Embedded AI in Smart Robots, Drones, IoT: Teaching Emerging Technologies in the Classroom

Debasis Bhattacharya, JD, DBA
University of Hawaii Maui College
debasisb@wawaii.edu
http://maui.hawaii.edu/cybersecurity
Rajiv Malkan
Lone Star College, TX
July 25, 2019

UNIVERSITY of HAWAI'I
MAUI COLLEGE

LONE STAR COLLEGE

Agenda

- What is AI, Machine Learning and Deep Learning?
- Embedded AI Devices – NVIDIA JETSON, NANO and AMZN Deep Lens
- Teaching AI and Cybersecurity Across Disciplines
- Case Studies from UH – AI Drone, Sentiment Analysis and Medical Imaging
- Case Studies from Lone Star College, TX
- Q&A, Discussion

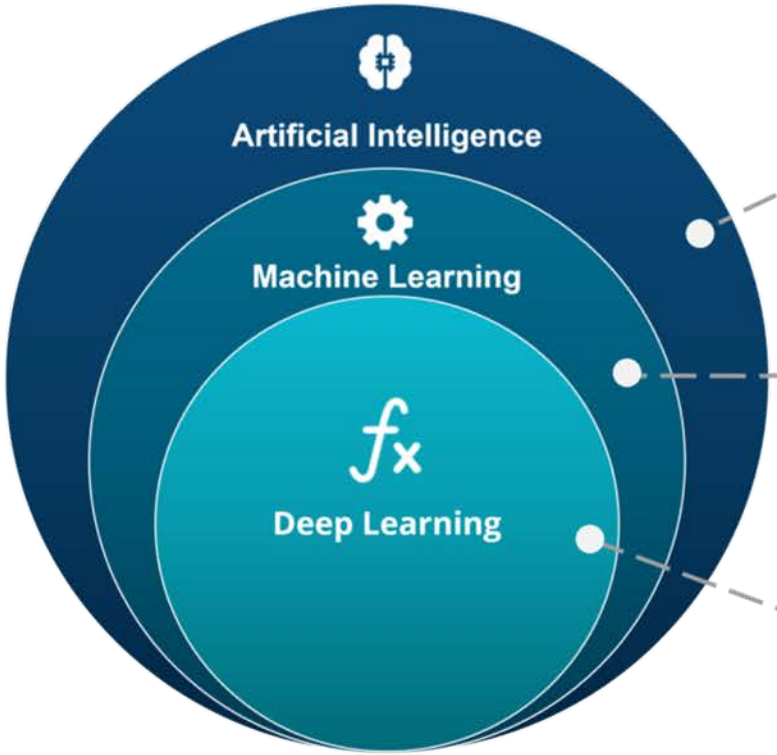

ARTIFICIAL INTELLIGENCE
A technique which enables machines to mimic human behaviour

MACHINE LEARNING
Subset of AI technique which use statistical methods to enable machines to improve with experience

DEEP LEARNING
Subset of ML which make the computation of multi-layer neural network feasible

Source: Edureka

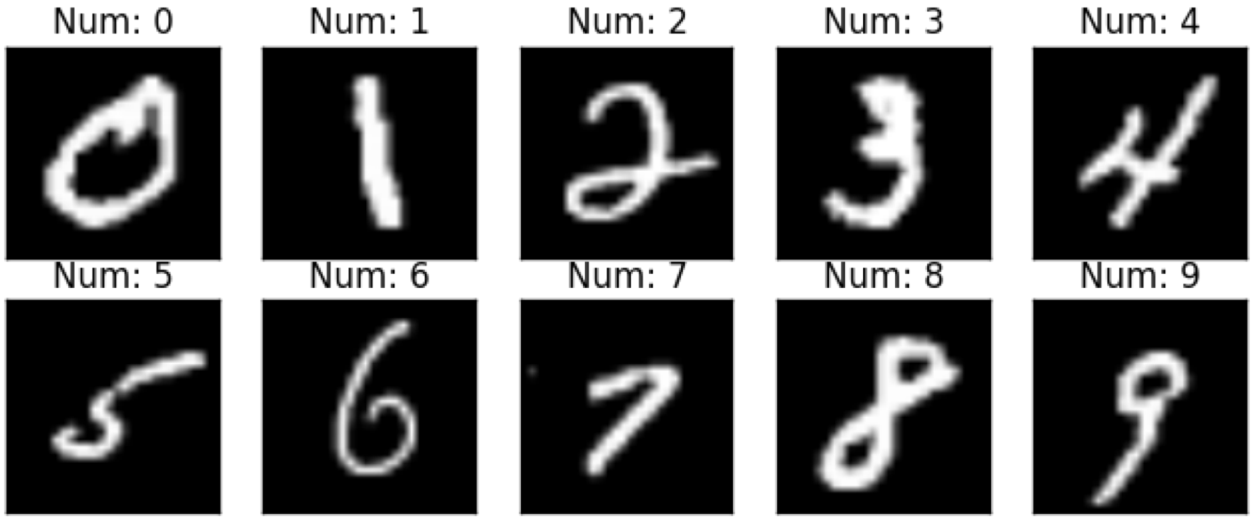

Source: TensorFlow for R - RStudio

Biological Neuron versus Artificial Neural Network

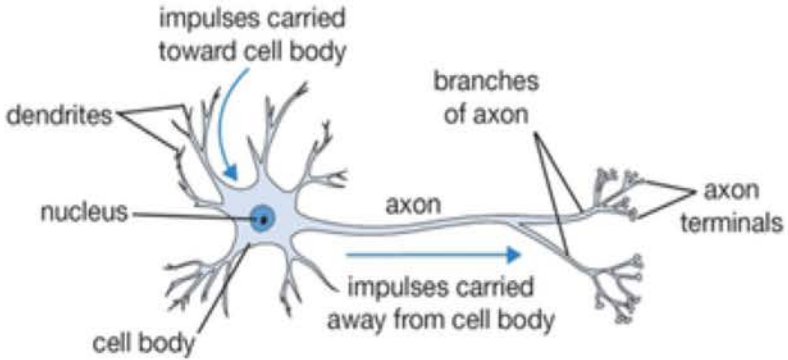

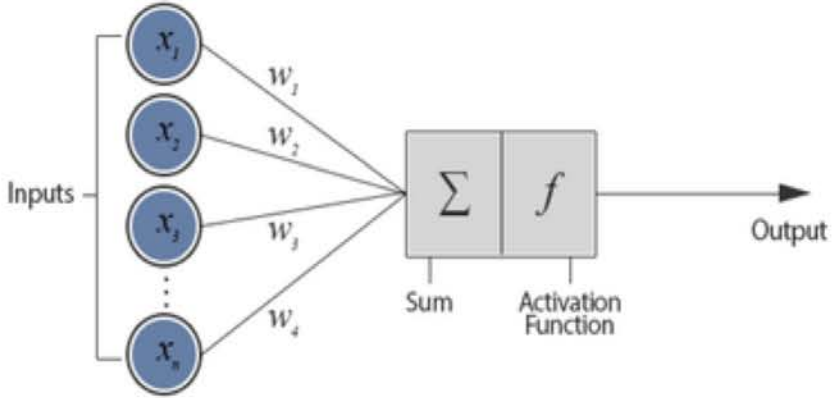

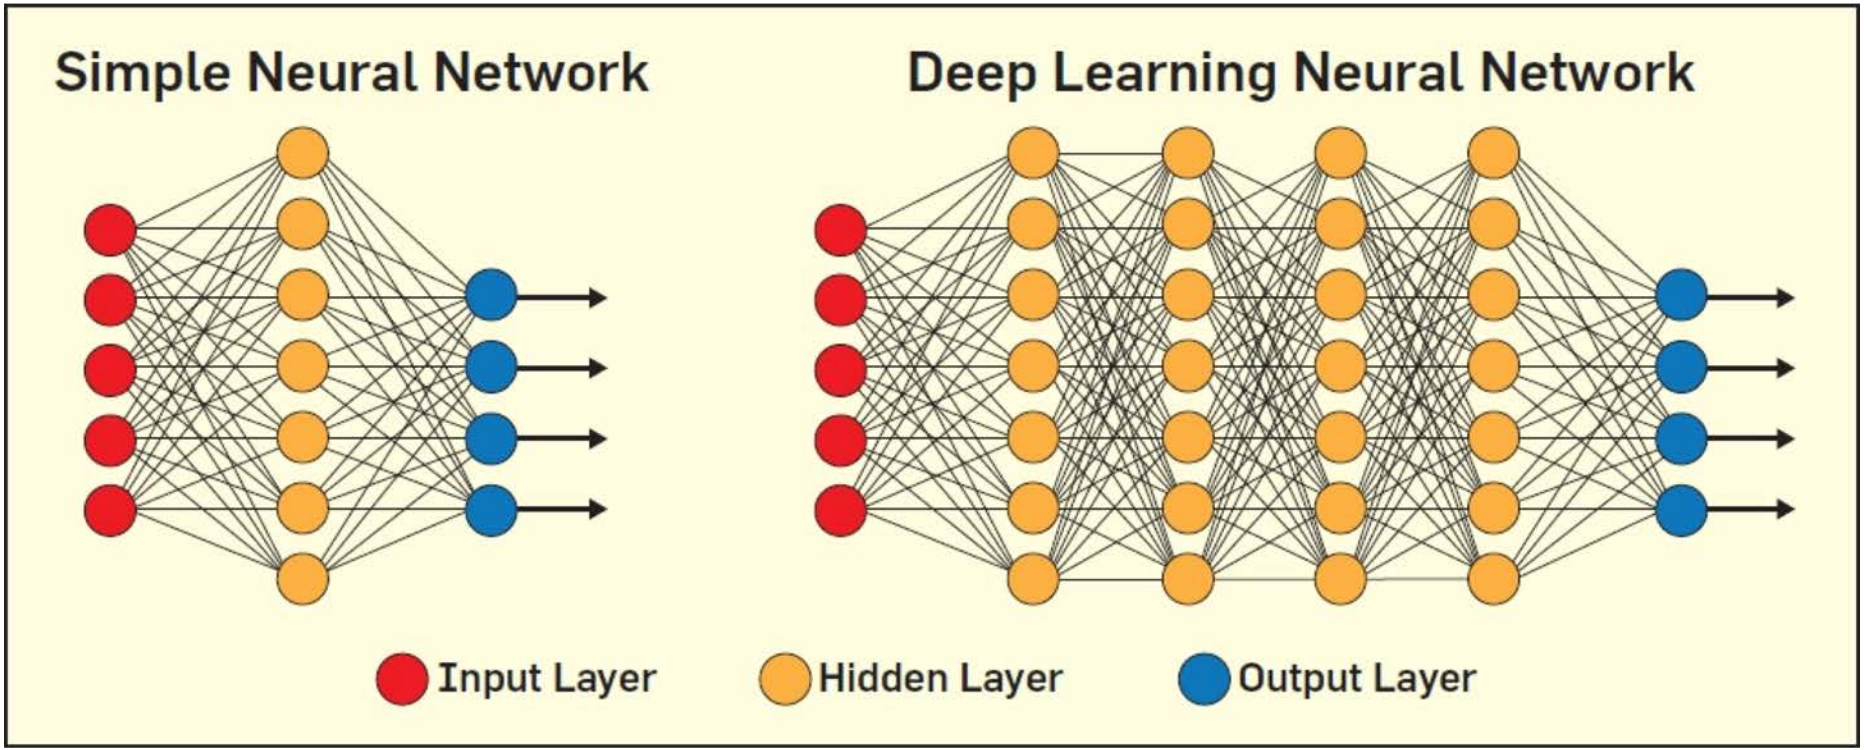

Source: ACM - Computers in Entertainment - Association for Computing Machinery

Convolutional Neural Network - ConvNet

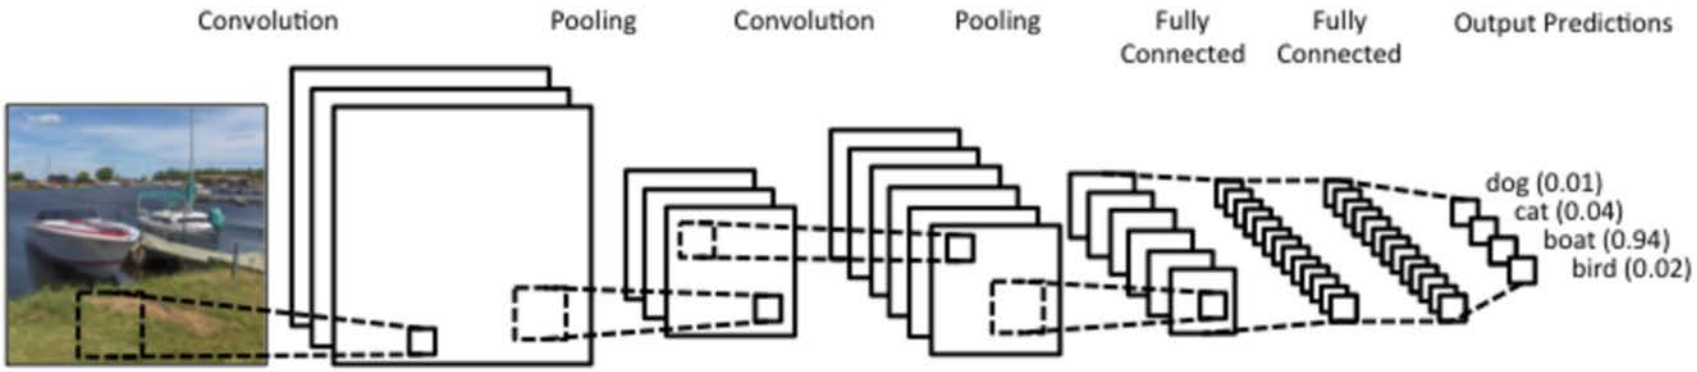

Source: https://ujjwalkarn.me/2016/08/11/intuitive-explanation-convnets/

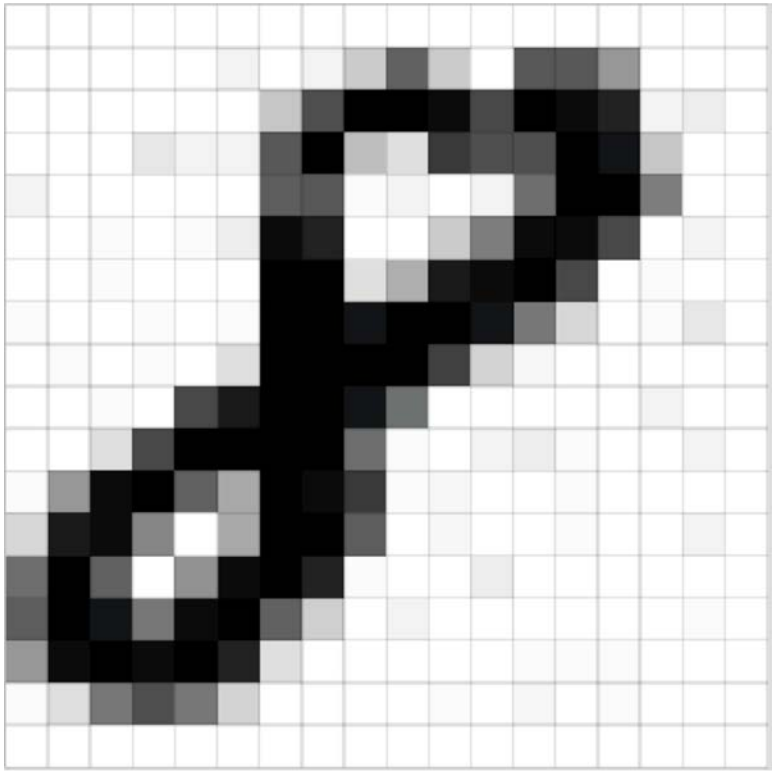

Image - Matrix of Pixels

Channels - Color images have 3 channels - RGB. Each Pixel Value ranges from 0 to 255

Grayscale Image - One channel, values 0 (white) to 255 (black)

Source: https://ujjwalkarn.me/2016/08/11/intuitive-explanation-convnets/

Convolution - Extract Features

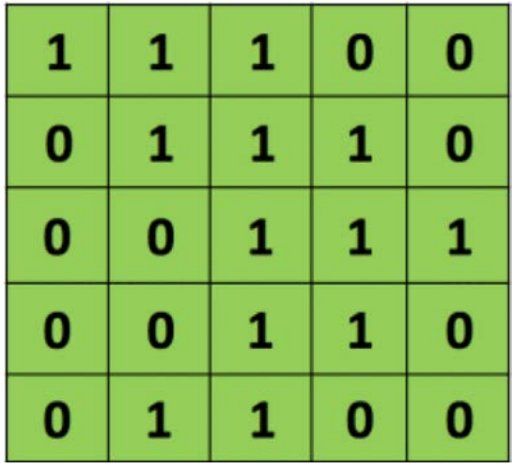

Input Image, 5x5

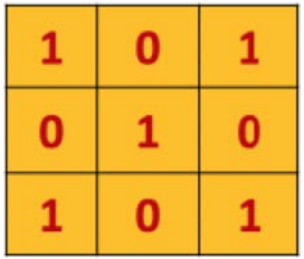

3x3 Filter, or Kernel or Feature Detector

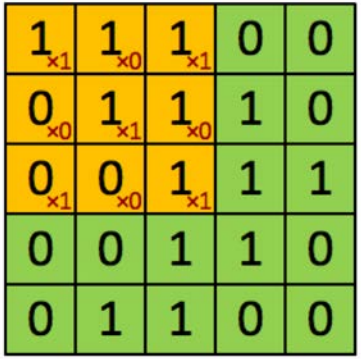

Image

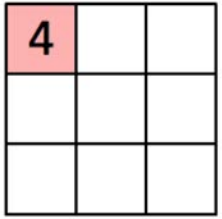

Convolved Feature

Move 3x3 Filter over 5x5 Input Image, one pixel at a time (STRIDE) and compute matrix multiplication. Convolved Feature or Feature Map

Source: https://ujjwalkarn.me/2016/08/11/intuitive-explanation-convnets/

Operation	Filter	Convolved Image
Identity	$\begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}$	
Edge detection	$\begin{bmatrix} 1 & 0 & -1 \\ 0 & 0 & 0 \\ -1 & 0 & 1 \end{bmatrix}$	
	$\begin{bmatrix} 0 & 1 & 0 \\ 1 & -4 & 1 \\ 0 & 1 & 0 \end{bmatrix}$	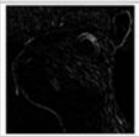
	$\begin{bmatrix} -1 & -1 & -1 \\ -1 & 8 & -1 \\ -1 & -1 & -1 \end{bmatrix}$	
Sharpen	$\begin{bmatrix} 0 & -1 & 0 \\ -1 & 5 & -1 \\ 0 & -1 & 0 \end{bmatrix}$	
Box blur (normalized)	$\frac{1}{9}\begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}$	
Gaussian blur (approximation)	$\frac{1}{16}\begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix}$	

Different values of the Filter or Kernel will create different Feature Maps from the Input Image

Examples of Features - Edges, Blurs, Sharpen

Parameters

Depth - # of Filters
Stride - movement across image
Zero-Padding - apply filter to edges

Source: https://ujjwalkarn.me/2016/08/11/intuitive-explanation-convnets/

Fully Connected Layer (FCNs)

Rectified and Pooled Feature Maps are connected together into on FCN
Training Data Set is used to Classify Image
SoftMax Activation Function used to create vector of values between 0 and 1

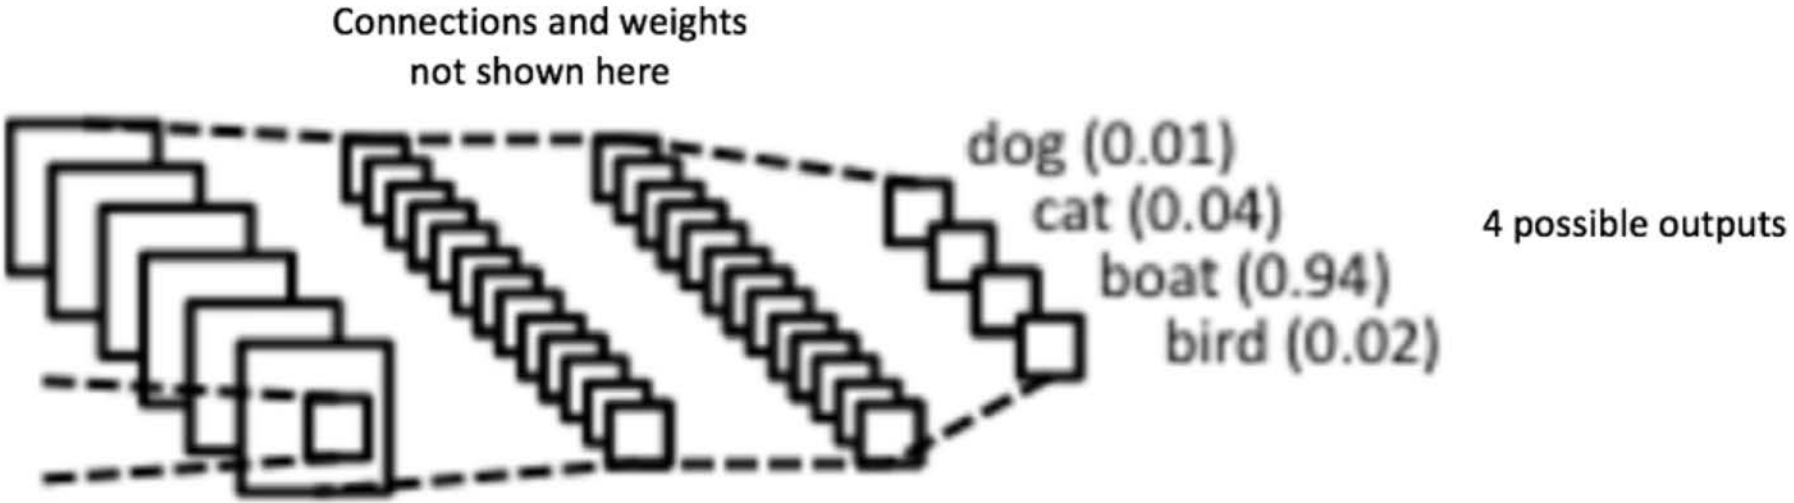

Source: https://ujjwalkarn.me/2016/08/11/intuitive-explanation-convnets/

Forward Propagation

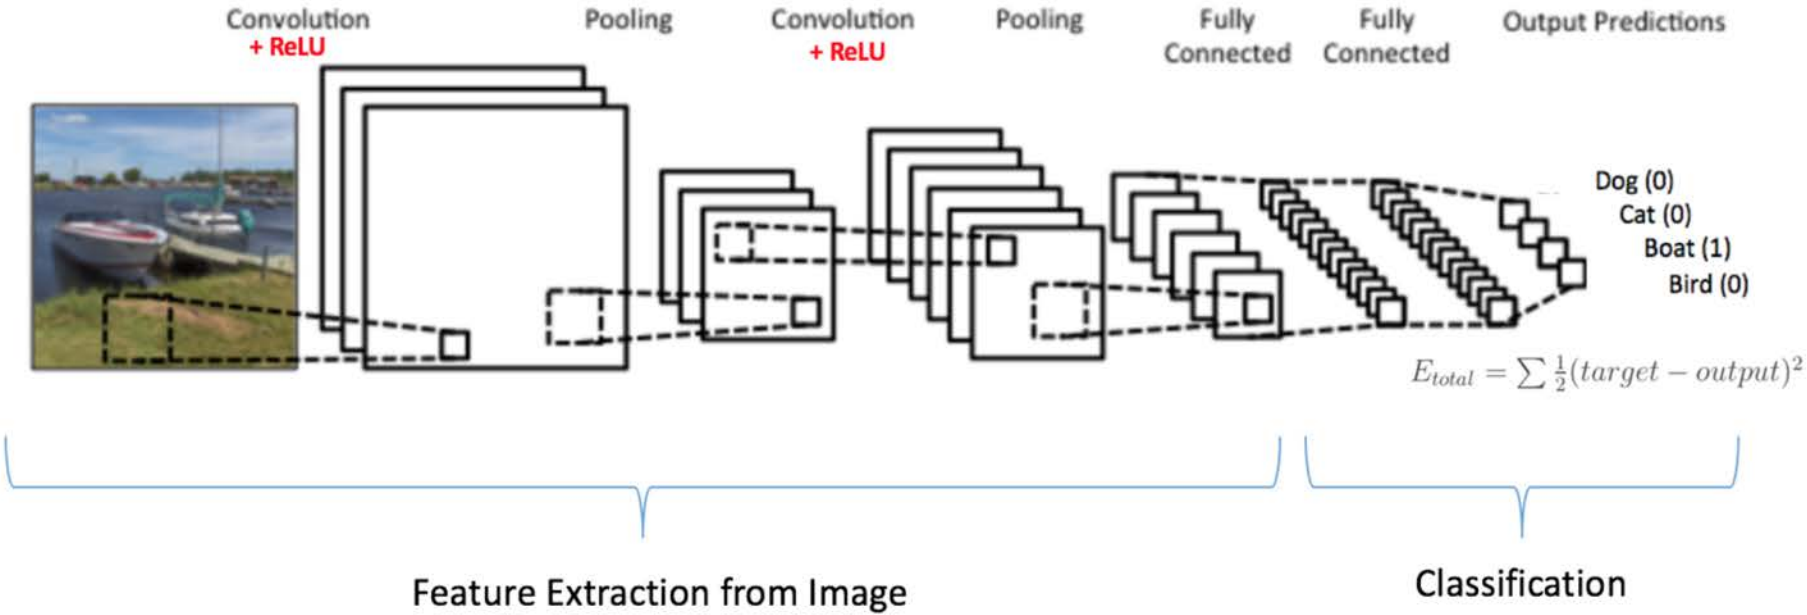

Source: https://ujjwalkarn.me/2016/08/11/intuitive-explanation-convnets/

Backward Propagation

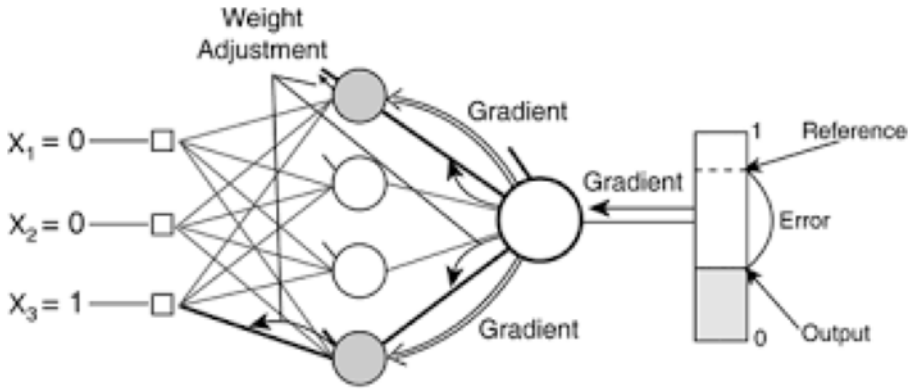

Calculate Gradient of error
Use Gradient Descent to update filter weights
Reduce Output Error or **Training Loss**
Epoch = Forward + Backward Propagation
Hyperparameter = Learning Rate
Validation Data -> Forward Propagation Only
Minimize Training Loss & Validation Loss
Control Overfitting using Dropouts
Allow for Generalization of New Test Data

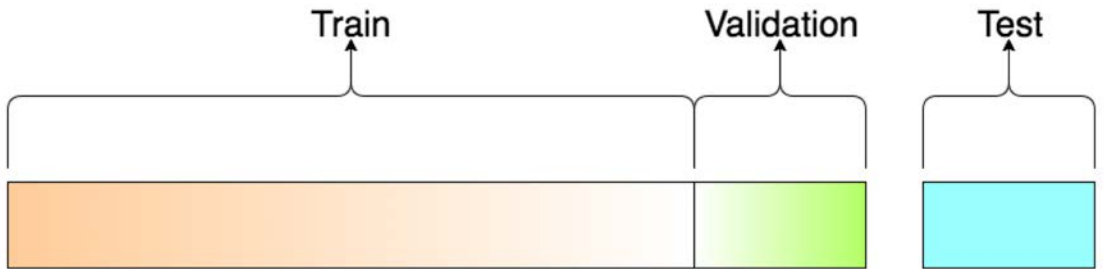

Source: https://ujjwalkarn.me/2016/08/11/intuitive-explanation-convnets/

Object Detection with YOLO

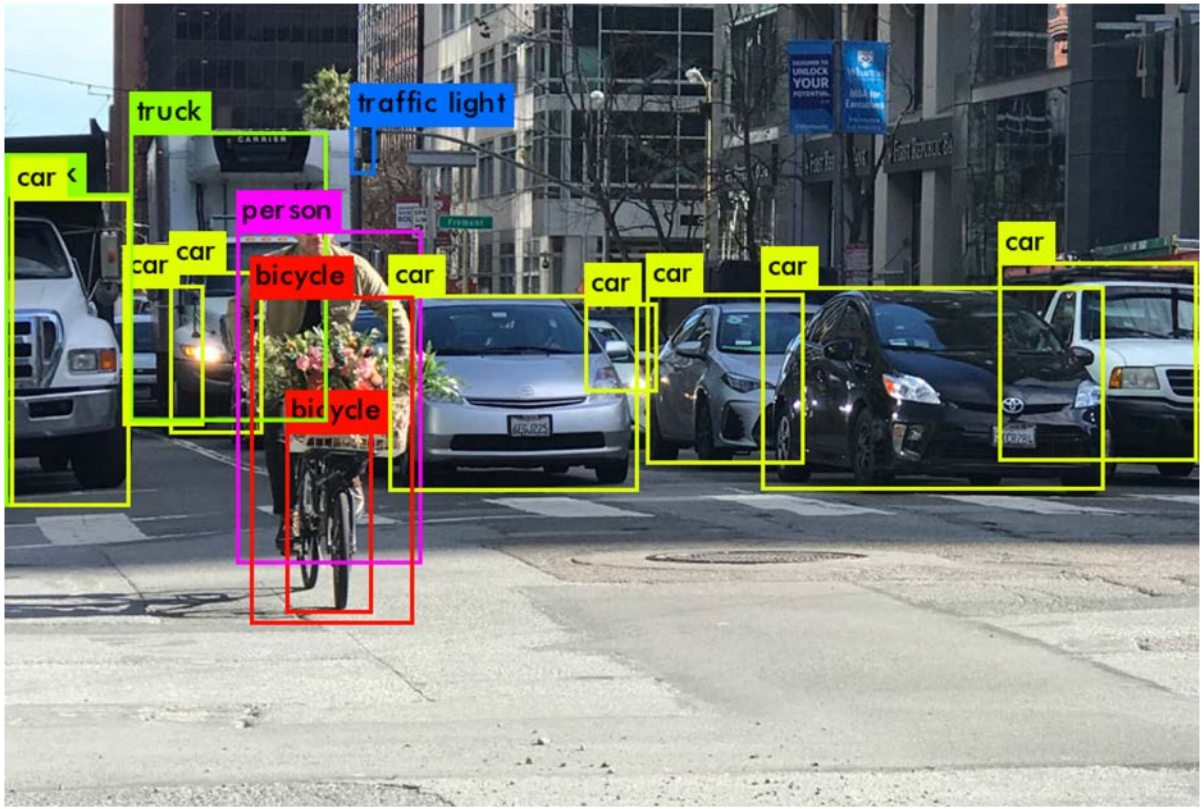

Bounding Boxes Detect and Classify Objects

Joseph Redmon link on YOLO at Ted Talk here

Source : Medium

Image Segmentation using CNN

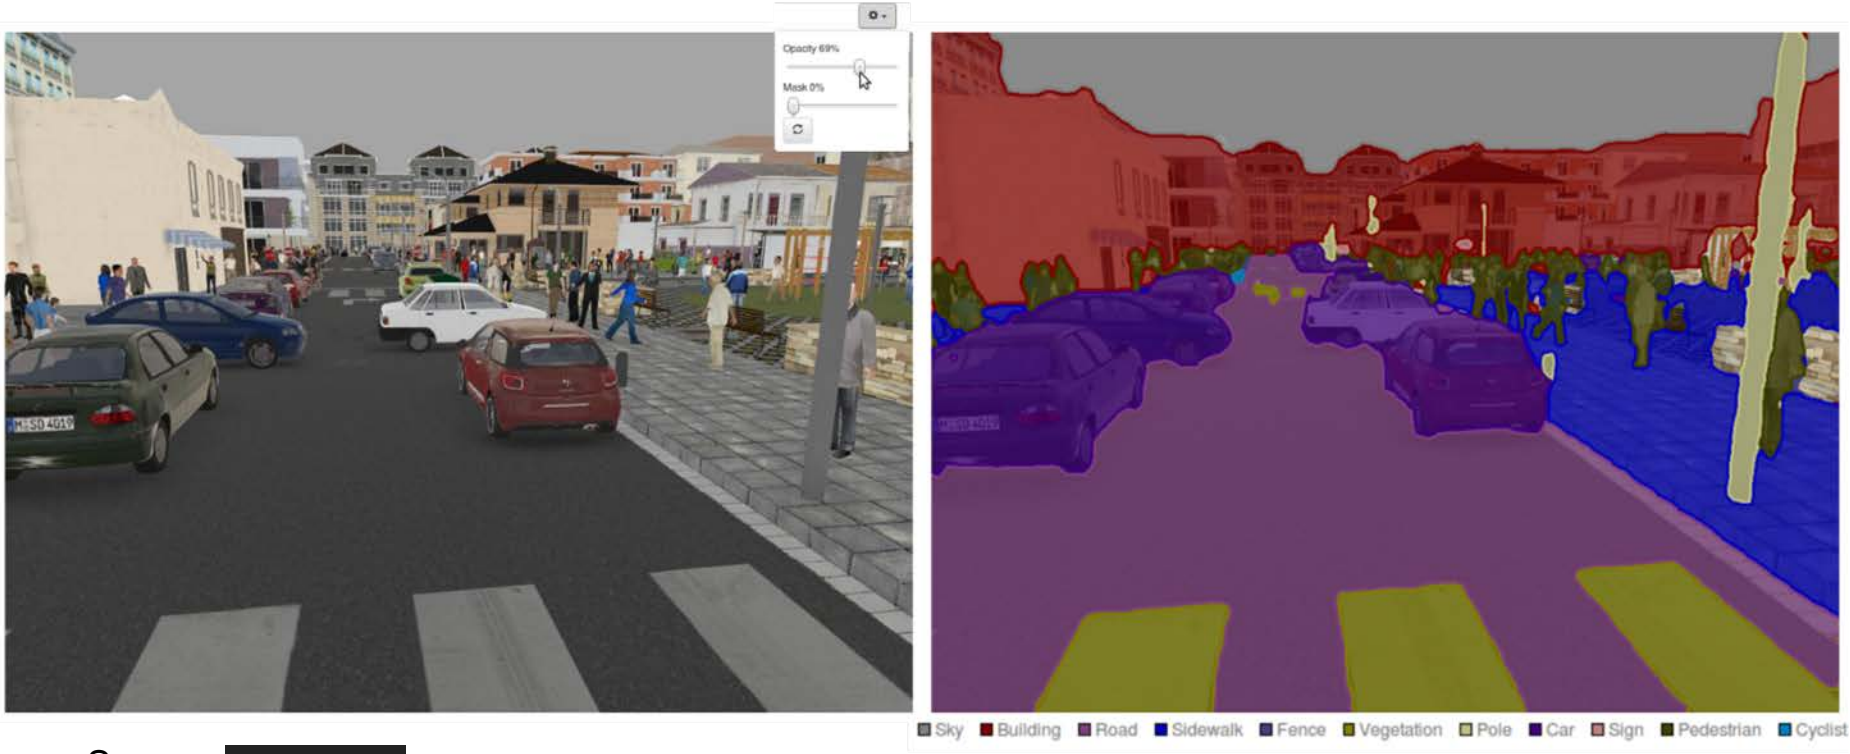

Source: Good Audience

Harness AI at the Edge with the Jetson TX2 Developer Kit

The Jetson TX2 Developer Kit gives you a fast, easy way to develop hardware and software for the Jetson TX2 AI supercomputer on a module. It exposes the hardware capabilities and interfaces of the developer board, comes with design guides and other documentation, and is pre-flashed with a Linux development environment. It also supports NVIDIA Jetpack—a complete SDK that includes the BSP, libraries for deep learning, computer vision, GPU computing, multimedia processing, and much more.

The Tech Specs

JETSON TX2 MODULE

- NVIDIA Pascal™ Architecture GPU
- 2 Denver 64-bit CPUs + Quad-Core A57 Complex
- 8 GB L128 bit DDR4 Memory
- 32 GB eMMC 5.1 Flash Storage
- Connectivity to 802.11ac Wi-Fi and Bluetooth-Enabled Devices
- 10/100/1000BASE-T Ethernet

JETSON CAMERA MODULE

- 5 MP Fixed Focus MIPI CSI Camera

BUTTONS

- Power On/Off
- Reset
- Force Recovery
- User-Defined

I/O

- USB 3.0 Type A
- USB 2.0 Micro AB (supports recovery and host mode)
- HDMI
- M.2 Key E
- PCI-E x4
- Gigabit Ethernet
- Full-Size SD
- SATA Data and Power
- GPIOs, I2C, I2S, SPI, CAN*
- TTL UART with flow control
- Display Expansion Header*
- Camera Expansion Header*
    - *I/O expansion headers: refer to product documentation for header specification.

POWER OPTIONS

- External 19V AC Adapter

KIT CONTENTS

- NVIDIA Jetson TX2 Developer Board
- AC Adaptor
- Power Cord
- USB Micro-B to USB A Cable
- USB Micro-B to Female USB A Cable
- Rubber Feet (4)
- Quick Start Guide
- Safety Booklet
- Antennas to Connect to Wi-Fi-Enabled Devices (2)

NVIDIA Jetson Nano enables the development of millions of new small, low-power AI systems. It opens new worlds of embedded IoT applications, including entry-level Network Video Recorders (NVRs), home robots, and intelligent gateways with full analytics capabilities.

TECHNICAL SPECIFICATIONS

GPU	NVIDIA Maxwell™ architecture with 128 NVIDIA CUDA® cores
CPU	Quad-core ARM® Cortex®-A57 MPCore processor
Memory	4 GB 64-bit LPDDR4
Storage	16 GB eMMC 5.1 Flash
Video Encode	4K @ 30 (H.264/H.265)
Video Decode	4K @ 60 (H.264/H.265)
Camera	12 lanes (3x4 or 4x2) MIPI CSI-2 DPHY 1.1 (18 Gbps)
Connectivity	Gigabit Ethernet
Display	HDMI 2.0 or DP1.2 | eDP 1.4 | DSI (1 x2) 2 simultaneous
UPHY	1 x1/2/4 PCIE, 1x USB 3.0, 3x USB 2.0
I/O	1x SDIO / 2x SPI / 4x I2C / 2x I2S / GPIOs -> I^2C, I^2S
Size	69.6 mm x 45 mm
Mechanical	260-pin edge connector

The world's first deep learning enabled video camera for developers

AWS DeepLens helps put machine learning in the hands of developers, literally, with a fully programmable video camera, tutorials, code, and pre-trained models designed to expand deep learning skills.

Pre-Order (2019 Edition)

Register your DeepLens

A new way to learn machine learning

AWS DeepLens allows developers of all skill levels to get started with deep learning in less than 10 minutes by providing sample projects with practical, hands-on examples which can start running with a single click.

Custom built for deep learning

AWS DeepLens was designed with deep learning in mind. With over 100 GFLOPS of compute power on the device, it can process deep learning predictions on HD video for real time.

Build custom models with Amazon SageMaker

Models trained in Amazon SageMaker can be sent to AWS DeepLens with just a few clicks from the AWS Management Console.

Broad framework support

AWS developers can run any deep learning framework, including TensorFlow and Caffe. AWS DeepLens comes pre-installed with a high performance, efficient, optimized inference engine for deep learning using Apache MXNet.

Integrated with AWS

AWS DeepLens integrates with Amazon Rekognition for advanced image analysis, Amazon SageMaker for training models, and with Amazon Polly to create speech-enabled projects. The device also connects securely to AWS IoT, Amazon SQS, Amazon SNS, Amazon S3, Amazon DynamoDB, and more.

Fully programmable

AWS DeepLens is easy to customize and is fully programmable using AWS Lambda. The deep learning models in DeepLens even run as part of an AWS Lambda function, providing a familiar programming environment to experiment with.

Deep Learning Tools

CNN Architectures

- LeNet
- AlexNet
- ZF Net
- GoogLeNet
- ResNets

- cuDNN
- DIGITS
- DetectNet
- TensorRT
- RAPIDS
- JETSON

Frameworks

- TensorFlow
- Caffe
- MS CNTK
- PyTorch/Torch
- MXNet
- Chainer
- Keras
- Theano
- Deeplearning4j
- PaddlePaddle
- MATLAB

Teaching AI and Cybersecurity Across Disciplines

- Electronics
- Healthcare
- Hospitality and Tourism
- Business, Finance, Accounting
- Criminal Justice
- Computer Science
- Mathematics
- Etc.

- *How do you teach the essence of AI across the disciplines?*

Case Study: Sentiment Analysis in Hospitality/Tourism

- Sentiment Analysis may be performed as an application of Machine Learning (ML) to large bodies of text, such as those found in large consumer review datasets, in order to determine sentiment (positive, negative, sarcastic, etc.) and gain feedback.
- The use of Machine Learning techniques in this endeavor allows for much larger quantities of data to be processed than would be practical for human evaluators working directly with the data.
- With recent advances in Machine Learning in the form of new and powerful frameworks, it is relatively simple to set up a machine to perform analysis on text in a way previously confined to the domain of commonsense, human interpretation of opinions, feelings, etc.
- Sample Code - https://github.com/UHMC/nifty-sentiment-analysis

Case Study: Neural Network for Medical Image Processing

- Medical imaging is becoming an increasingly popular application of Machine Learning. Medical practitioners can use software to obtain diagnosis or second opinions on X-Ray images, lowering the chances of a missed threat.
- In this assignment, students will be able to set up a model to train using the Deep Learning Tool Kit and Tensorflow.
- Applicable for Students in various disciplines – CompSci, Healthcare etc.
- Source Code - https://github.com/UHMC/nifty-medical-imaging

Best Practices from Lone Star College, TX

http://www.lonestar.edu/land-survey-aas.htm


pepper

Pepper the robot is now working at a bank in New York City

UPS Enhances Driver Safety Training With Virtual Reality

US 08/15/17

Atlanta, GA

- ***VR experience to debut at nine U.S. Integrad® facilities this year***
- ***UPS IT team created content that displays on VR headsets***
- ***UPS Integrad technology supports company's safety culture***

UPS (NYSE:UPS) today said it will start training student delivery drivers to spot and identify road hazards using Virtual Reality (VR) headsets that vividly simulate the experience of driving on city streets while teaching a more memorable classroom lesson.

THETA

Product ▾ Topics ▾ Support ▾ Download Buy ▾ Solutions ▾


VIVE VR SYSTEM
• Fully immersive first-person experience
• Featuring SteamVR Tracking
NOW $ 499 *
With 6-month VIVEPORT Infinity membership, $78 value. Code included in box**.
**Offer subject to change.
Test Your PC
Buy Now
PayPal CREDIT
12 Easy Payments as low as $41.58/mo for est. total payments of $499.00
Subject to credit approval.

Questions? Comments

Debasis Bhattacharya, JD, DBA
University of Hawaii Maui College
debasisb@wawaii.edu
http://maui.hawaii.edu/cybersecurity
Rajiv Malkan
Lone Star College, TX
July 25, 2019

UNIVERSITY of HAWAI'I
MAUI COLLEGE

LONE STAR COLLEGE