

Machine Learning for Embedded Software Engineers

Class 4: Machine Vision with OpenMV

April 25, 2019
Jacob Beningo

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

- Introduction to OpenMV
- The OpenMV hardware
- An example
 - Smile detection
 - Training with Caffe
 - Converting the model
 - Deploying the model




Presented by:

OpenMV


OpenMV is a project that brings machine vision into a simple Arduino like environment for “simple” applications.

- Can write scripts in Python
- Expandable I/O for interfacing
- Built-in camera module
- Expandable camera options
- Custom IDE for developing and deploying scripts




The OpenMV Hardware

- Arm Cortex-M7 Processor
 - STM32H743VI
 - 400 MHz
 - 1MB RAM
 - 2 MB Flash
- OV7725 image sensor
 - 640x480 16-bit RGB565 @ 60 FPS
 - 640x480 8-bit Grayscale images
- 2.8 mm lens



OpenMV Expansion Hardware

LCD



Wi-Fi



Servo




Pan and Tilt



Motor



Flit Camera Adapter



Getting Started with OpenMV

- Use a screw driver to remove the two lens mount screw from the lens mount on your OpenMV Cam.
- Apply some isopropyl alcohol to a small part of the cloth.
- Rub the wet part of the cloth on the camera IC gently. Any dirt spots on the camera IC will be microscopic to the human eye so just try to generally rub down the top of the camera IC. Note that the top of the camera IC is glass.
- After cleaning the camera IC make sure the alcohol has evaporated completely and that no cloth strands were left behind. Note that we're using isopropyl alcohol versus water since isopropyl alcohol evaporates quickly and doesn't leave anything behind.
- Use the screw driver to re-attach the lens mount. Make sure that the set screw on the lens mount points off the top/back of the OpenMV Cam.

Source: docs.openmv.io

The OpenMV IDE

The screenshot displays the OpenMV IDE interface. The main window shows a Python script named 'helloworld 1.py' with the following code:

```
1 import sensor, image, time
2
3 # Setup Camera
4 sensor.reset()
5 sensor.set_pixformat(sensor.RGB565)
6 sensor.set_framesize(sensor.QQVGA)
7 sensor.skip_frames(10)
8 threshold = (10, 90, -80, -30, 20, 50)
9 clock = time.clock()
10
11 # Find blobs
12 while(True):
13     clock.tick()
14     img = sensor.snapshot()
15     for b in img.find_blobs([threshold]):
16         img.draw_rectangle(b[0:4])
17         print("====\nBlob %s" % str(b))
18     print("FPS %d" % clock.fps())
19
```

The right side of the IDE shows a live camera feed with several colored blobs (red, blue, green) detected and outlined in white. Below the camera feed is a histogram of the LAB color space, showing the distribution of colors in the image. The histogram is divided into three sections: L (Lightness), A (Alpha), and B (Beta). The L section shows a peak around 88, the A section shows a peak around -17, and the B section shows a peak around -7. The histogram also displays various statistical values for each section.

The bottom of the IDE shows the Serial Terminal output, which matches the print statements in the code:

```
====
Blob (88, 33, 46, 47, 833, 115, 50, 0.898667, 1, 1)
FPS 15
====
Blob (24, 27, 34, 28, 413, 37, 37, 2.638215, 1, 1)
====
Blob (88, 33, 47, 48, 842, 116, 50, 0.9409514, 1, 1)
FPS 15
====
Blob (24, 27, 34, 28, 408, 37, 37, 2.625929, 1, 1)
====
Blob (88, 33, 47, 48, 839, 115, 50, 0.9321314, 1, 1)
FPS 15
====
Blob (24, 27, 34, 28, 410, 37, 37, 2.628374, 1, 1)
====
Blob (88, 33, 46, 47, 854, 116, 50, 0.9656991, 1, 1)
FPS 15
```

The bottom status bar indicates the Firmware Version: 2.0.0 - [latest], Serial Port: COM5, Drive: I:/, and FPS: 15.2.

Smile Recognition

- 1) Download the smile dataset
- 2) Train a model using Caffe
- 3) Quantize the model
- 4) Convert the model to a binary
- 5) Deploy to the OpenMV module
- 6) Run the model

Download the Smile Dataset

<https://github.com/hromi/SMILEsmileD>

Code Issues 0 Pull requests 0 Projects 0 Wiki Insights

open source smile detector haarcascade and associated positive & negative image datasets

4 commits 1 branch 0 releases 1 contributor

Branch: master New pull request Create new file Upload files Find File Clone or download

hromi added sorter miniapp for manual sorting&labeling of positive&negative... Latest commit 6afac66 on Oct 28, 2010

SMILEs	say hello to the world of smiles :)	9 years ago
appz	added sorter miniapp for manual sorting&labeling of positive&negative...	9 years ago
smileD	say hello to the world of smiles :)	9 years ago
README	added sorter miniapp for manual sorting&labeling of positive&negative...	9 years ago

README

SMILEsmileD - smile detector open source project v0.5

SMILEs - contains positive and negative sample sets with associated idx files
smileD - XML cascades for OpenCV's cvHaarDetectObjects function
appz - ad hoc applications useful for extending of SMILEsample and smileDetector testing

The Smile Dataset

- Contains
 - 3,000 positive images
 - 9,000 negative images
- The skewed dataset could result in bias!

Augment the images using the following script:

https://github.com/openmv/openmv/blob/master/tools/augment_images.py

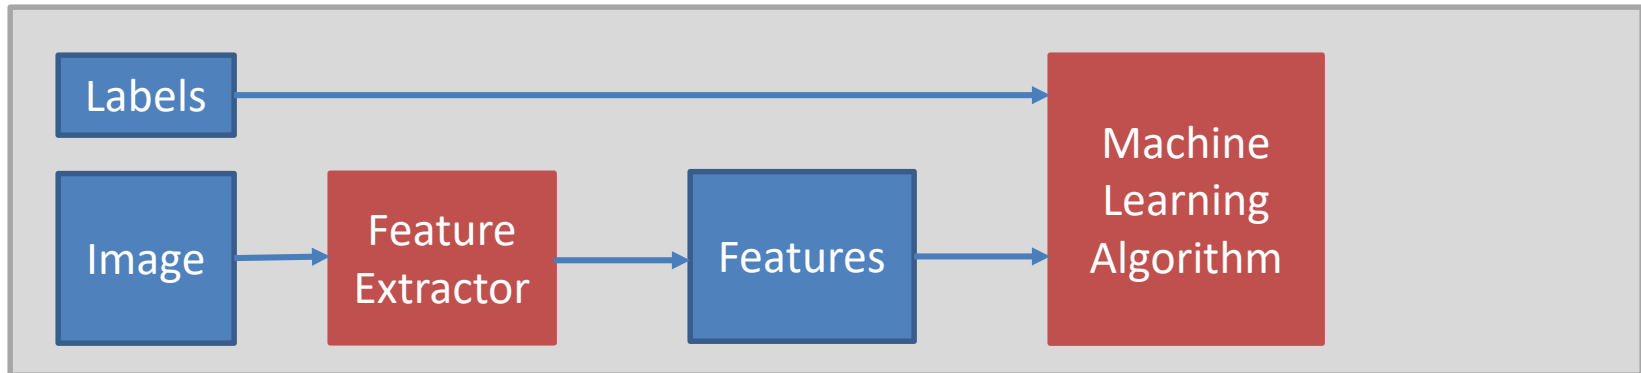
The Smile Dataset

- Download the script or copy it to a file
- Use the following to run the script:

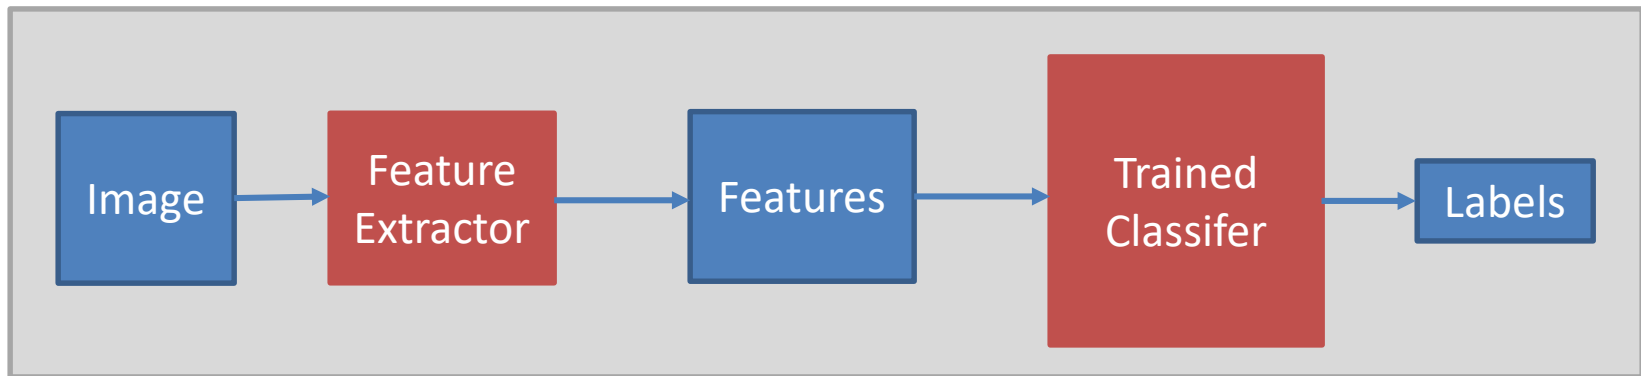
```
Python augment_images.py -input /Smiles_POSITIVE_DIR -  
output /Smiles_POSITIVE_2 -count 4
```

Training using Caffe

Training Phase



Prediction Phase



Training using Caffe

Caffe is a deep learning framework created with expression, speed and modularity in mind.

- Developed by Berkeley AI Research
- Download Caffe from <https://caffe.berkeleyvision.org/>

4 Steps to Training a CNN

- 1) Prepare the data (convert to Caffe friendly format)
- 2) Define the model (create config file, .prototxt)
- 3) Define the solver (solver parameters, .prototxt)
- 4) Model training (execute the model)

Quantizing the Model

After model training, need to quantize the model

- Convert weights and activations from floating point to fixed point
- Arm provides a script for Caffe at https://github.com/ARM-software/ML-examples/blob/master/cmsisnn-cifar10/nn_quantizer.py
- Run the script:

```
python2 nn_quantizer.py --model  
models/smile/smile_train_test.prototxt --weights  
models/smile/smile_iter_*.caffemodel --save  
models/smile/smile.pkl
```

Converting the Model to Binary

Need to output code for each layer in the NN with dimensions and weights.

Download the NN Converter Script:

- https://github.com/openmv/openmv/blob/master/ml/cmsisnn/nn_convert.py

```
python2 nn_convert.py --model models/smile/smile.pkl --mean  
/path/to/mean.binaryproto --output smile.network
```


Deploying on the OpenMV

Start by loading the NN into memory:

```
# Load Smile Detection network
```

```
net = nn.load('/smile.network')
```

```
# Load Face Detection Haar Cascade
```

```
face_cascade = image.HaarCascade("frontalface", stages=25)
```

```
print(face_cascade)
```

Deploying on the OpenMV

Next, capture an image and find the faces:

```
# Capture snapshot
```

```
img = sensor.snapshot()
```

```
# Find faces.
```

```
objects = img.find_features(face_cascade, threshold=0.75,  
scale_factor=1.25)
```

Deploying on the OpenMV

Pass each image to the neural network:

for r in objects:

```
# Resize and center detection area
```


```
r = [r[0]+10, r[1]+25, int(r[2]*0.70), int(r[2]*0.70)]
```

```
out = net.forward(img, roi=r, softmax=True)
```

```
img.draw_string(r[0], r[1], ':)') if (out[0] > 0.8) else ':(', color=0,  
scale=2)
```

The Results

```
smile_detection.py - OpenMV IDE
File Edit Tools Window Help
smile_detection.py* Line: 36, Col: 23 Frame Buffer Record Zoom Disable
8 sensor.skip_frames(time=1000)
9 sensor.set_auto_gain(False)
10 sensor.set_auto_exposure(False)
11
12 # Load smile detection network
13 net = nn.load('/smile.network')
14
15 # Load Face Haar Cascade
16 face_cascade = image.HaarCascade("frontalface", stages=20)
17
18 # FPS clock
19 clock = time.clock()
20 while (True):
21     clock.tick()
22
23     # Capture snapshot
24     img = sensor.snapshot().lens_corr(2)
25
26     # Find faces.
27     objects = img.find_features(face_cascade, threshold=0.75,
28
29     # Detect smiles
30     for r in objects:
31         img.draw_rectangle(r)
32         out = net.forward(img, roi=r, softmax=True)
33         img.draw_string(r[0], r[1]+r[3]-20, ':)') if out.index
34
35     # Print FPS.
36     print(clock.fps())
```



References and Resources

- <https://openmv.io/pages/download>
- <https://openmv.io/blogs/news/deep-learning-on-a-cortex-m7-camera-3ma-deep-learning>
- <http://adilmoujahid.com/posts/2016/06/introduction-deep-learning-python-caffe/>

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