Machine Learning for Embedded Software Engineers

Class 3: Machine Learning Applications: Vision and Speech

April 24, 2019

Jacob Beningo

DesignNews

CEC CONTINUING EDUCATION CENTER

Presented by: Digi-Key ELECTRONICS

Course Overview

Topics:

- Introduction to Machine Learning
- Machine Learning Architectures for Embedded Systems
- **Machine Learning Applications: Vision and Speech**
- Machine Vision with OpenMV
- Near Real-time Machine Learning using Coral

Session Overview

- Machine Learning in Embedded
- Machine Vision
- Keyword Spotting
- Best Practices

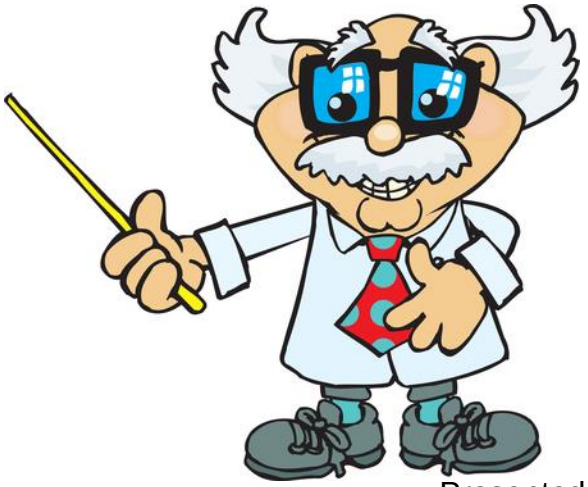

Presented by:

Machine Learning in Embedded

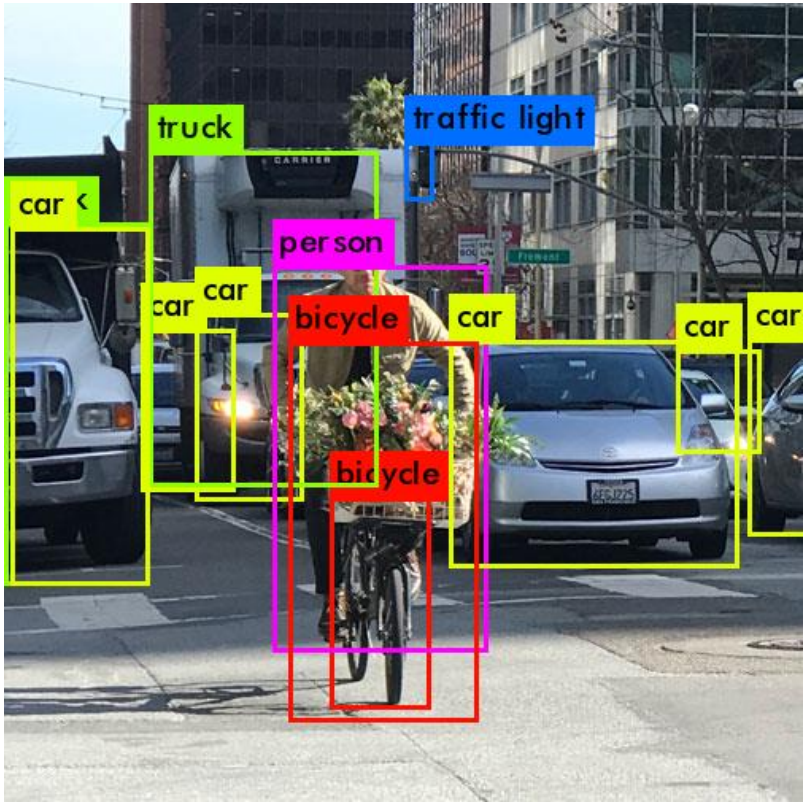

Source: Papers with Code

Source: Amazon.com

Machine Vision

Machine vision (MV) is the technology and methods used to provide automatic image-based inspection and analysis for applications such as:

- Automatic hardware inspections
- Process control
- Robot guidance
- Text recognition
- Facial recognition
- Object recognition

Machine Vision

High-end systems and platforms

- Stand-alone, vision only systems
- Executing on vision optimized processors
- Expensive

Mid-Range systems and platforms

- Hybrid system doing both vision and I/O tasks
- General purpose application processor
- Low to Mid-range costs
- i.e. Raspberry Pi

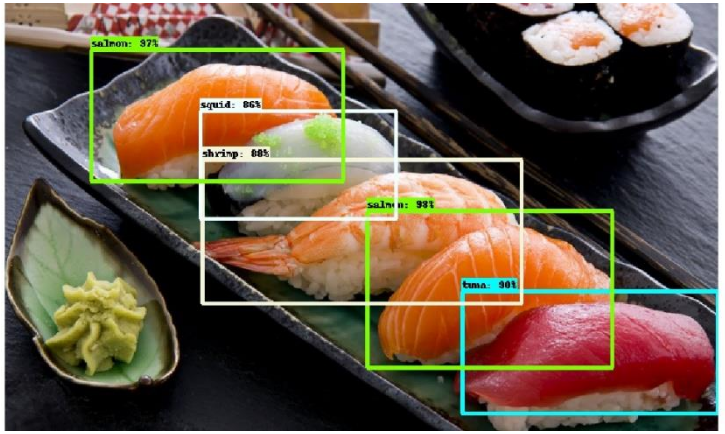

Machine Vision

Image Source:
Amazon.com

Image Source:
Kwiksource.com

Image Source:
Safewise

Machine Vision

Image Source:
Amazon.com

Image Source:
Kwiksource.com

Image Source:
Safewise

Machine Vision

Source: aimee.io

Machine Vision Architectures

Machine Vision Architectures

Google's Vision API

What can an online API do?

- Insights from images
  - 2 Cats
  - 1 Dog
  - 1 Squirrel
- Extract text
  - Optical Character Recognition (OCR)
- Face detection
- Content moderation
  - SafeSearch

Intelligence in the Cloud

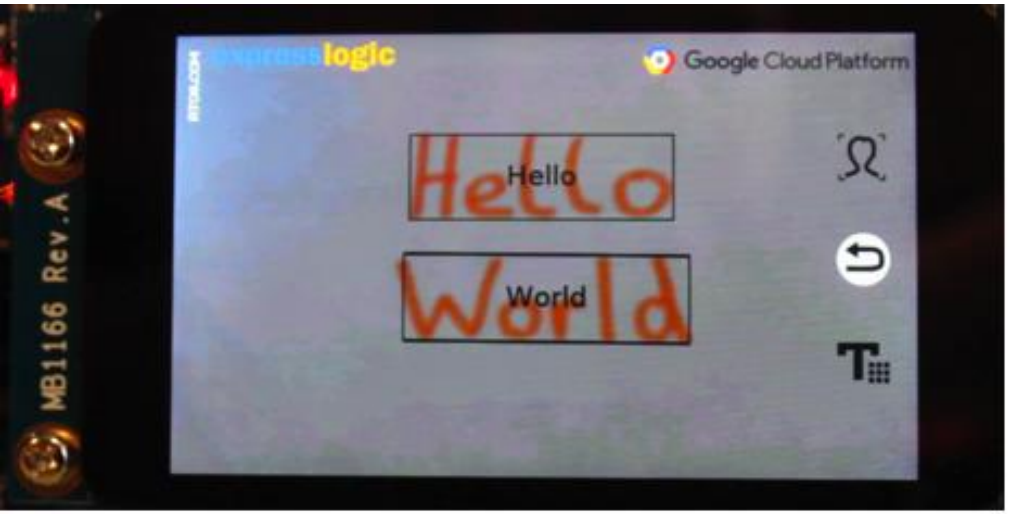

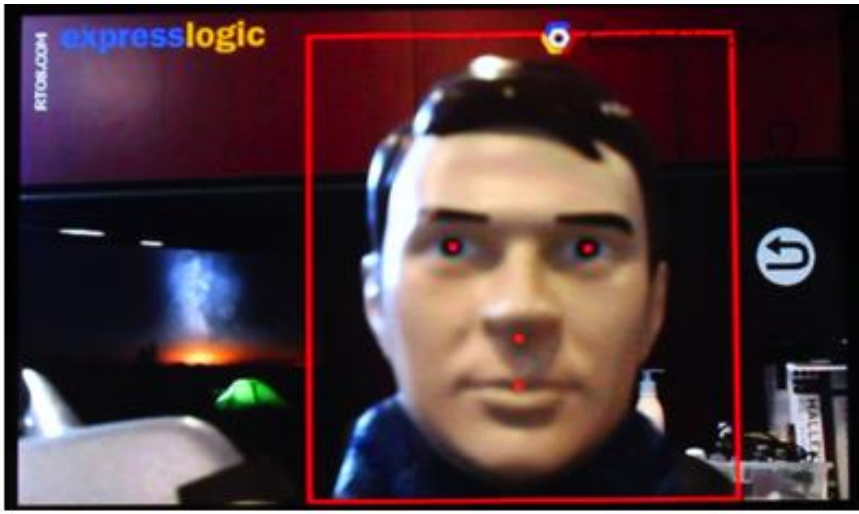


```
Output:                                                    Log file: Off
DHCP client running...
IP address: 10.0.0.221
Network mask: 255.255.255.0
Google Vision API is ready
DNS Lookup for GOOGLE APIs Domain:vision.googleapis.com
GOOGLE APIs Domain IPv4 address: 172.217.0.10
TCP session established.
TLS session established.
Text[(258, 111), (541, 111), (541, 203), (258, 203)]: Hello
Text[(252, 240), (558, 242), (558, 336), (252, 334)]: World
```

Vision in the MCU

What do you need to do machine learning at the edge?

- Hardware Floating Unit (FPU)
- ML Libraries
- Enough CPU cycles
- Training Dataset
  - 5,000 labeled examples per category for acceptable performance
  - 10,000,000 labeled examples to achieve human performance
- Time and patience

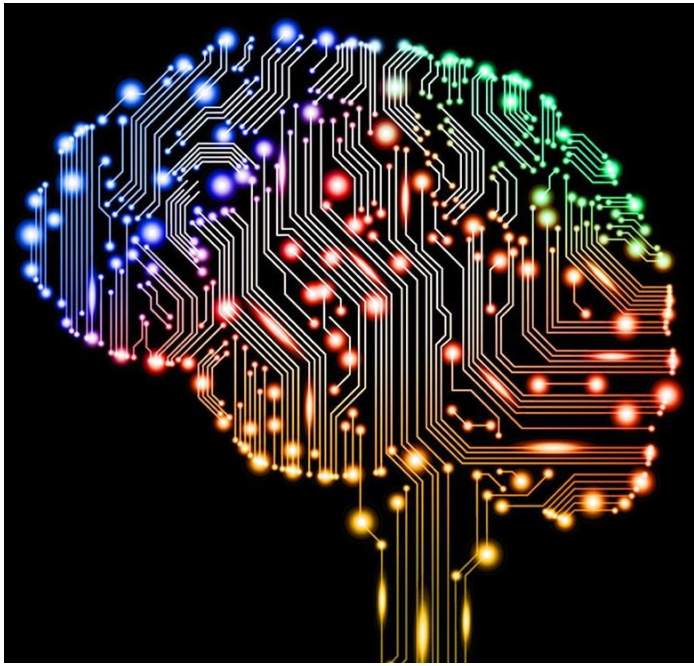

Vision in the MCU

OpenMV Cam with a Cortex-M7

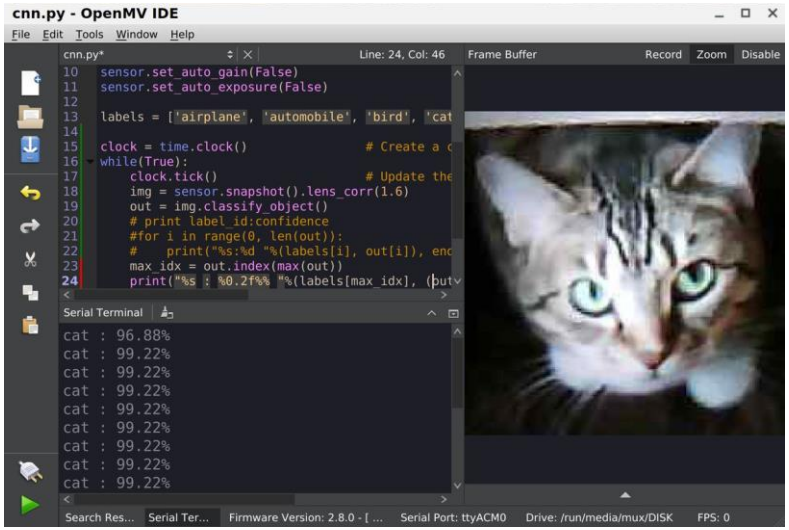

Video :
https://www.youtube.com/watch?v=PdWi_fvY9Og

Speech Recognition – Keyword Spotting

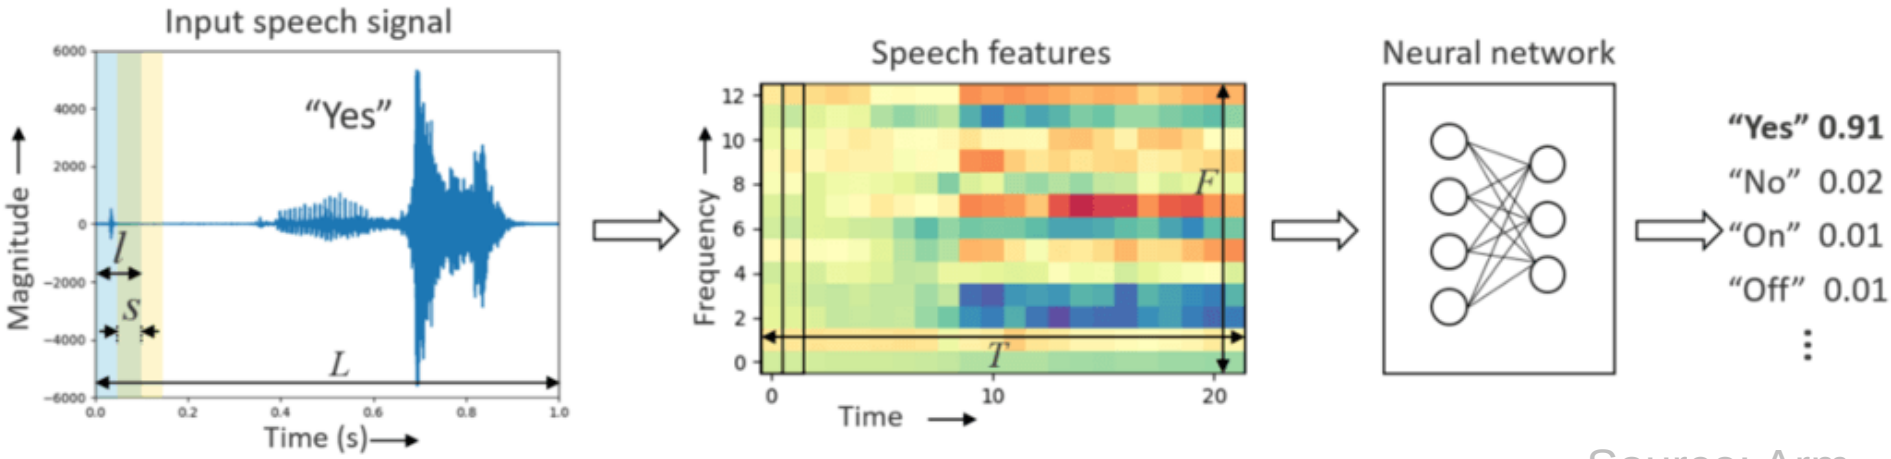

Source: Arm

Speech Recognition – Keyword Spotting

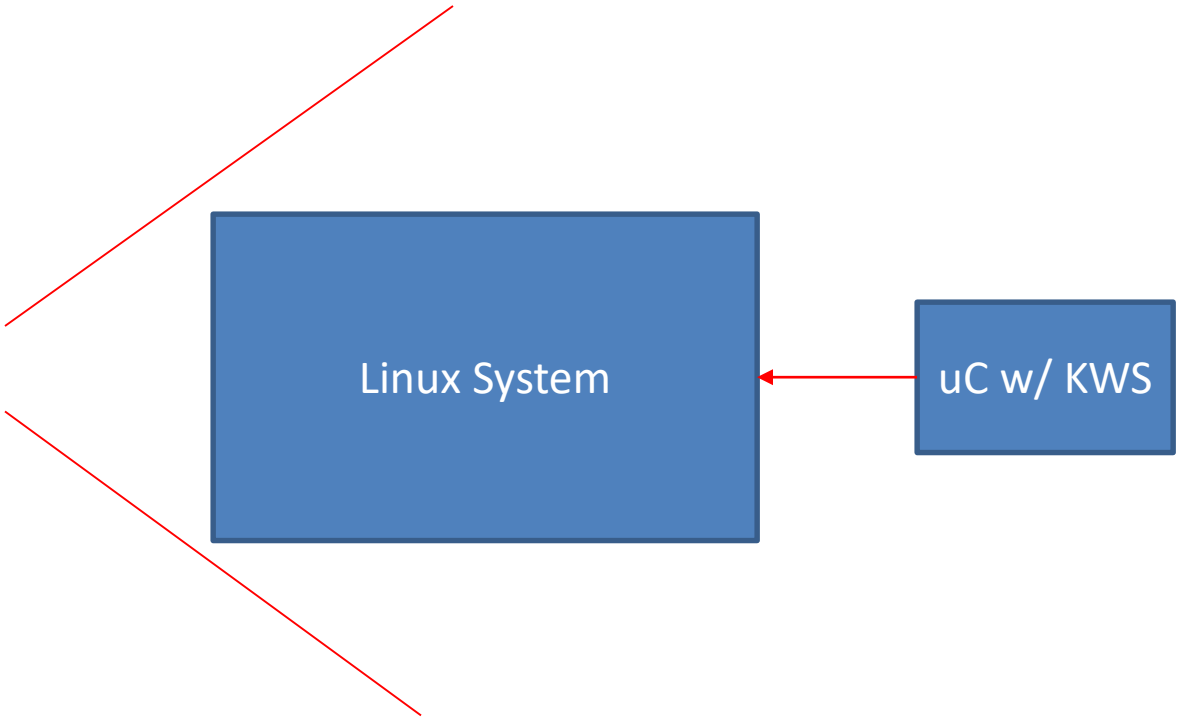

Source: Amazon.com

Speech Recognition – Keyword Spotting

Networks used:

- Deep Neural Networks (DNN)
- Convolution Neural Network (CNN)
- Recurrent Neural Network (RNN)
- Convolution Recurrent Network (CRNN)
- Depthwise Separable Convolution Network (DS-CNN)

https://community.arm.com/developer/ip-products/processors/b/processors-ip-blog/posts/high-accuracy-keyword-spotting-on-cortex-m-processors

Machine Learning Best Practices

1. Don’t do it yourself! Leverage an existing platform to accelerate development
2. For resource constrained, connected systems, off-load the vision to the cloud
3. Price shouldn’t be the only factor considered
4. Make sure you are using the right data
5. Try multiple tools to see which one best fits your application and team
6. Use 80% of your images for training and the last 20% for validating the algorithms
7. To be successful, understand the fundamental machine learning algorithms
8. Leverage a safety and security certified RTOS to minimize the security threats
9. Explore CMSIS-NN and the white papers that surround it
10. Start early, don’t wait to the last minute to learn how machine vision works.

Additional Resources

- Download Course Material for
  - C/C++ Doxygen Templates
  - Example source code
  - Blog
  - YouTube Videos
- Embedded Bytes Newsletter
  - http://bit.ly/1BAHYXm

From www.beningo.com under

- Blog > CEC – Machine Learning for Embedded Software Engineers