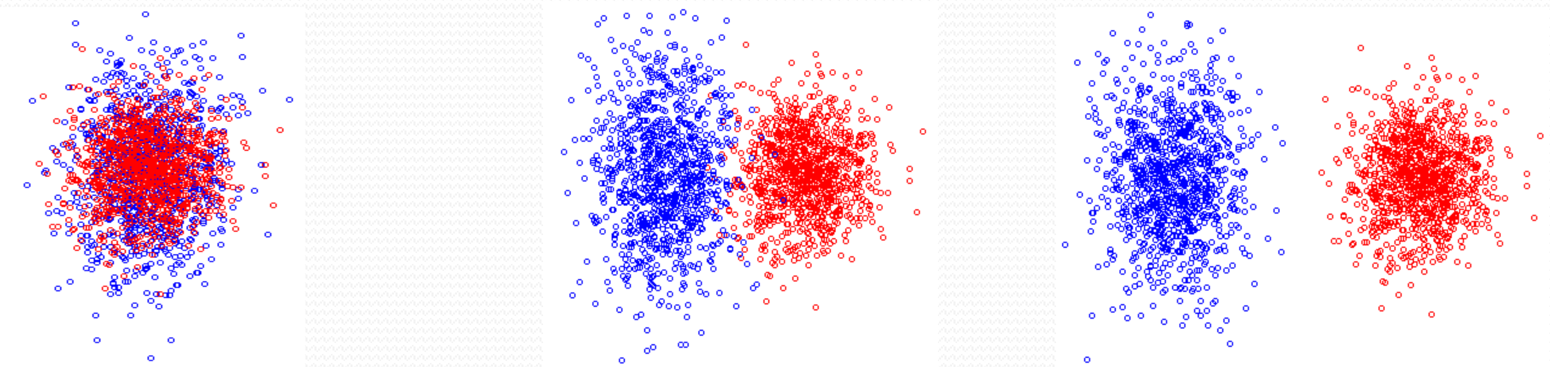

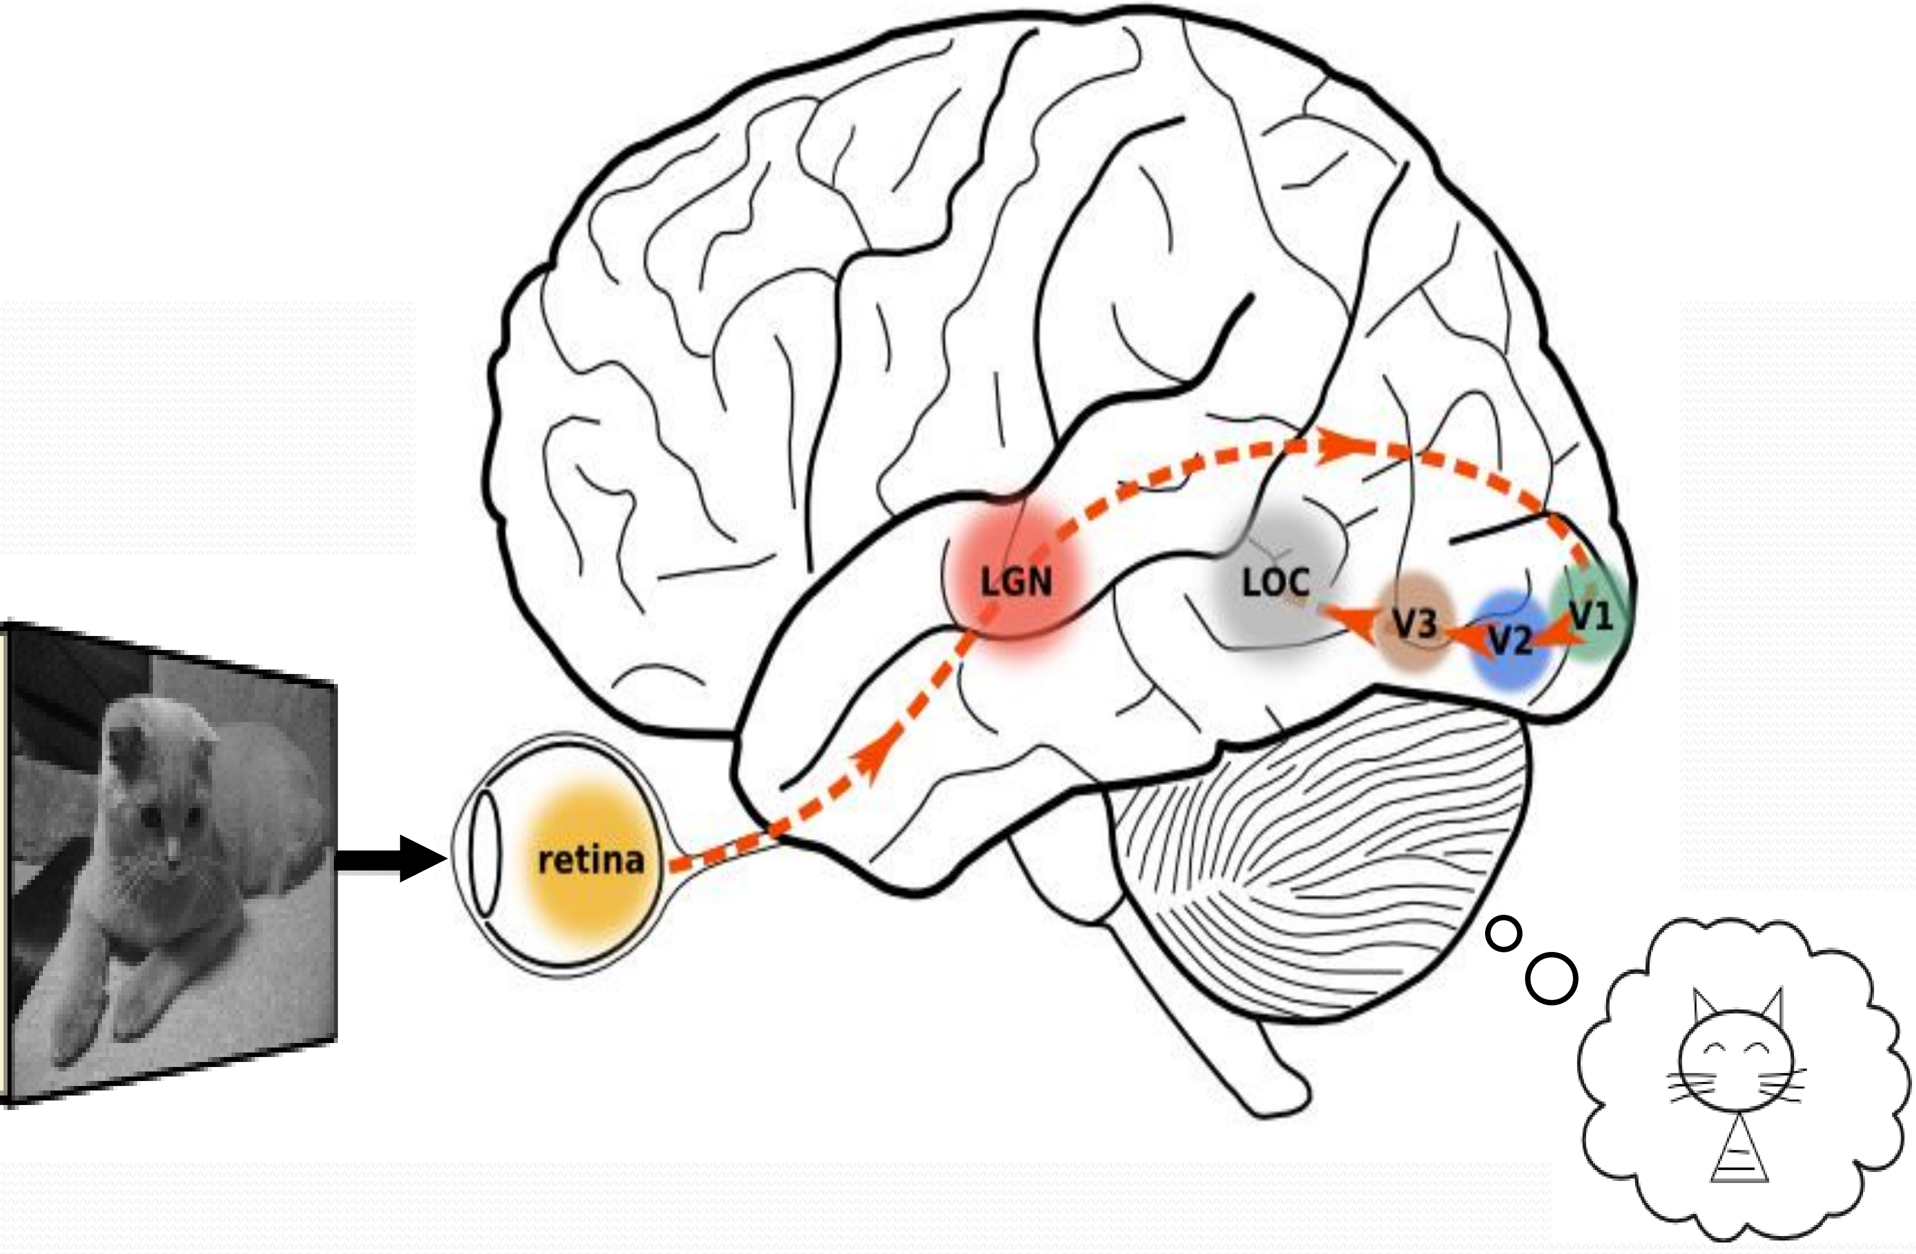
LGN
LOC
V3
V2
V1
retina

Hierarchical Image Processing in the Brain

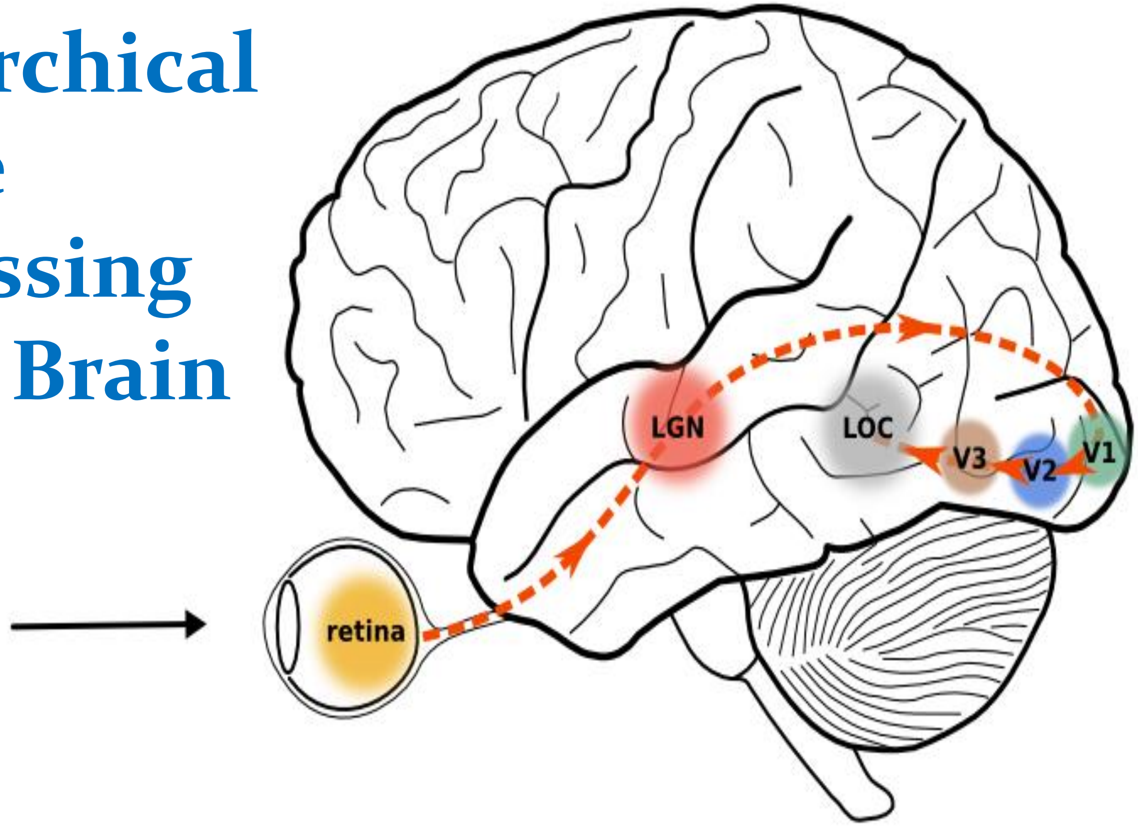

Deep Learning: Neural-Inspired AI

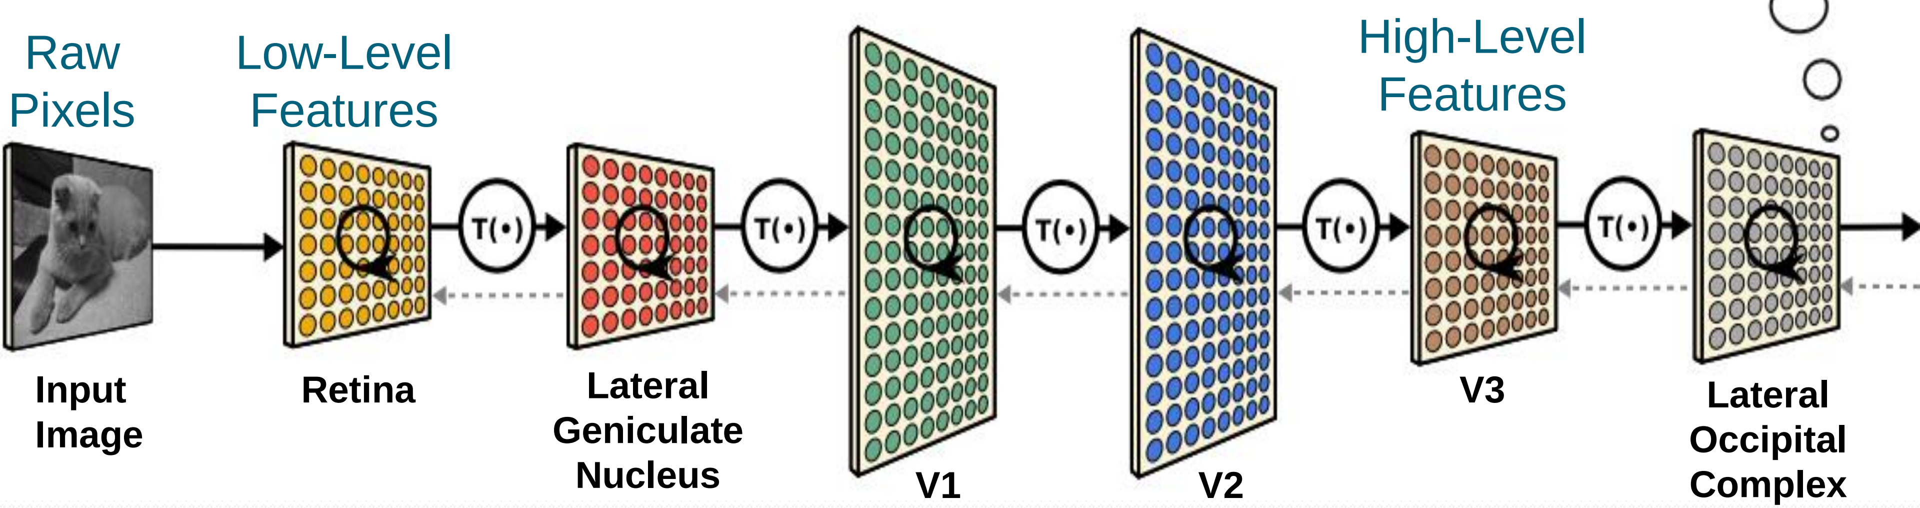

Fully-Connected Feed-Forward Artificial Neural Network

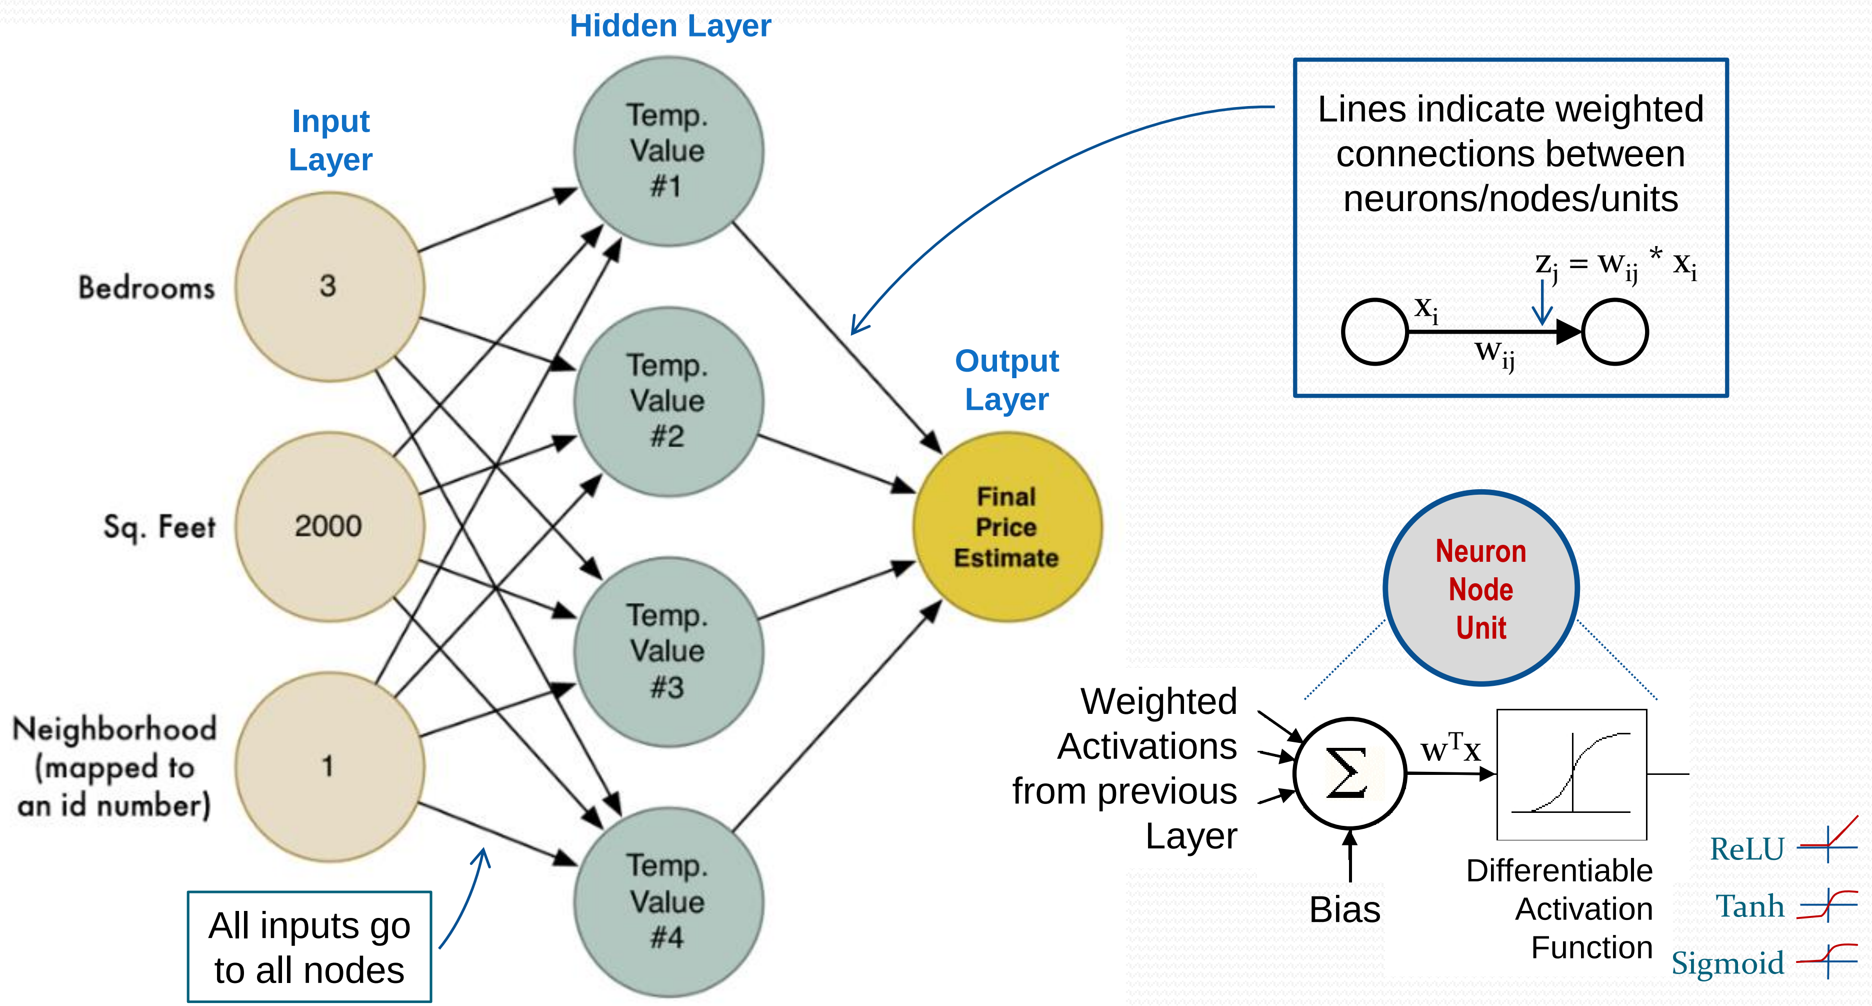

Fully-Connected Feed-Forward Artificial Neural Network

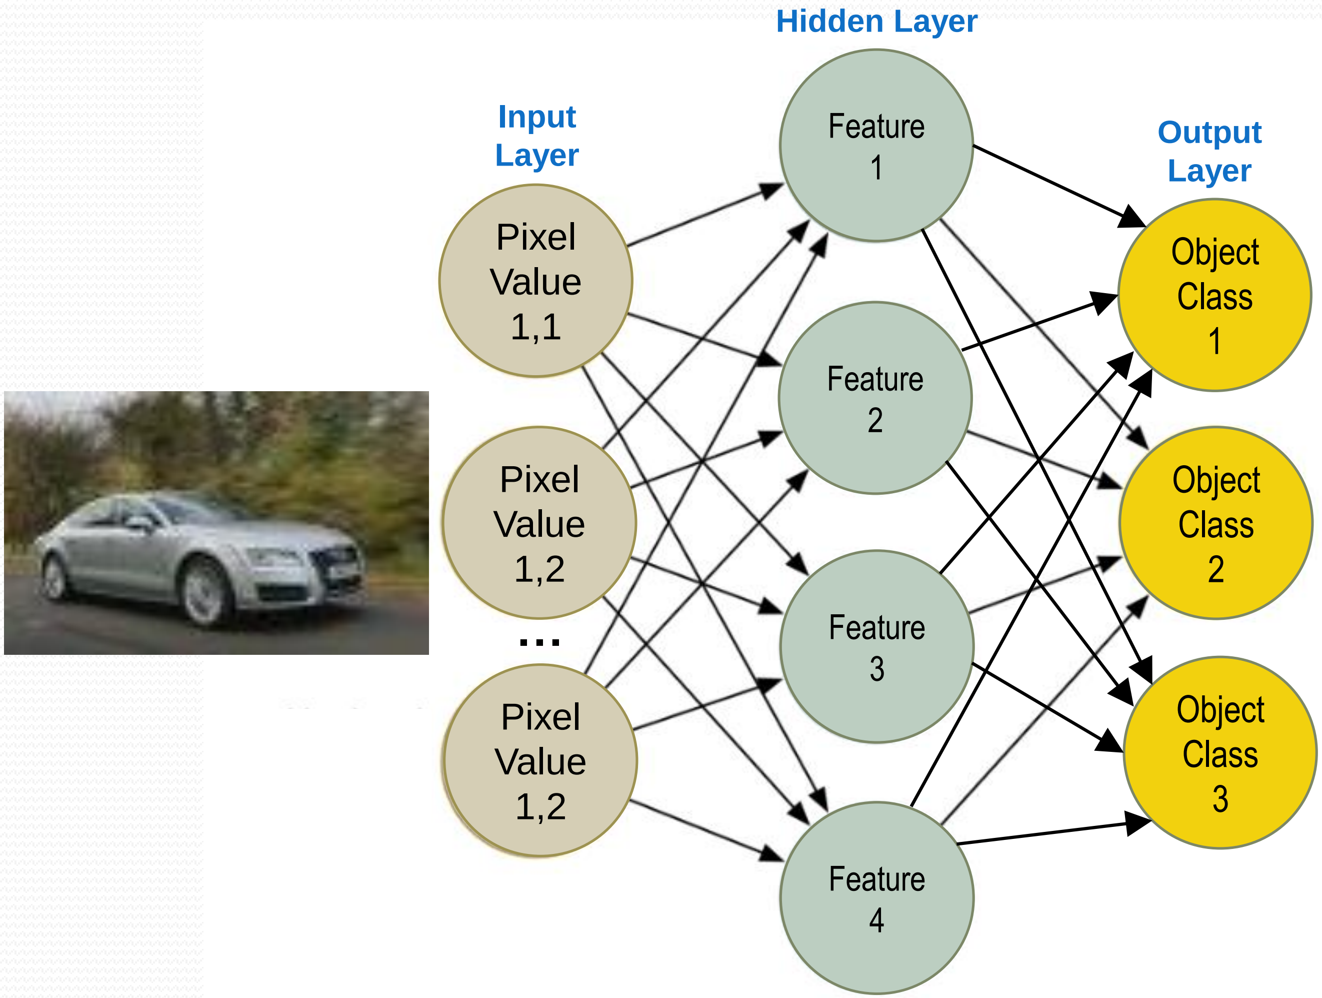

Convolutional Neural Network (CNN)

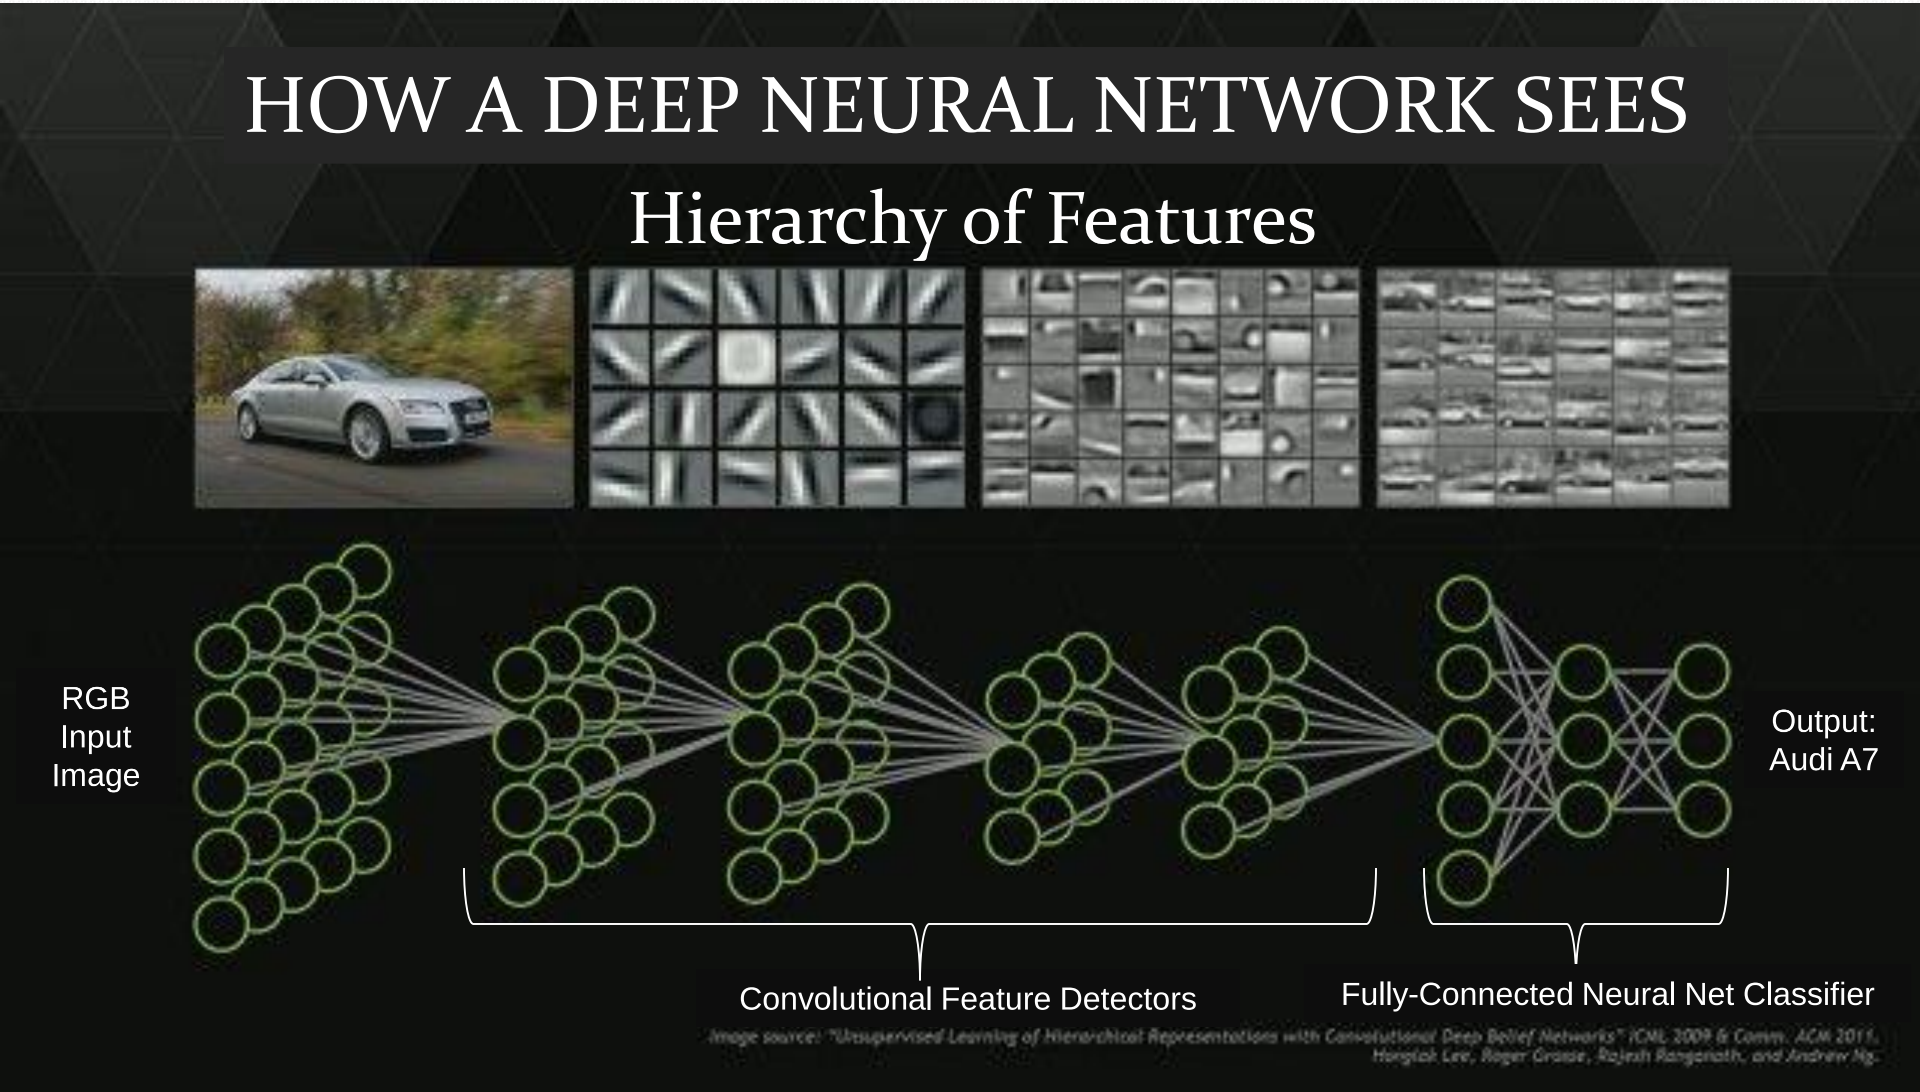

CNN Operation

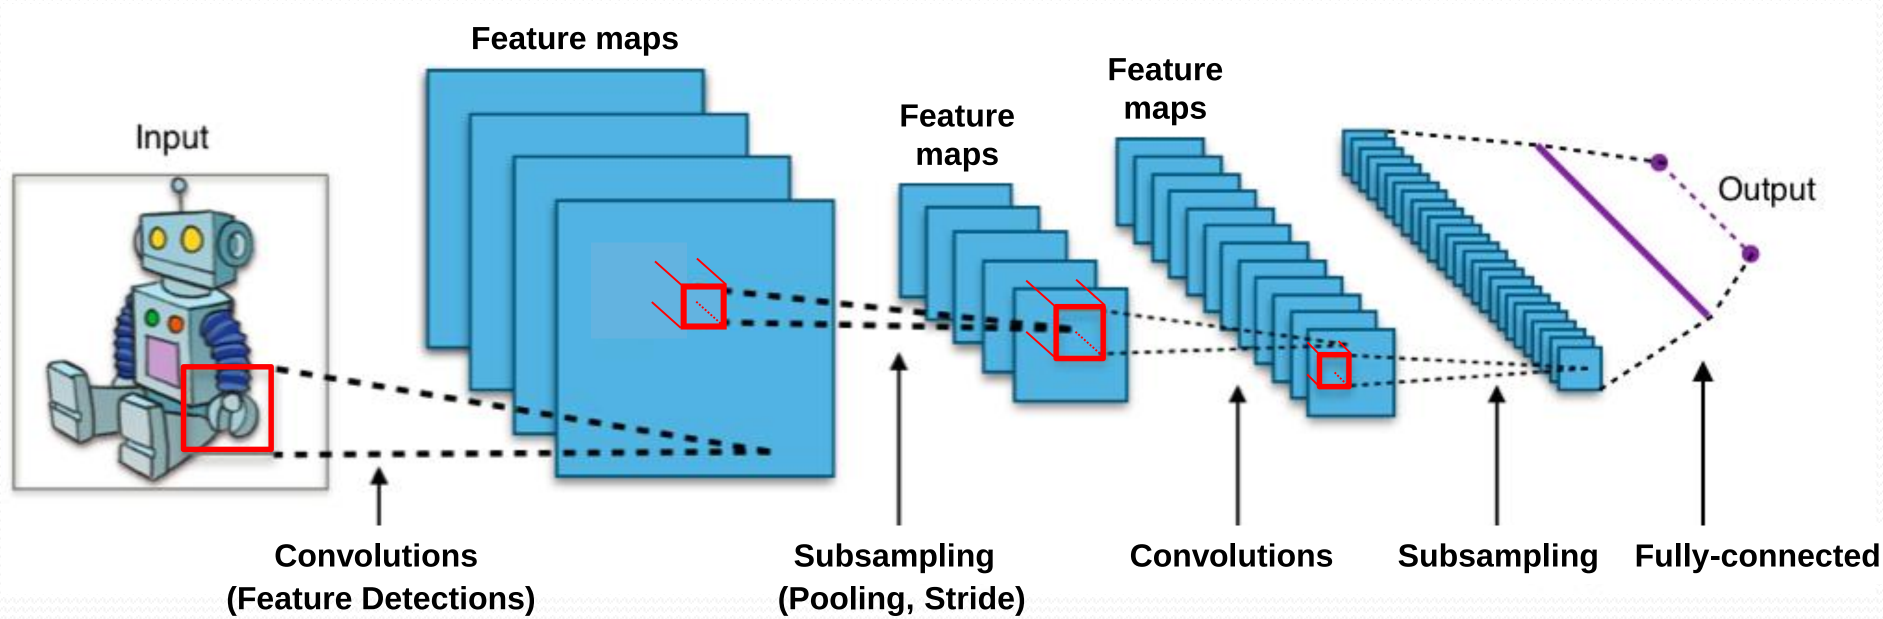

- Capture (Training) or Detect (Testing) spatial structure (features)
  - Convolution is used to find features in signals (template matching)

Let f be the signal and g be a feature template/filter/kernel

$$(f^*g)(t) = \int_{-\infty}^{\infty} f(\tau)g(t-\tau)d\tau$$

$$(f^*g)(t) = \int_{m=-M}^{M} f(m)g(t-m)$$

CNN Training

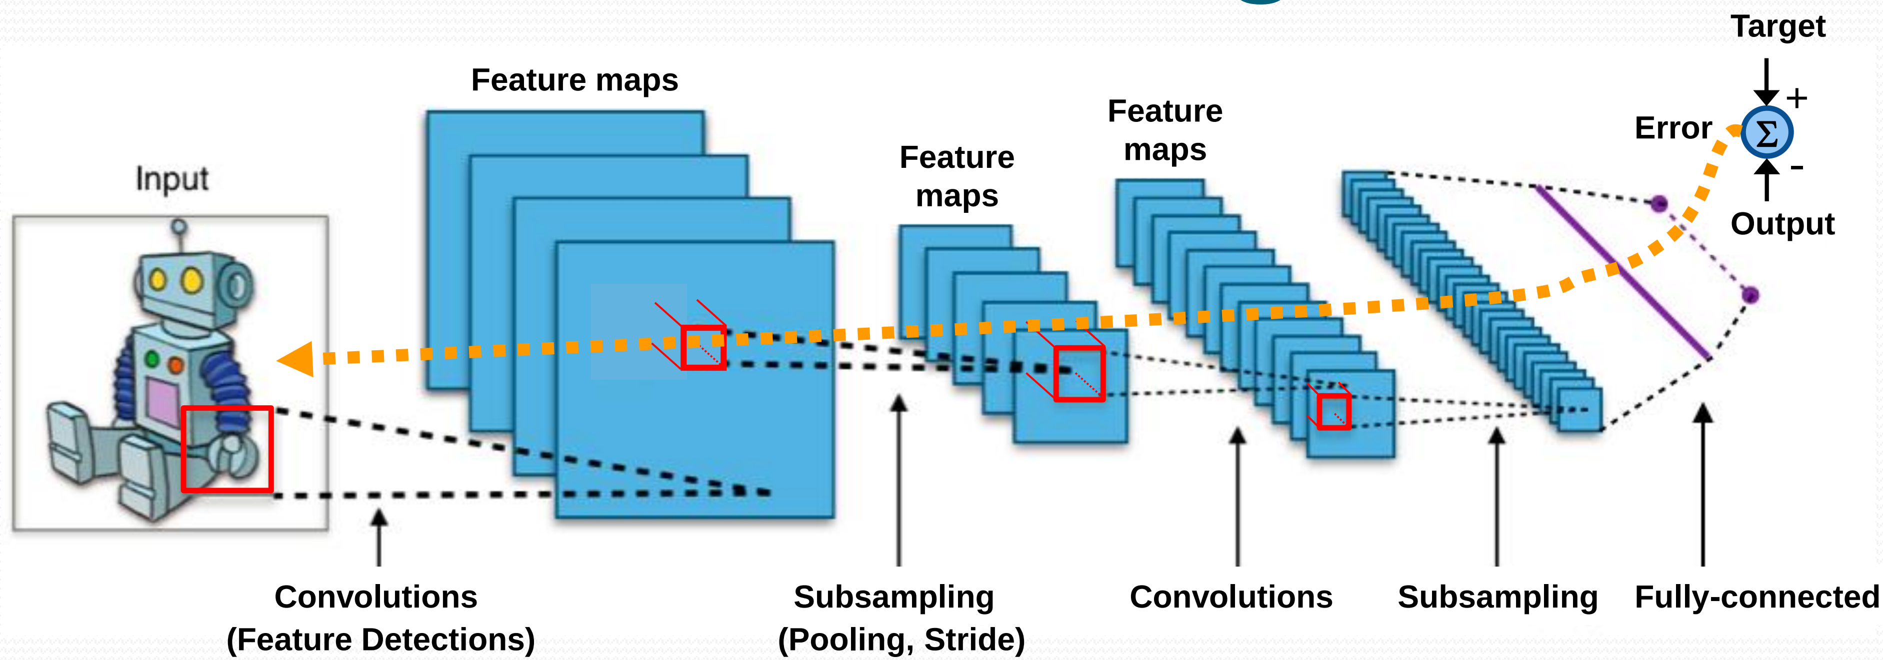

- Convolution filters are the weights in a CNN that can be trained.
- Filter values (weights) are initialized with random values and update via back-propagation.

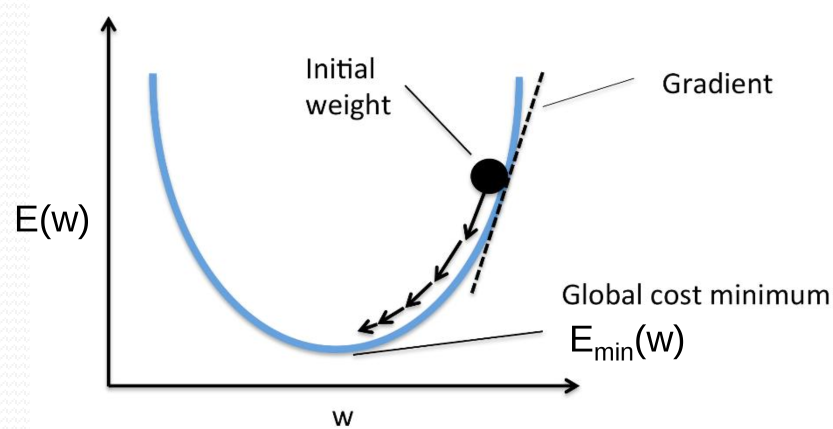

CNN Training

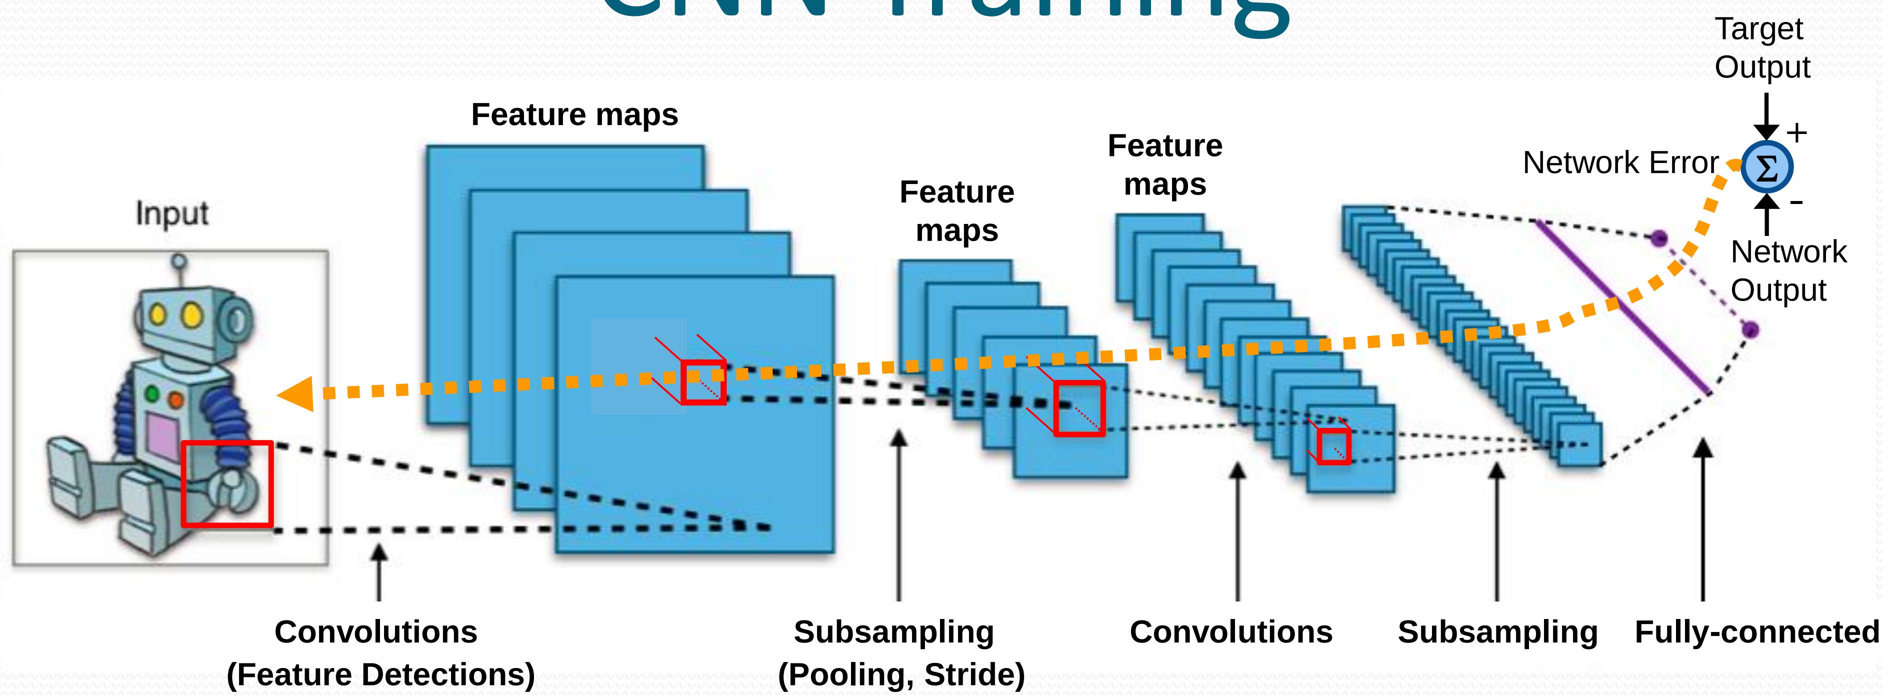

- Convolution filters are the weights in a CNN that can be trained.
- Filter values (weights) are initialized with random values and update via back-propagation.

<u>Weight update during Training</u>

$$w(t+1) = w(t) - \eta \Delta w(t)$$

$$\Delta w = \frac{\partial E}{\partial w}$$

Finding Features in an Image

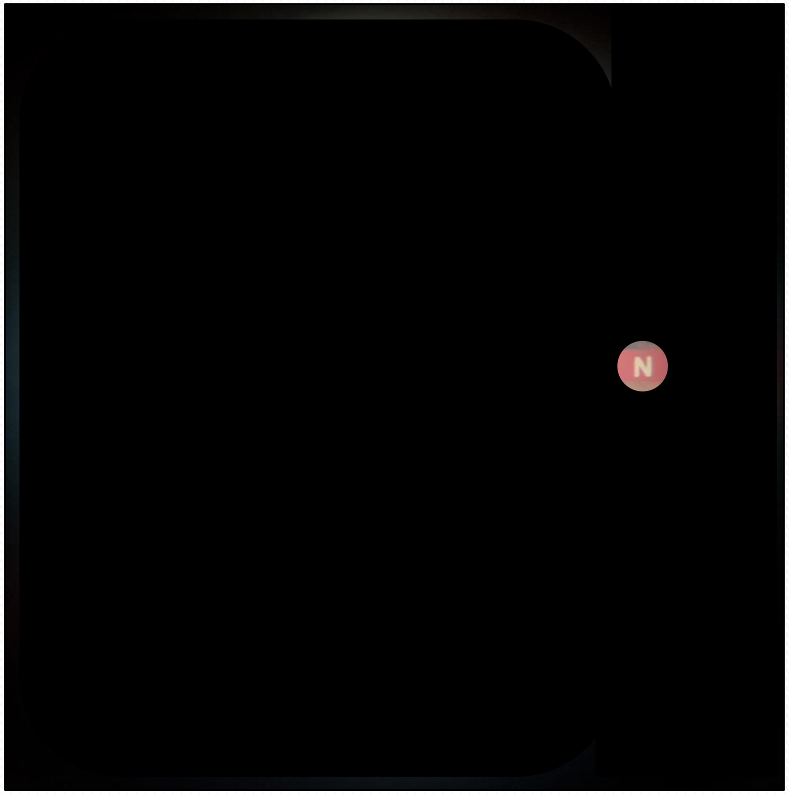

Finding Features in an Image

Finding Features in an Image

Finding Features via Convolution

Convolution Produces Feature Maps

Stride

Output of Convolution will be a
7 x 9 “Pixel” Feature Map

Unsupervised Learning

Reinforcement Learning

Supervised Learning

Transfer Learning

Neural Nets – What’s Changed?

Neural Nets – What's Changed?

- Architecture – NOT REALLY

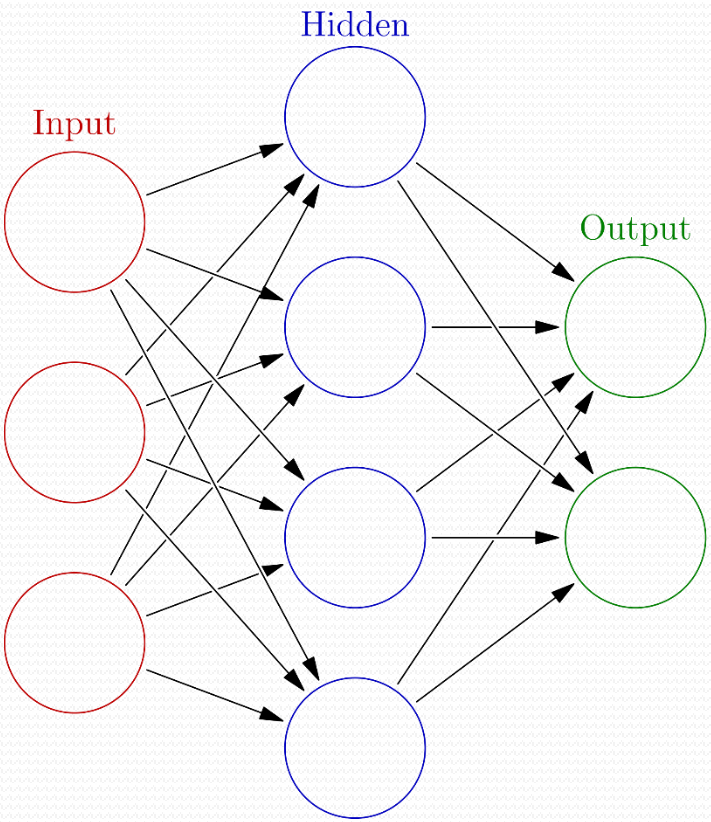

Feed-Forward Neural Network

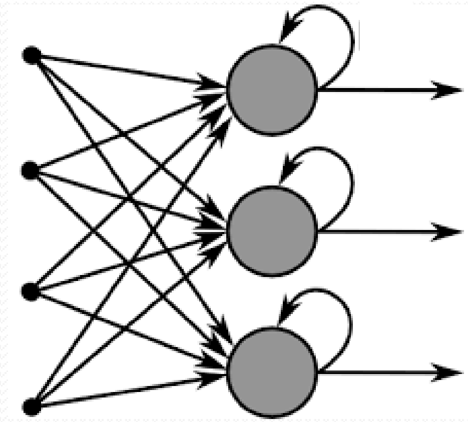

Recurrent Neural Network

Long Short-Term Memory Cell

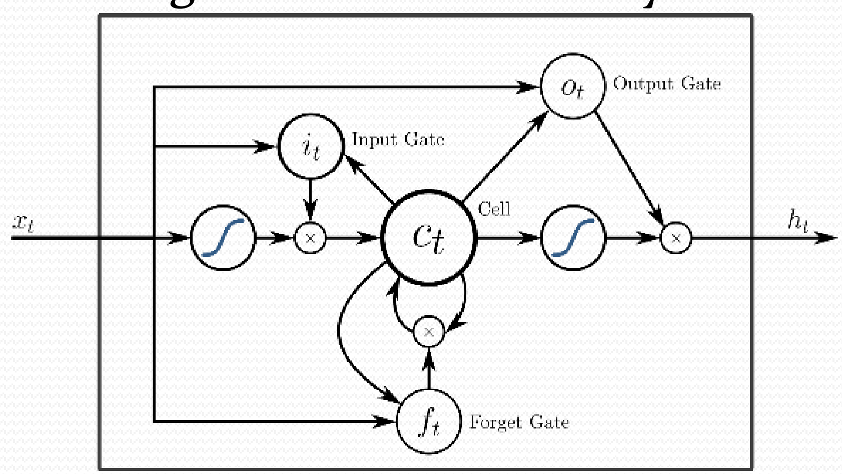

Neural Nets – What's Changed?

- Architecture – NOT REALLY
- Learning algorithms – NOT REALLY since BACKPROP
  - Better Optimizers, Adaptive learning rates, Minibatches

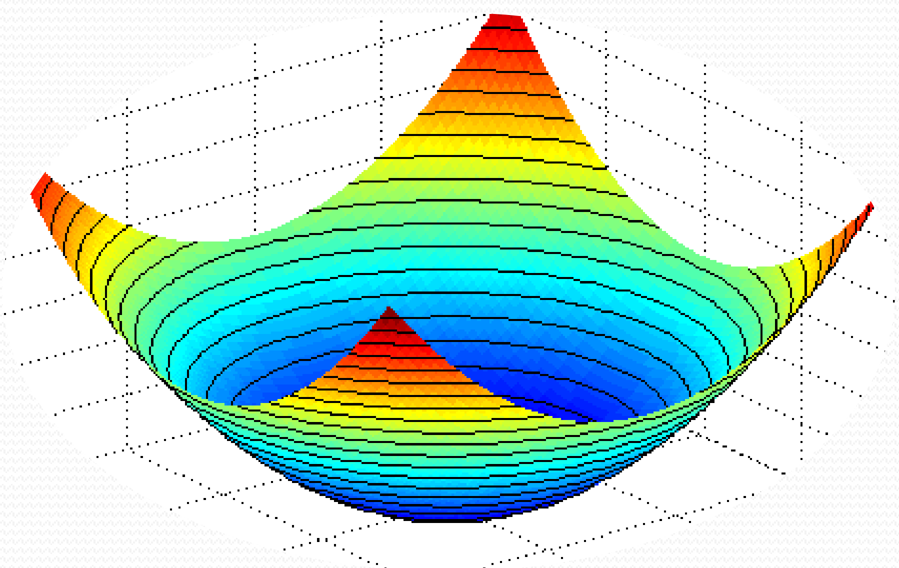

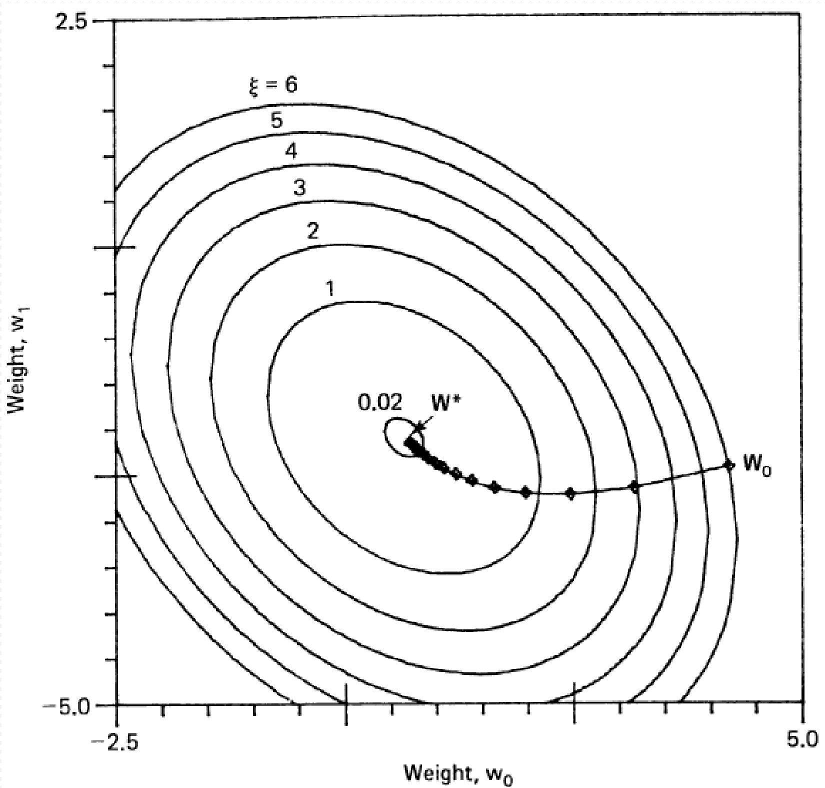

Neural Nets – What's Changed?

- Architecture – NOT REALLY
- Learning algorithms – NOT REALLY since BACKPROP
  - Better Optimizers, Adaptive learning rates, Minibatches
- Regularization techniques – YES
  - Dropout, Data augmentation, Ensemble methods

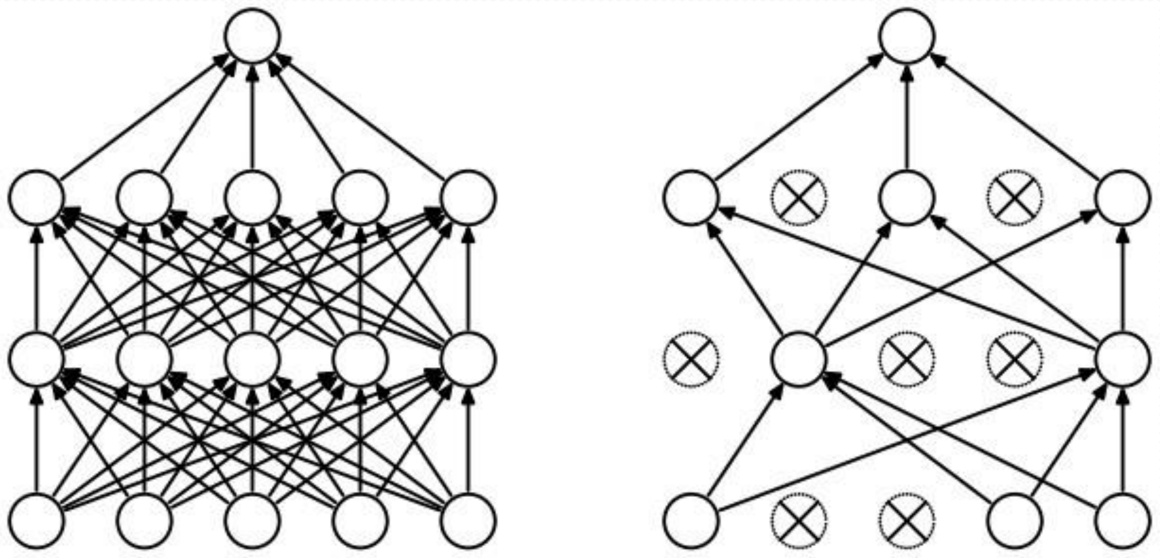

(a) Standard Neural Net (b) After applying dropout.

Neural Nets – What's Changed?

- Architecture – NOT REALLY
- Learning algorithms – NOT REALLY since BACKPROP
  - Better Optimizers, Adaptive learning rates, Minibatches
- Regularization techniques – YES
  - Dropout, Data augmentation, Ensemble methods
- Activation function – YES
  - Rectified Linear Unit (ReLU)

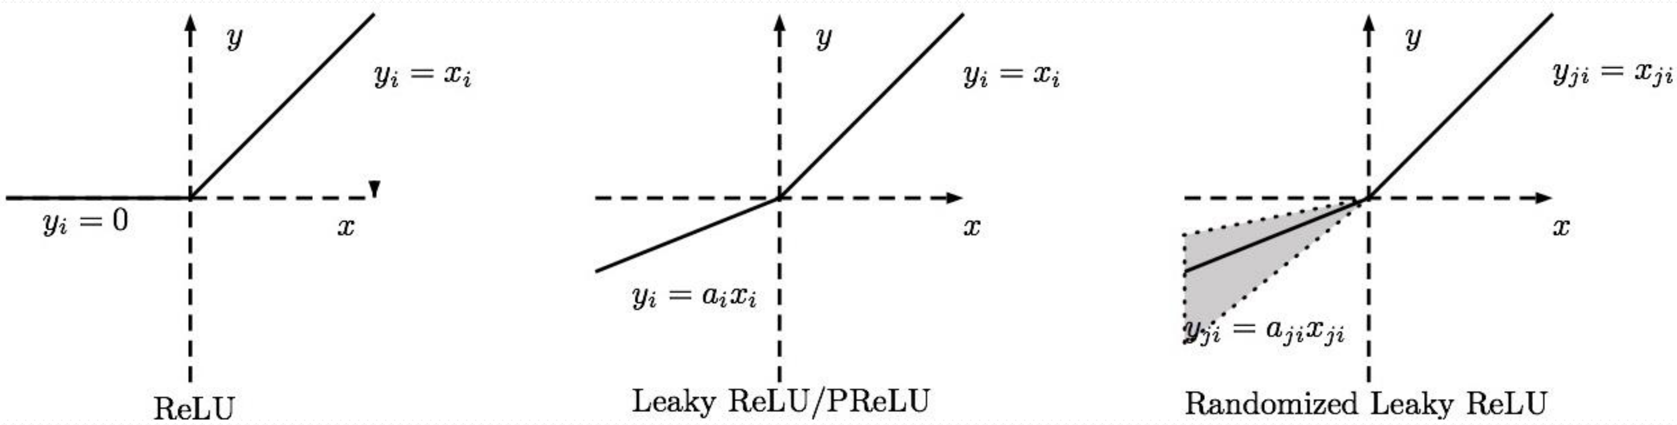

Neural Nets – What's Changed?

- Architecture – NOT REALLY
- Learning algorithms – NOT REALLY since BACKPROP
  - Better Optimizers, Adaptive learning rates, Minibatches
- Regularization techniques – YES
  - Dropout, Data augmentation, Ensemble methods
- Activation function – YES
  - Rectified Linear Unit (ReLU)
- Labelled data – YES

Music Datasets
- Piano-midi.de
- Nottingham
- MuseData
- JSB
- FMA

Natural Images
- MNIST
- NIST
- Perturbed NIST
- CIFAR10 / CIFAR100
- Caltech 101
- Caltech 256
- Caltech Silhouettes
- STL-10
- The Street View House Numbers (SVHN)
- NORB:
- Imagenet
- Pascal VOC
- Labelme
- COIL 20
- COIL100

Artificial Datasets
- Arcade Universe
- BabyAISchool:
  - BabyAIShapesDatasets
  - BabyAIImageAndQuestion
- DeepVsShallowComparisonICML2007)
  - MnistVariations
  - RectanglesData
  - ConvexNonConvex
  - BackgroundCorrelation

Faces
- Labelled Faces in the Wild
- Toronto Face
- Olivetti
- Multi-Pie
- Face-in-Action
- JACFEE
- FERET
- IndianFaceDatabase:
- The Yale Face Database

Text
- 20 newsgroups
- Reuters (RCV*)
- Penn Treebank
- Broadcast News
- Wikipedia Dataset
- Multidomain sentiment analysis

Speech
- TIMIT Speech Corpus
- Aurora

Recommendation Systems
- MovieLens
- Jester
- Netflix Prize
- Book-Crossing dataset

Misc
- “Musk”
- CMU Motion Capture Database
- Brodatz
- Million Song
- Merck Molecular Activity Challenge

Neural Nets – What's Changed?

- Architecture – NOT REALLY
- Learning algorithms – NOT REALLY since BACKPROP
  - Better Optimizers, Adaptive learning rates, Minibatches
- Regularization techniques – YES
  - Dropout, Data augmentation, Ensemble methods
- Activation function – YES
  - Rectified Linear Unit (ReLU)
- Labelled data – YES
- Compute power – YES
  - GPUs

Neural Nets – What's Changed?

- Architecture – NOT REALLY
- Learning algorithms – NOT REALLY since BACKPROP
  - Better Optimizers, Adaptive learning rates, Minibatches
- Regularization techniques – YES
  - Dropout, Data augmentation, Ensemble methods
- Activation function – YES
  - Rectified Linear Unit (ReLU)
- Labelled data – YES
- Compute power – YES
  - GPUs
- Software – YES

Sandia Work – Neurogenesis Deep Learning

PROBLEM: Existing Autoencoder (AE) can't represent novel data

SOLUTION: Add new node(s) to each level of an AE as needed to represent novel data

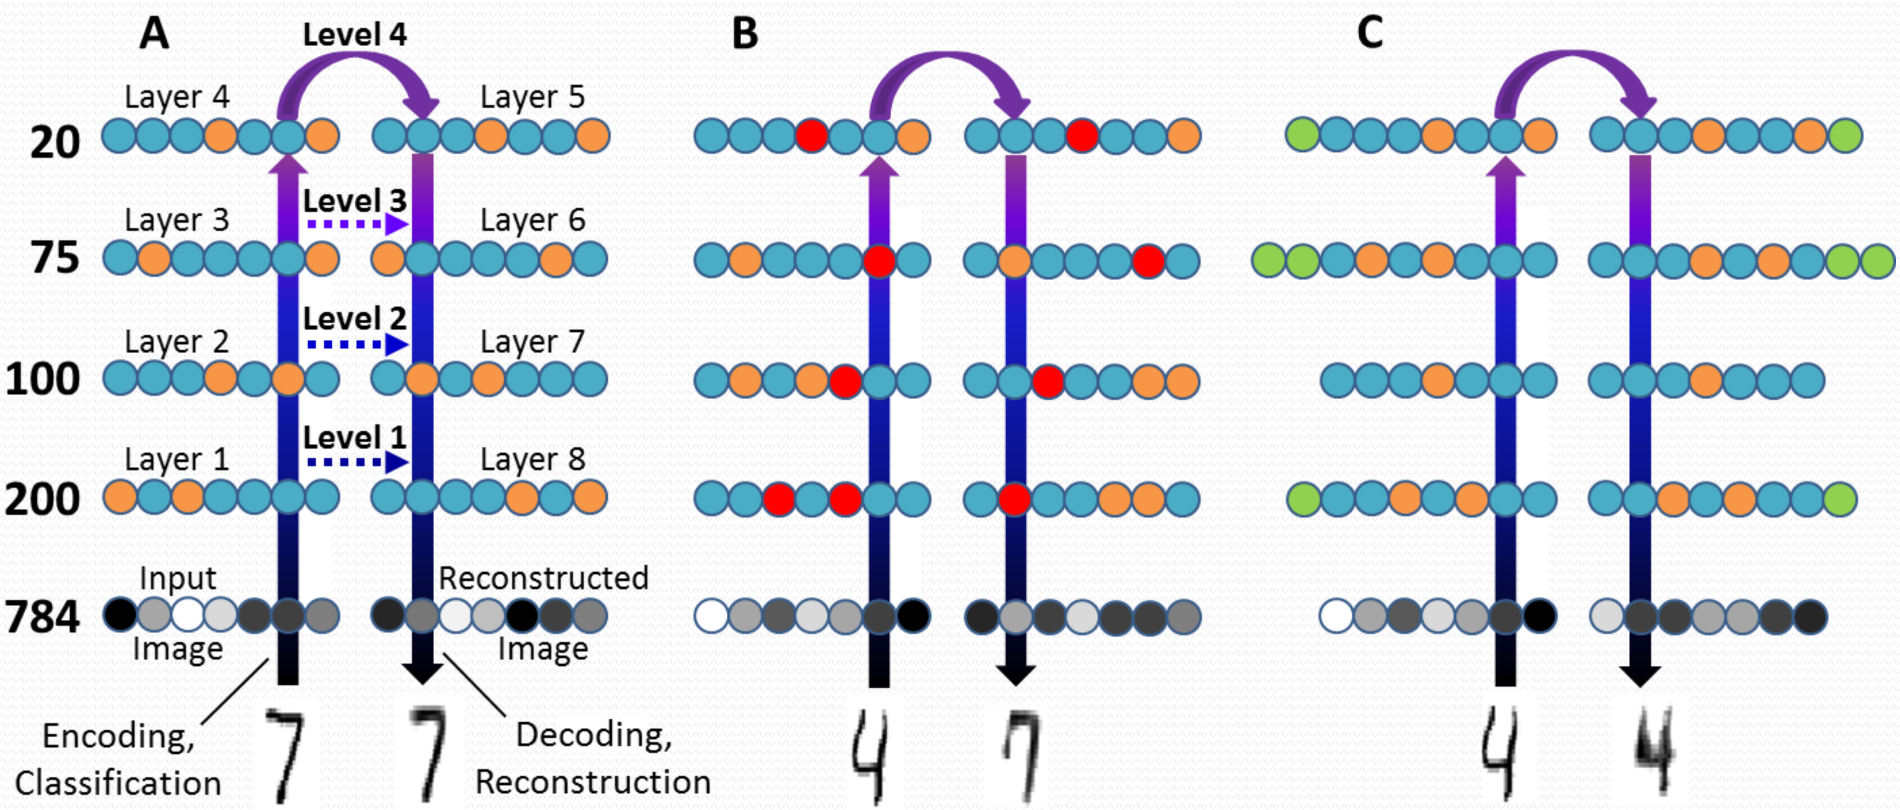

Sandia Work – Seismic Event Discrimination

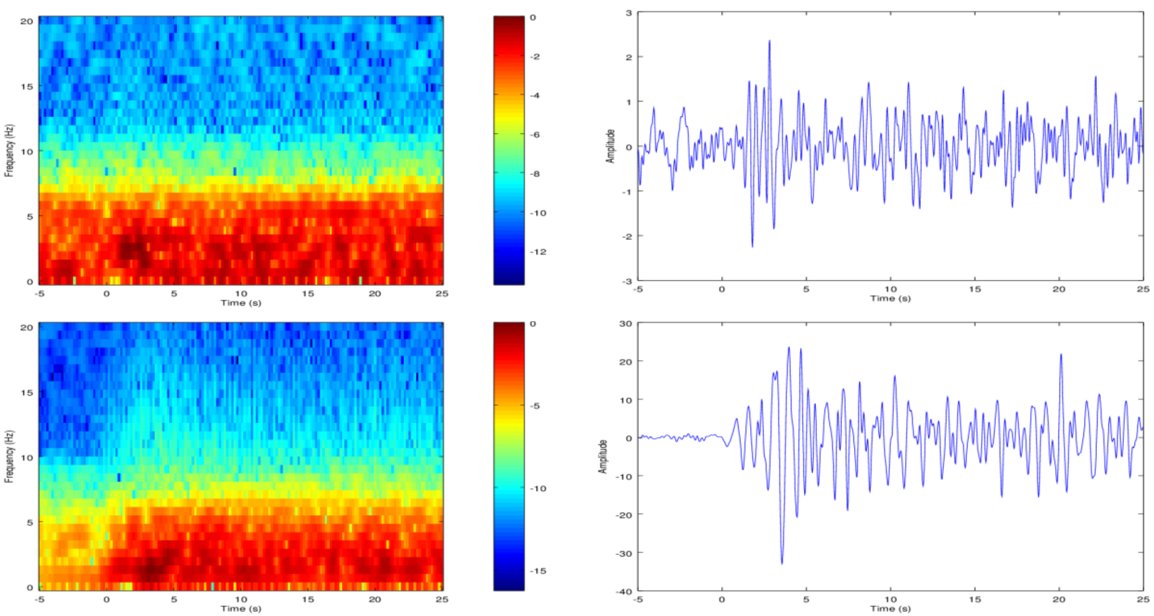

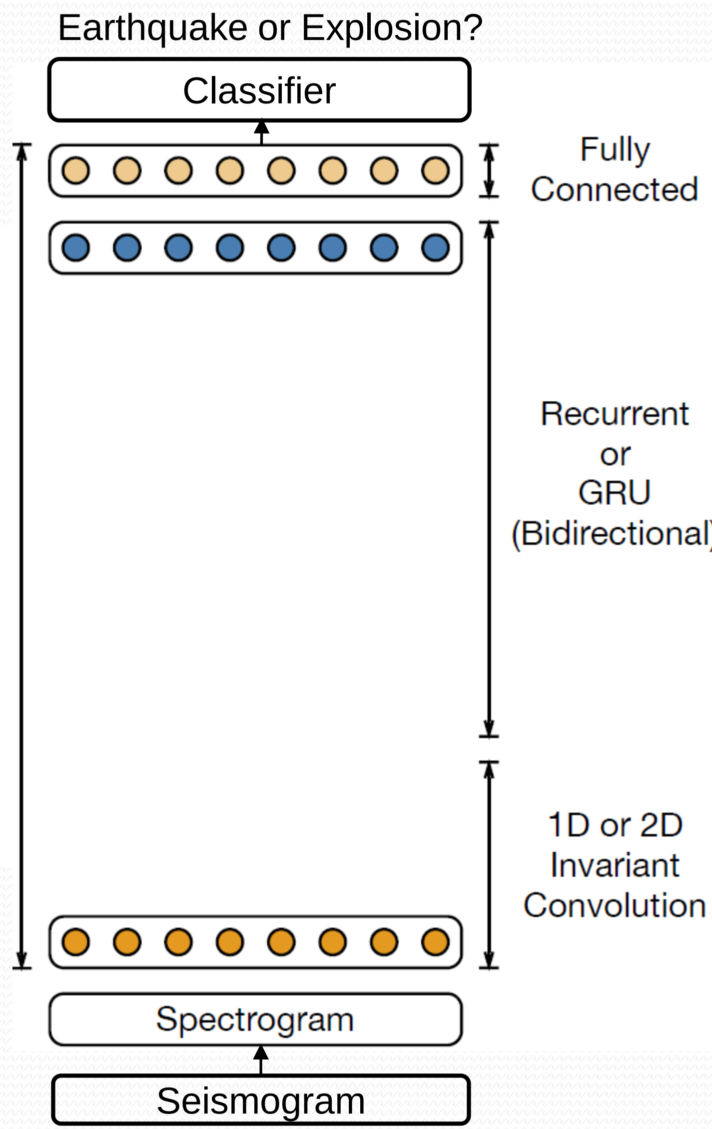

Sandia Work – Coordinated ARTMAP

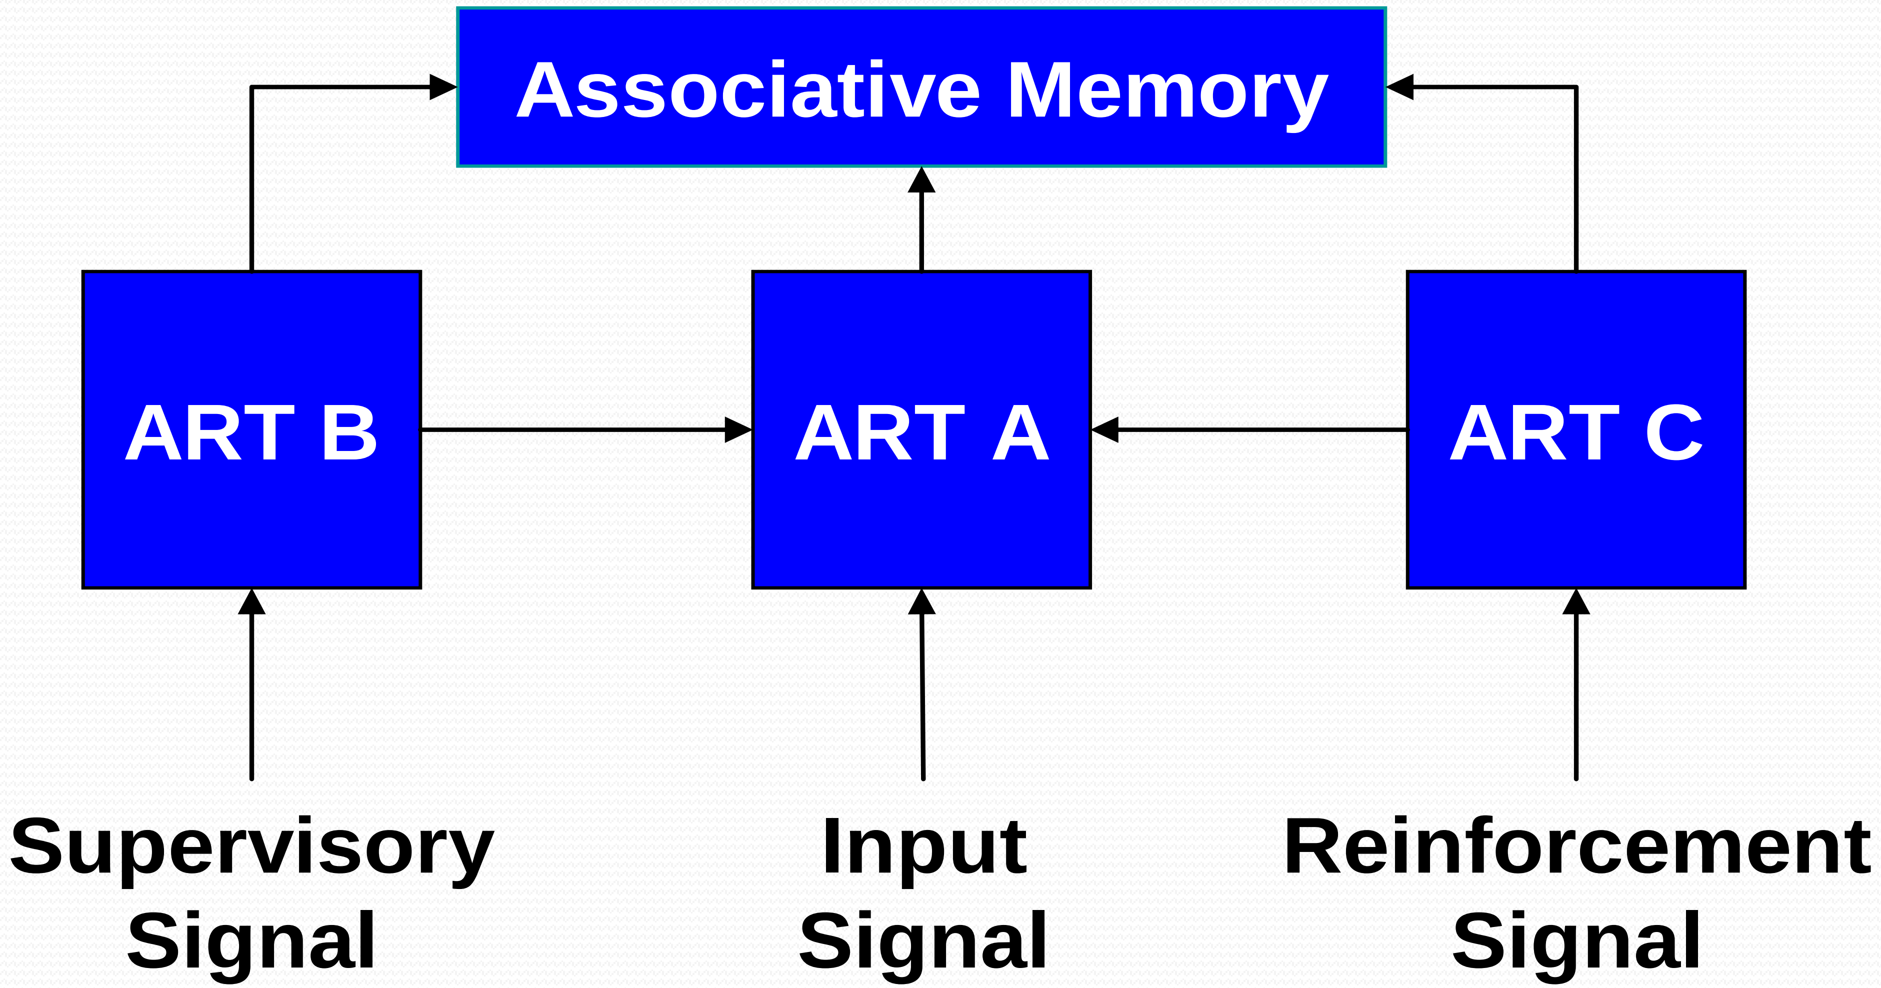

Sandia Work - Fine-grained Image Recognition using contour data

Sandia Work

- Secure DL
  - Does compressing a network help with universal adversarial manipulation?

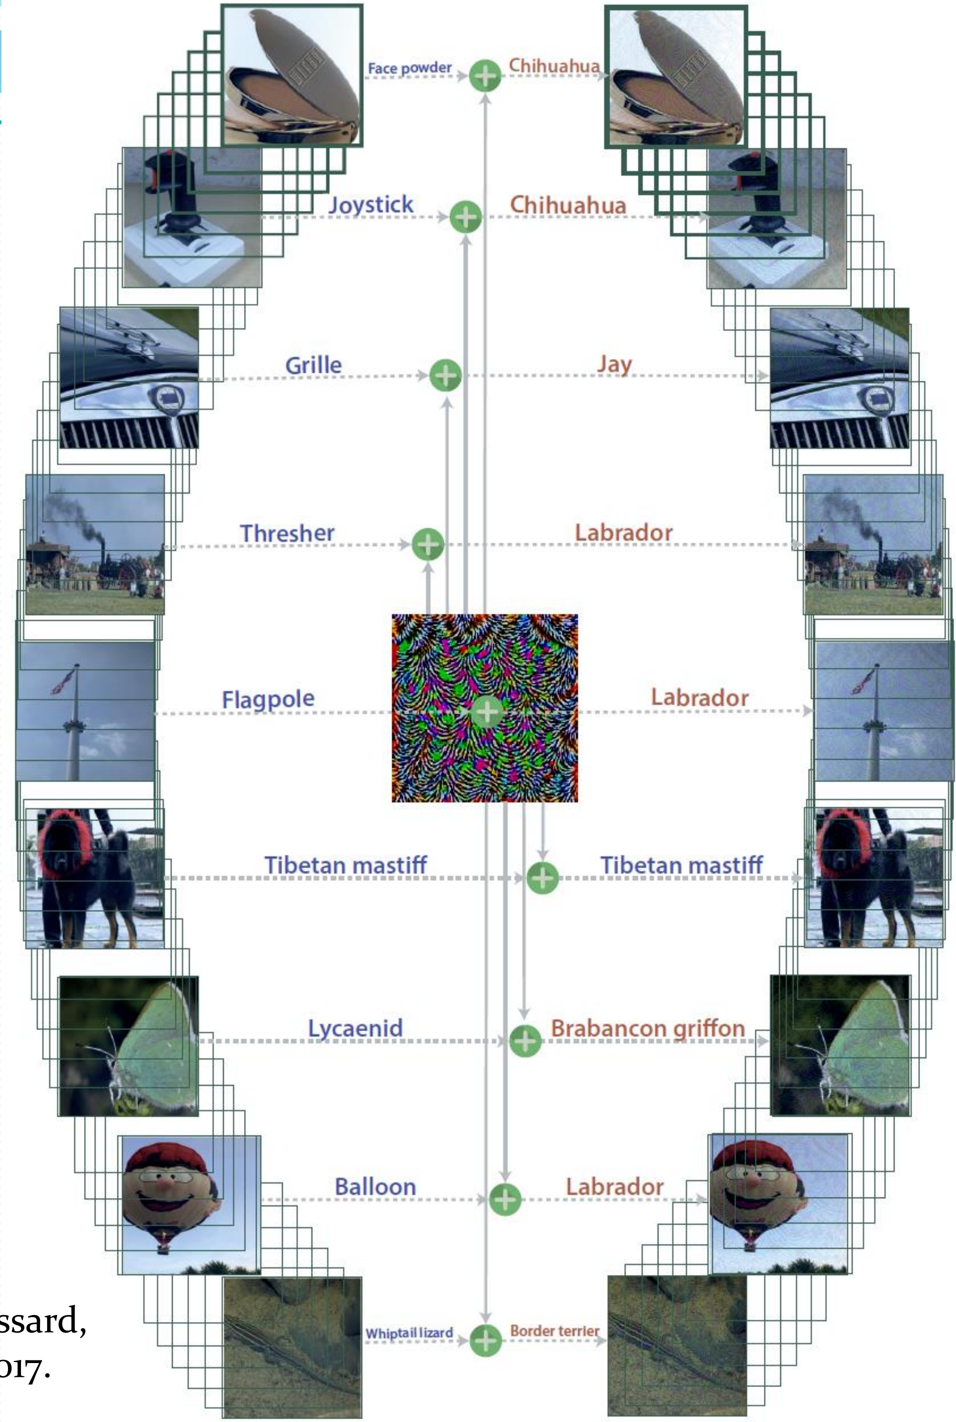

S.-M. Moosavi-Dezfooli, A. Fawzi, O. Fawzi, and P. Frossard, "Universal adversarial perturbations," in IEEE CVPR, 2017.

Sandia Work

- Object Tracking with CNN
  - Real Sandia Peak data augmented with simulated data

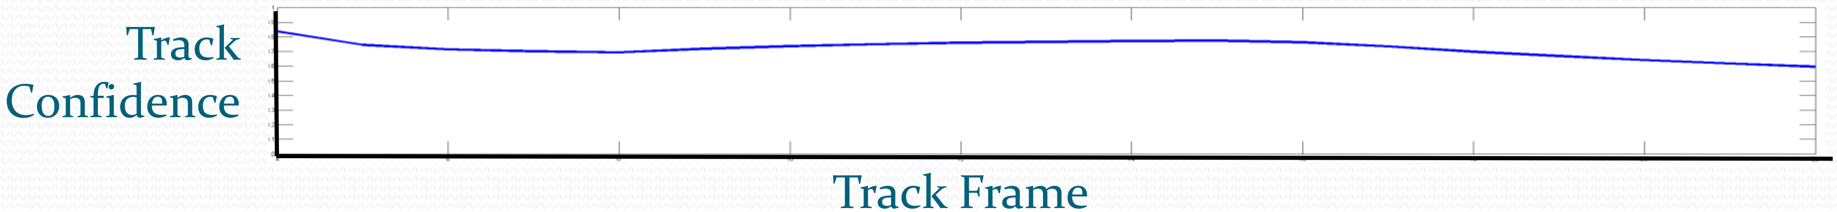

- Denoising

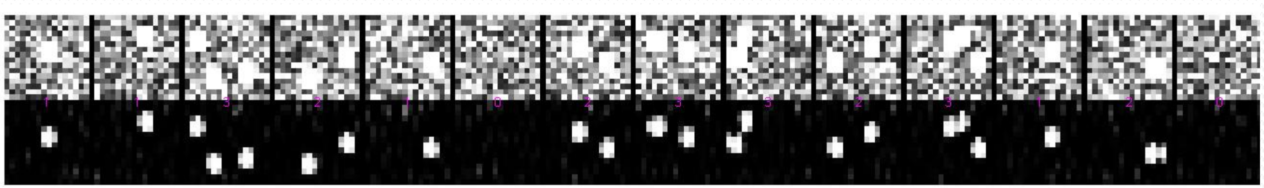

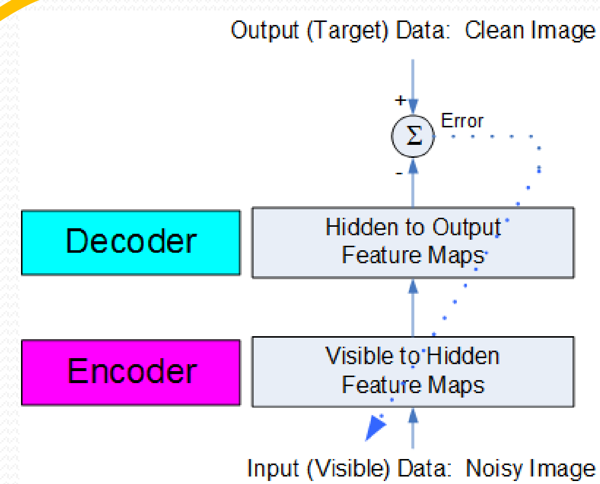

Sandia Work – Cyber Security

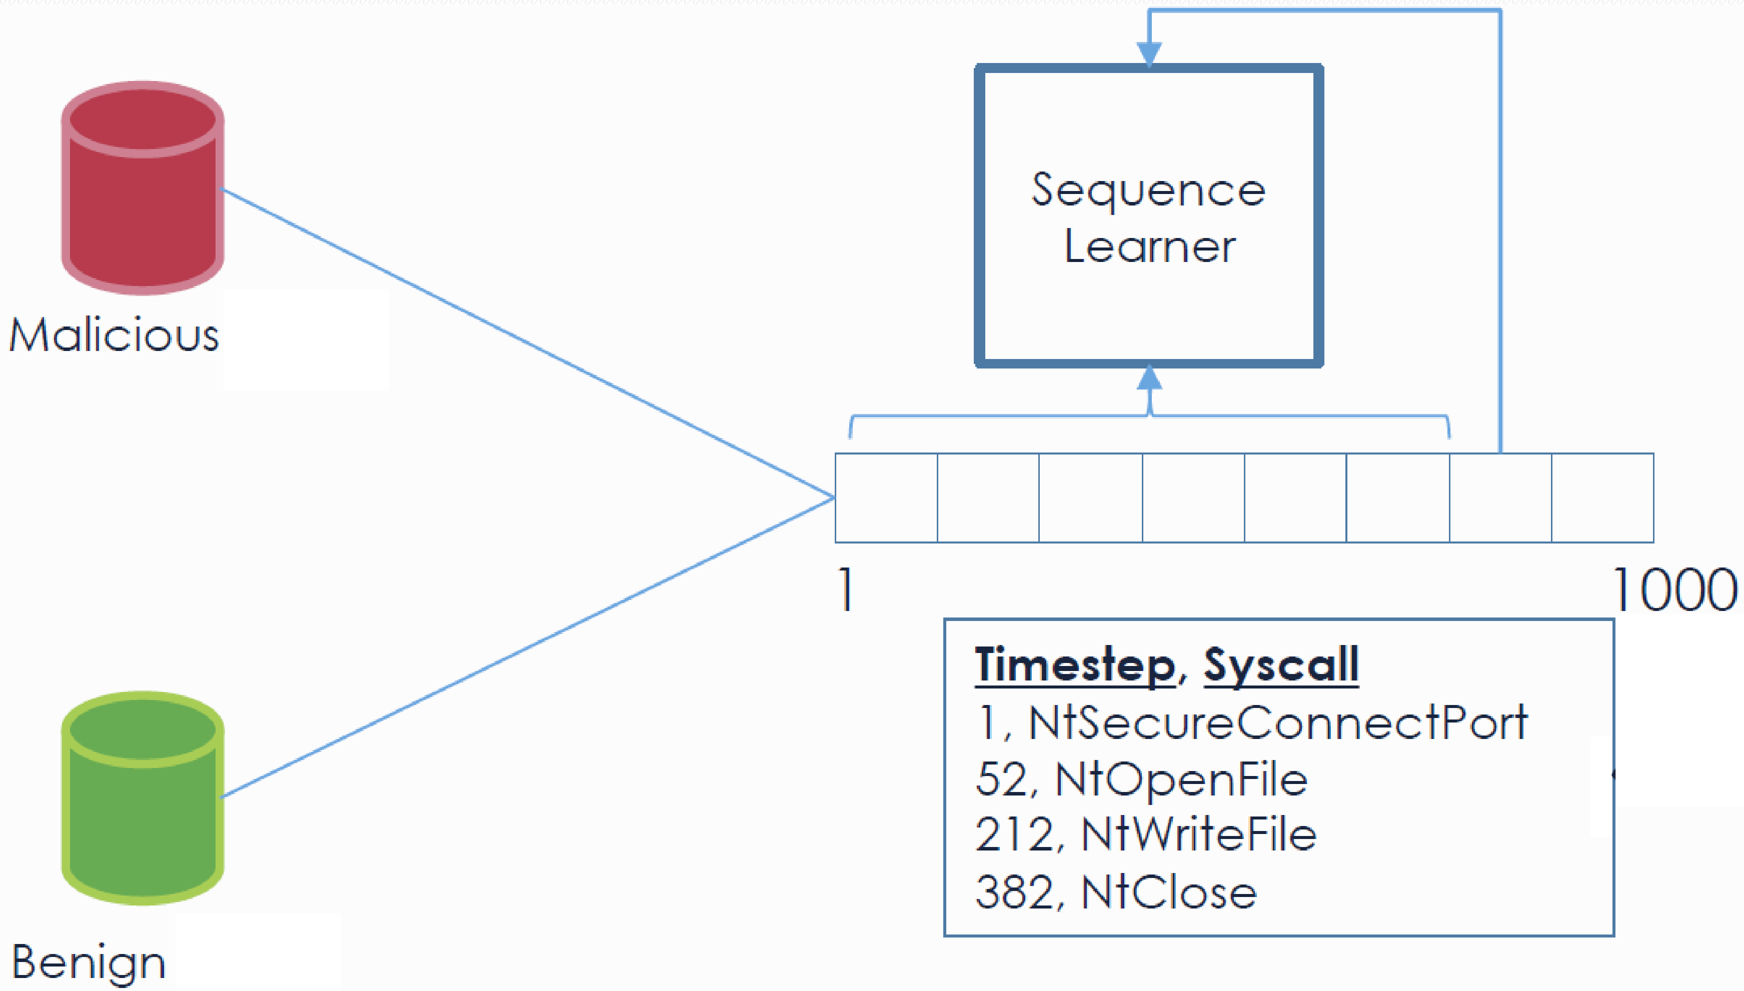

Sandia Work – Paraphrase Detection

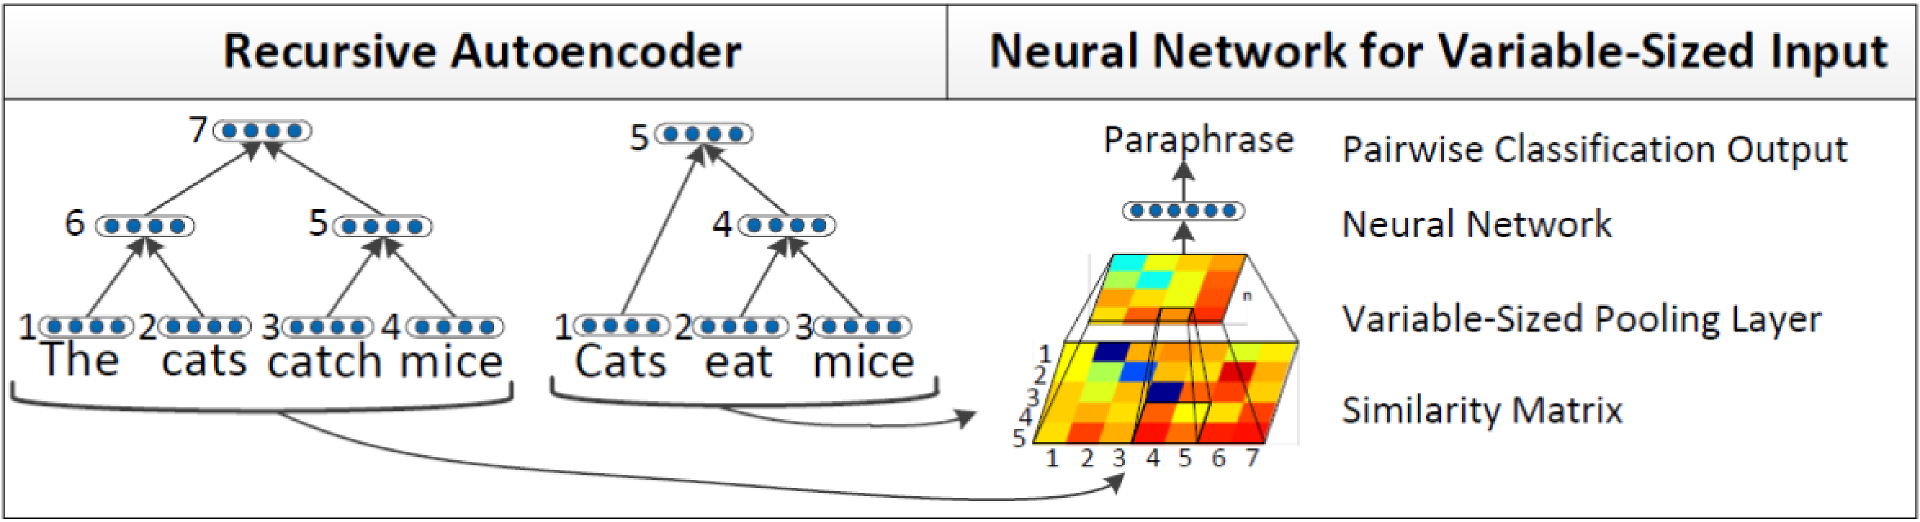

R. Socher, E. Huang, J. Pennington, A. Ng, and C. Manning. Dynamic Pooling and Unfolding Recursive Autoencoders for Paraphrase Detection, NIPS 2011.

Opportunities at Sandia

- Sandia Research Foundations
  - Bioscience
  - Computing and Information Science
  - Engineering Science
  - Geoscience
  - Materials Science
  - Nanodevices and Microsystems
  - Radiation Effects and High Energy Density Science

- National Security Issues
  - Trust
  - Explainability

- Problems of National Interest
  - Cybersecurity
  - Energy

Simulate at high level of neural fidelity

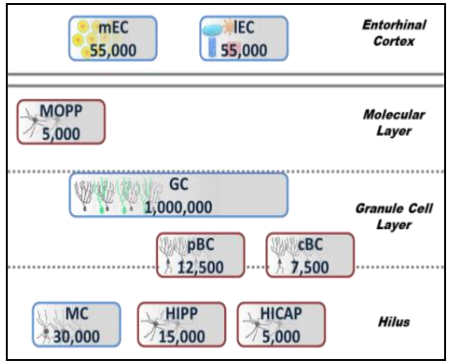

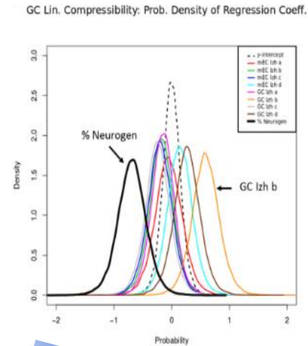

Formalize hippocampus / neurogenesis sparse coding algorithm

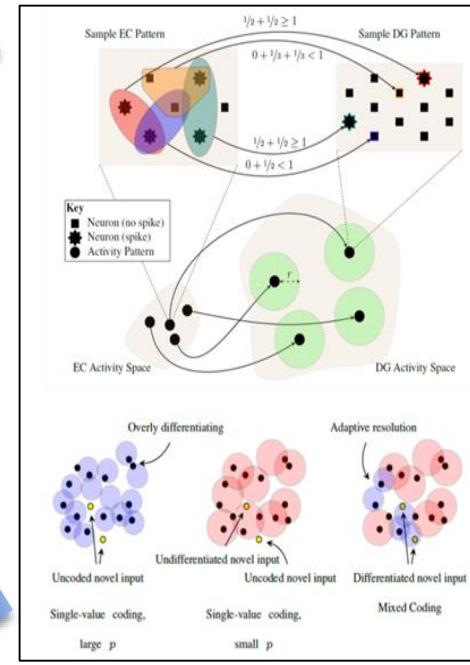

Dentate Gyrus with Green New Neurons

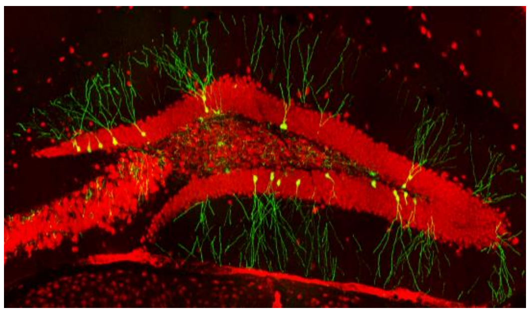

Identify critical aspects of computation

- **expansive scale**
- **sparse activation**
- **neurogenesis**

Adaptive Neural Hardware, Architectures, and Algorithms

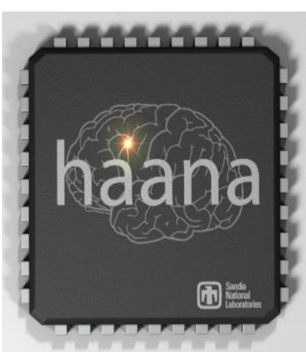

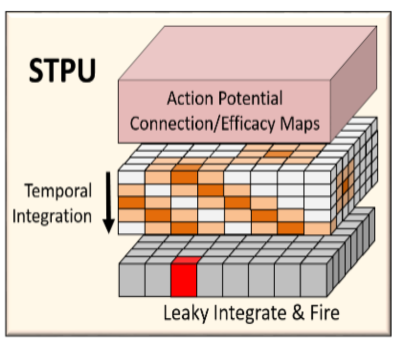

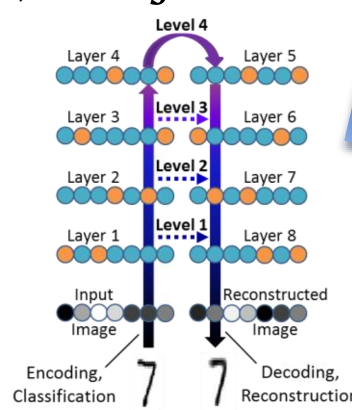

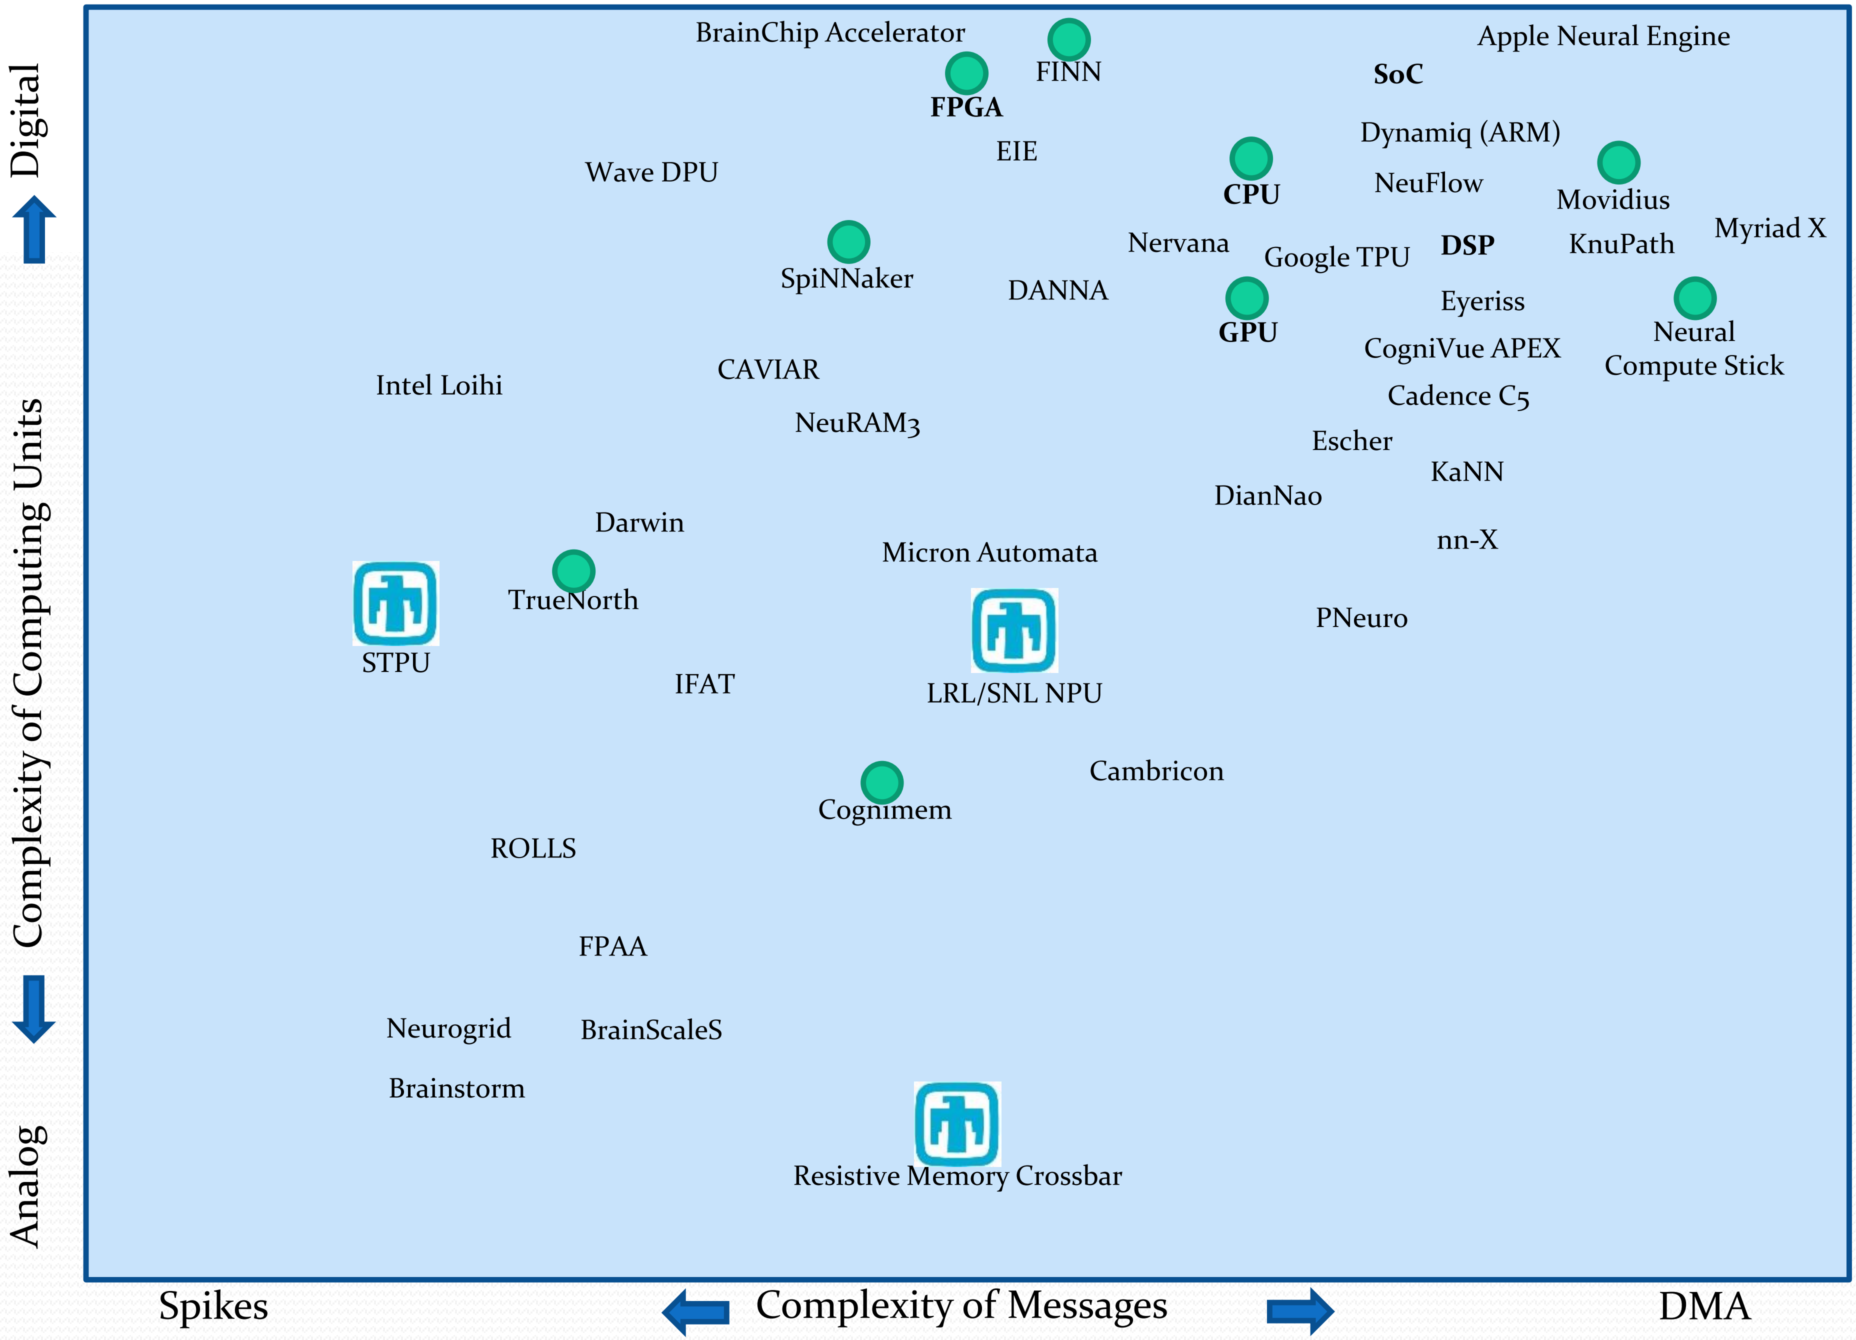

Digital
Complexity of Computing Units
Analog
BrainChip Accelerator
FINN
FPGA
Apple Neural Engine
SoC
Dynamiq (ARM)
EIE
Wave DPU
CPU
NeuFlow
Movidius
Myriad X
Nervana
Google TPU
DSP
KnuPath
SpiNNaker
DANNA
GPU
Eyeriss
Neural
Compute Stick
CogniVue APEX
CAVIAR
Intel Loihi
Cadence C5
NeuRAM3
Escher
KaNN
DianNao
Darwin
Micron Automata
nn-X
TrueNorth
STPU
PNeuro
IFAT
LRL/SNL NPU
Cambricon
Cognimem
ROLLS
FPAA
Neurogrid
BrainScaleS
Brainstorm
Resistive Memory Crossbar
Spikes
Complexity of Messages
DMA

MACHINE LEARNING & DEEP LEARNING

MLDL

Future of DL

- Intertwined Multimodal-Temporal Networks
- Semi-Supervised RL
- Low-shot learning
- Attention, Focus, Anticipation
  - What to ignore in streaming data
- Spiking algorithms
- Learnable hardware

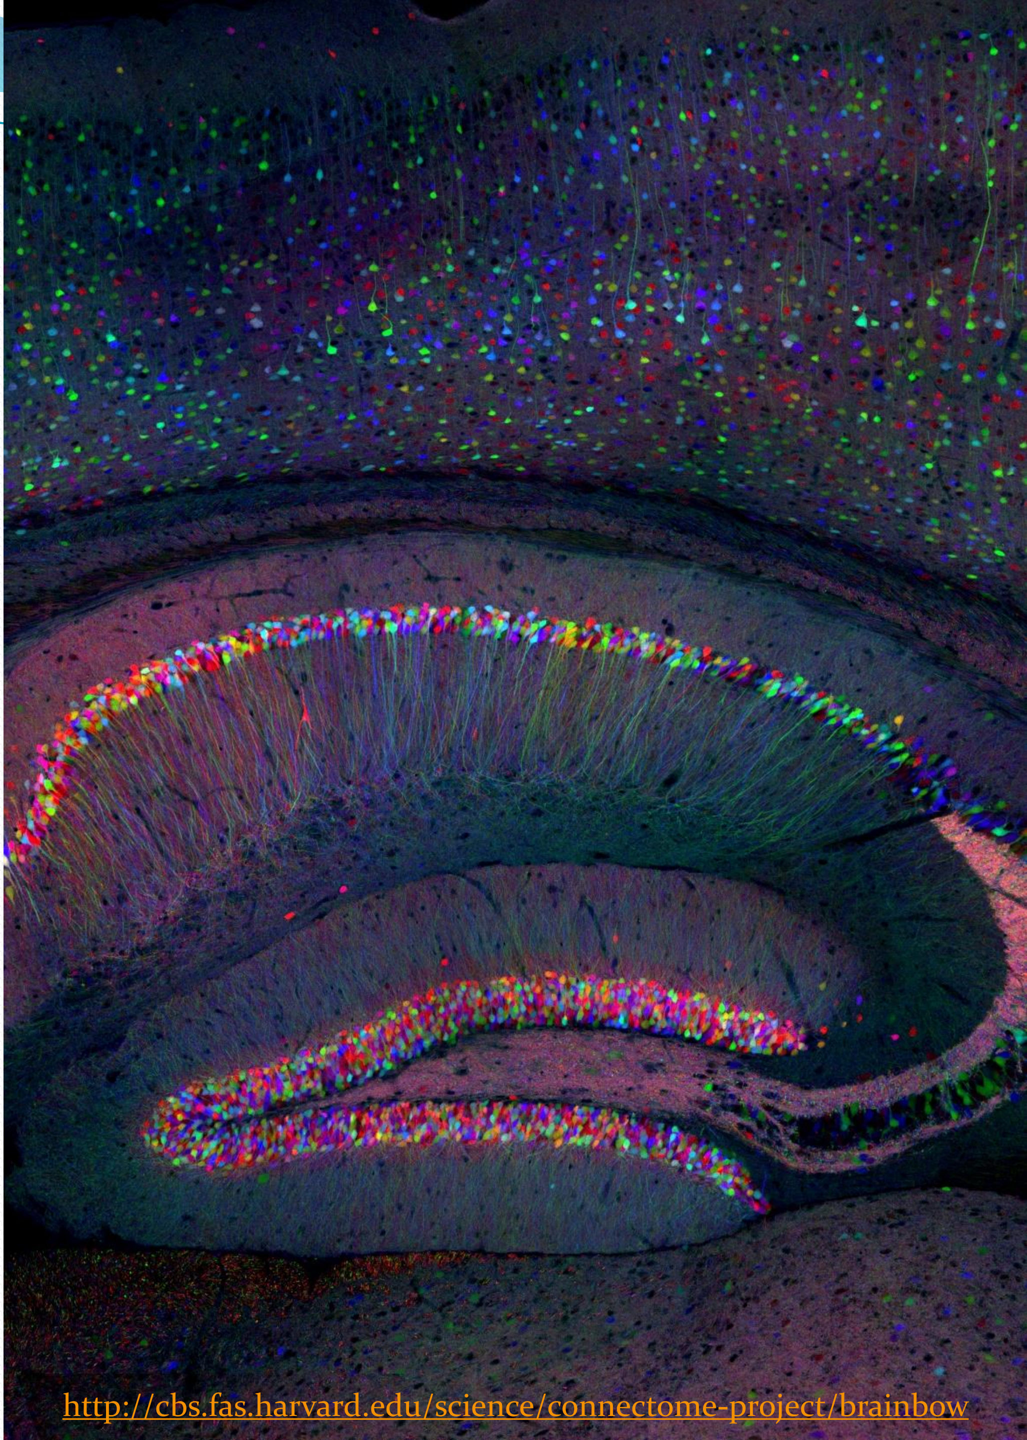

Turing Test Sequence*

- Can we tell the difference between a human and a robot?
  - T1 – Through conversations behind a curtain.
  - T2 – Through T1 and physical appearance.
  - T3 – Through T1, T2, and physical activity.
  - T4 – Through T1, T2, T3, and skin samples.
  - T5 – Through T1, T2, T3, T4, and brain scans.
  - T6 – Through T1, T2, T3, T4, T5, and surgery.
  - …

* Selmer Bringsjord, "What Robots Can and Can't Be"

AI = Tool in the hands of moral agents

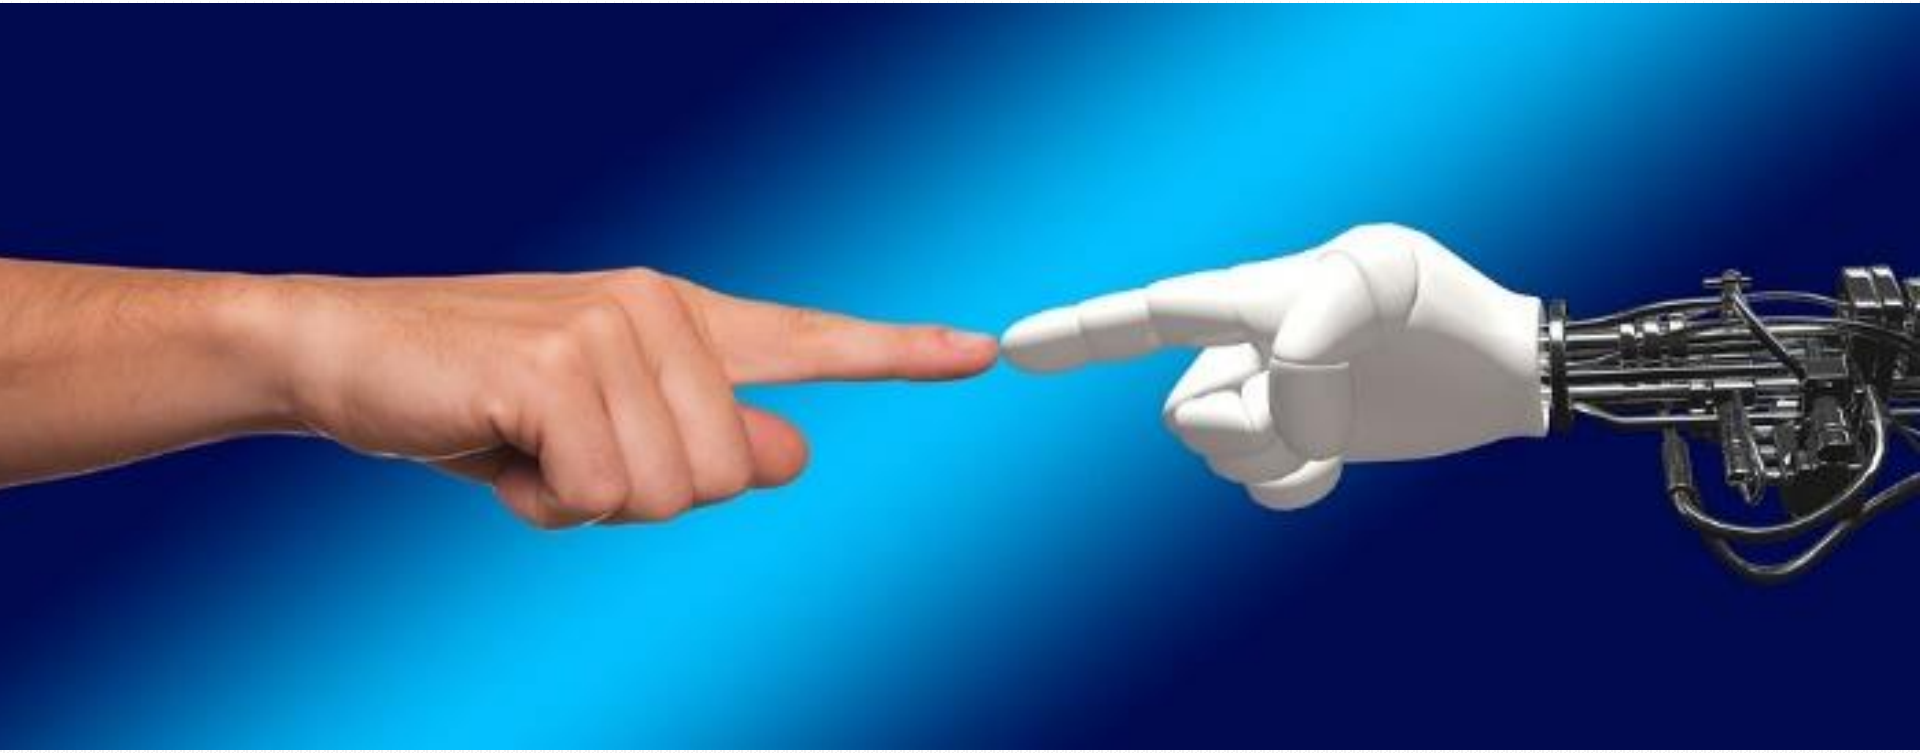

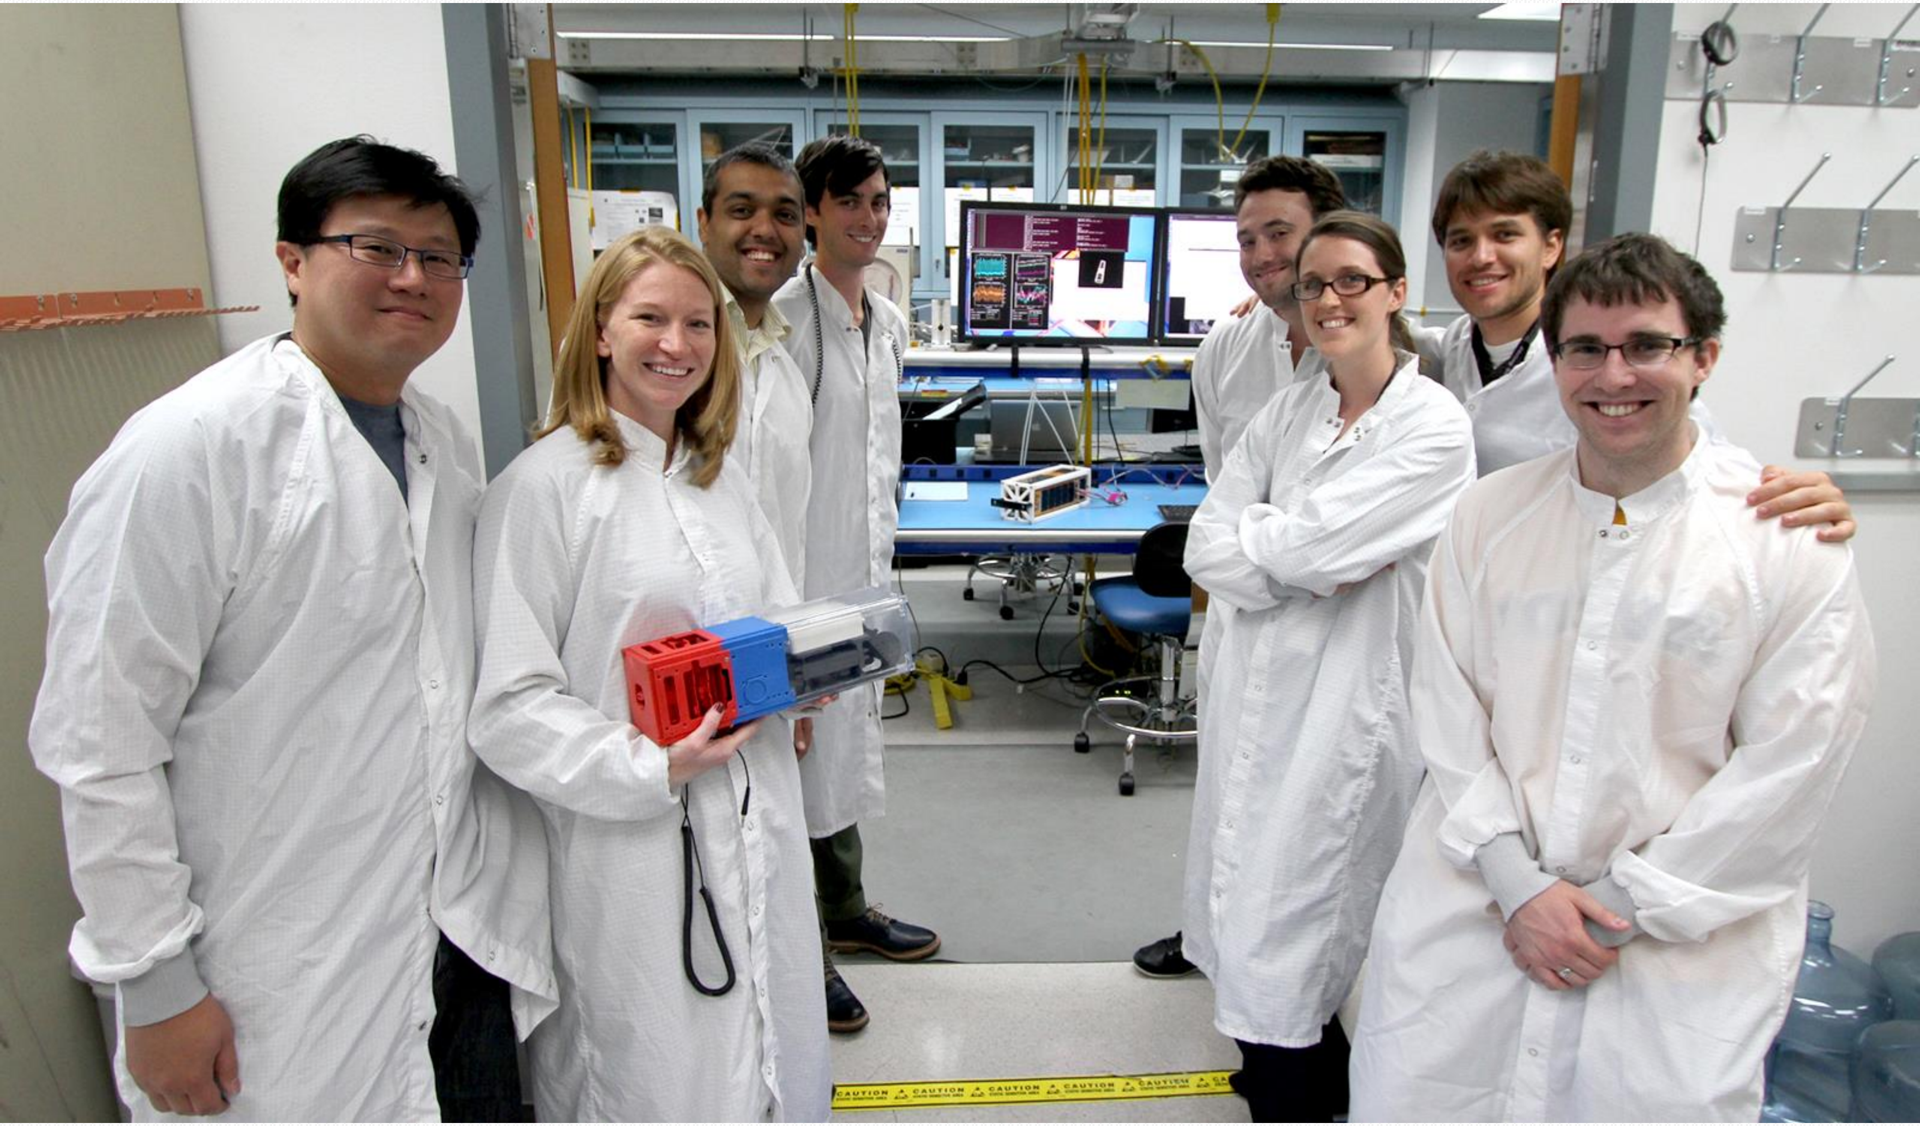

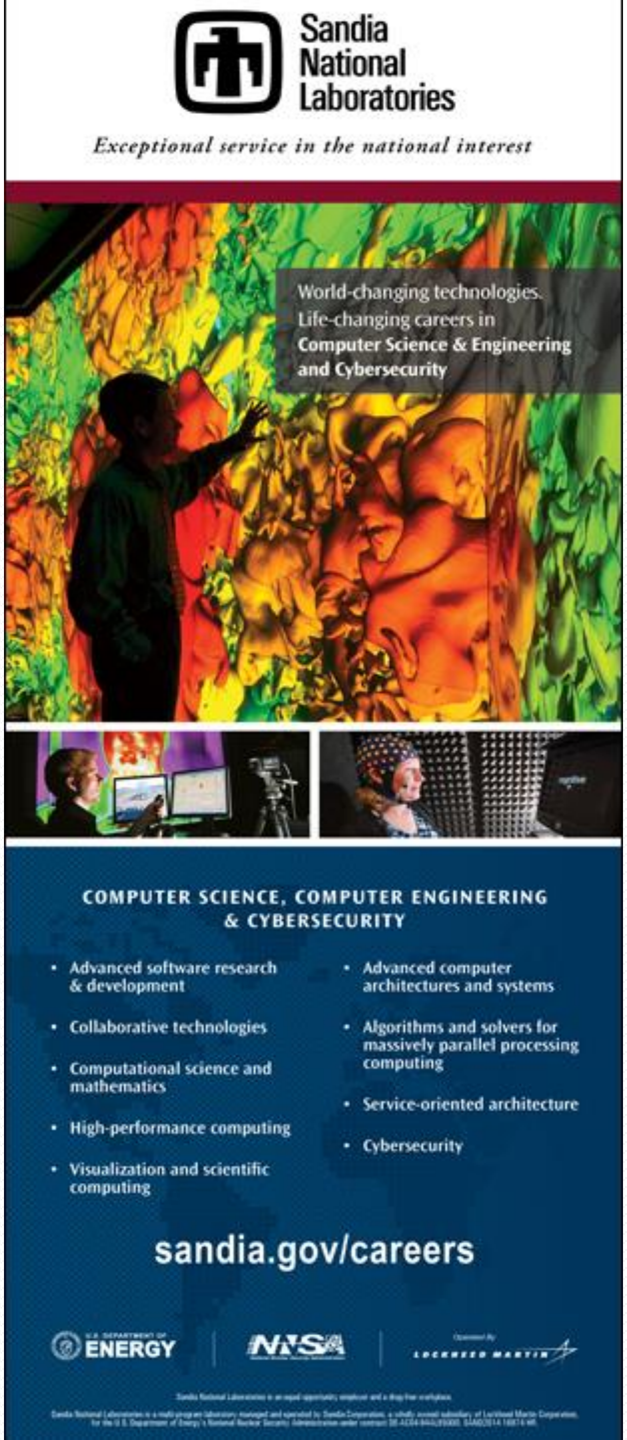
Sandia National Laboratories
Exceptional service in the national interest
World-changing technologies.
Life-changing careers in
Computer Science & Engineering and Cybersecurity
COMPUTER SCIENCE, COMPUTER ENGINEERING & CYBERSECURITY
• Advanced software research & development
• Collaborative technologies
• Computational science and mathematics
• High-performance computing
• Visualization and scientific computing
• Advanced computer architectures and systems
• Algorithms and solvers for massively parallel processing computing
• Service-oriented architecture
• Cybersecurity
sandia.gov/careers
ENERGY
NNSA
LOCKHEED MARTIN

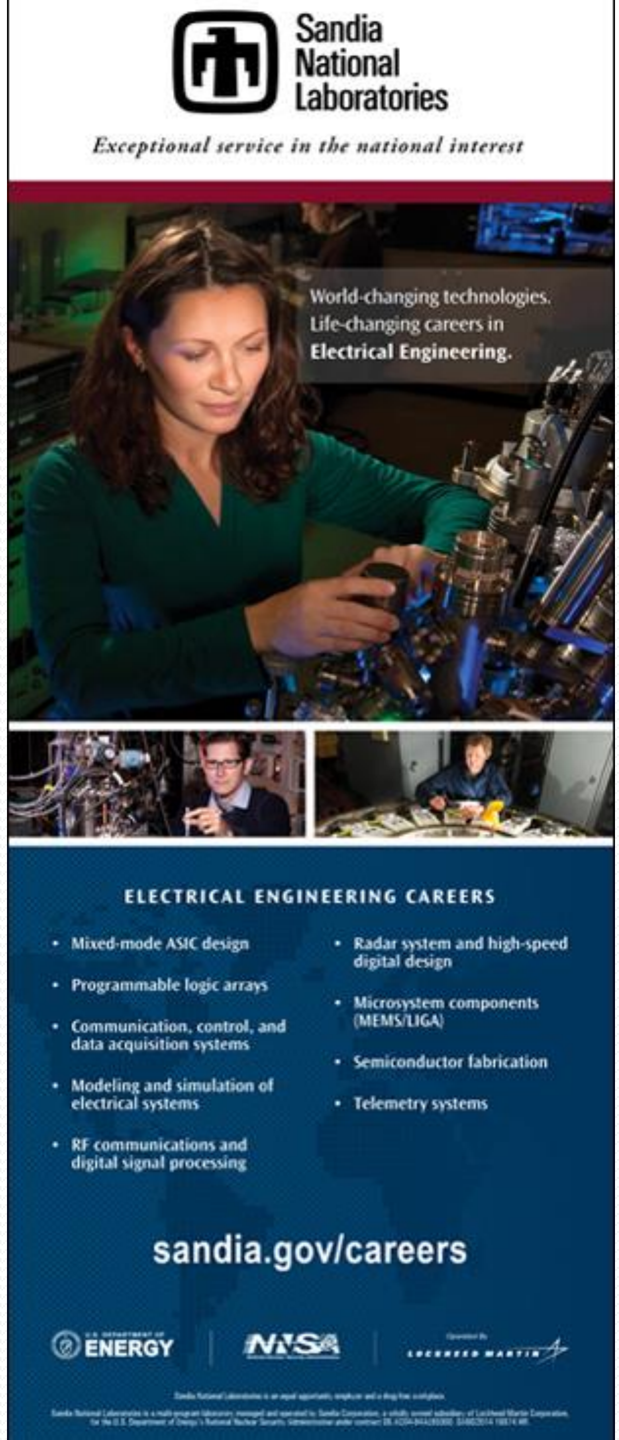
Sandia National Laboratories
Exceptional service in the national interest
World-changing technologies.
Life-changing careers in
Electrical Engineering.
ELECTRICAL ENGINEERING CAREERS
• Mixed-mode ASIC design
• Programmable logic arrays
• Communication, control, and data acquisition systems
• Modeling and simulation of electrical systems
• RF communications and digital signal processing
• Radar system and high-speed digital design
• Microsystem components (MEMS/LIGA)
• Semiconductor fabrication
• Telemetry systems
sandia.gov/careers
ENERGY
NNSA
LOCKHEED MARTIN

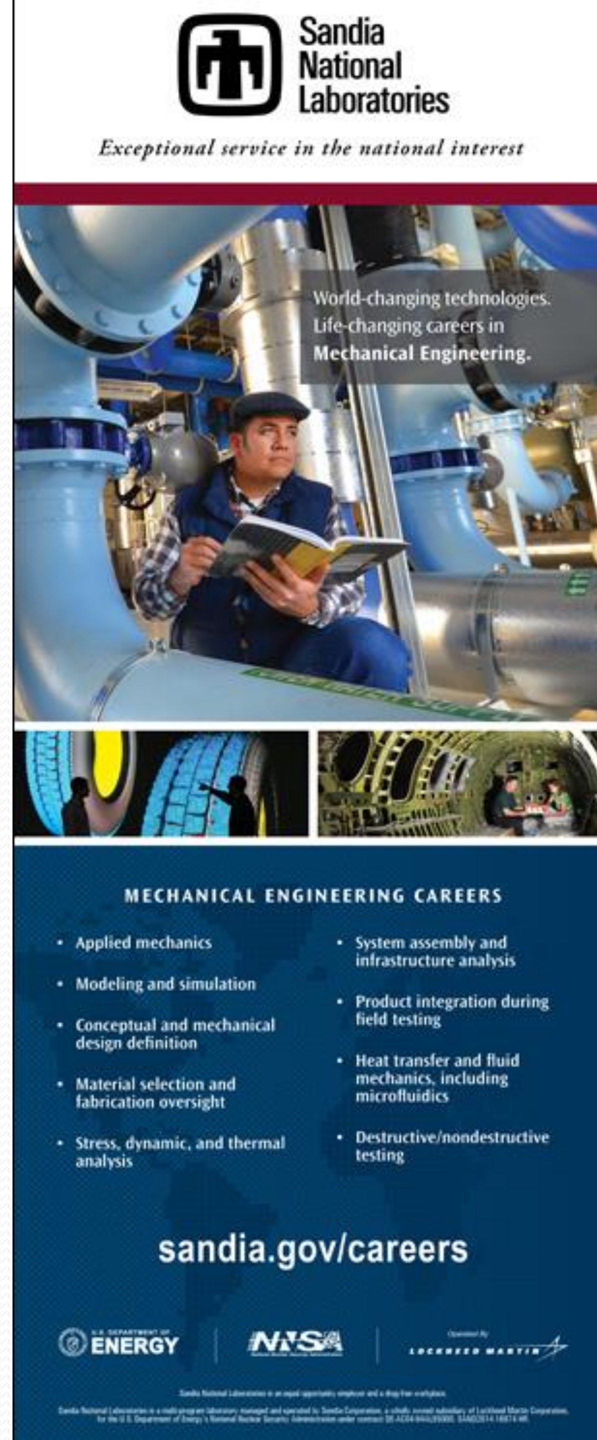
Sandia National Laboratories
Exceptional service in the national interest
World-changing technologies.
Life-changing careers in
Mechanical Engineering.
MECHANICAL ENGINEERING CAREERS
• Applied mechanics
• Modeling and simulation
• Conceptual and mechanical design definition
• Material selection and fabrication oversight
• Stress, dynamic, and thermal analysis
• System assembly and infrastructure analysis
• Product integration during field testing
• Heat transfer and fluid mechanics, including microfluidics
• Destructive/nondestructive testing
sandia.gov/careers
ENERGY
NNSA
LOCKHEED MARTIN