

ECE 105: Introduction to Electrical Engineering

Lecture 17



Intro to hardware ML

Yasser Khan

Rehan Kapadia

Why use NN when you can hardcode?



Deterministic



Temperature
measurement is
absolute

Probabilistic

"stop"



Voice can
vary

Self-driving car

Under the bonnet

How a self-driving car works

Signals from **GPS (global positioning system)** satellites are combined with readings from tachometers, altimeters and gyroscopes to provide more accurate positioning than is possible with GPS alone

Lidar (light detection and ranging) sensors bounce pulses of light off the surroundings. These are analysed to identify lane markings and the edges of roads

Video cameras detect traffic lights, read road signs, keep track of the position of other vehicles and look out for pedestrians and obstacles on the road


Radar sensor

Ultrasonic sensors may be used to measure the position of objects very close to the vehicle, such as curbs and other vehicles when parking

The information from all of the sensors is analysed by a **central computer** that manipulates the steering, accelerator and brakes. Its software must understand the rules of the road, both formal and informal

Radar sensors monitor the position of other vehicles nearby. Such sensors are already used in adaptive cruise-control systems

Source: *The Economist*



On-device machine learning applications in the single mW and below



Vibration and motion

Any 'signal'

Predictive maintenance, sensor fusion, accelerometer, pressure, lidar/radar, speed, shock, vibration, pollution, density, viscosity, etc.



Voice and sound

Recognition and creation

Keyword spotting, speech recognition, natural language processing, speech synthesis, sound recognition, etc.



Vision

Images and video


Object detection, face unlock, object classification etc.

Sub-mW computing vs 100s of mW computing



Single Board Computer

- More powerful (faster processor, more memory)
- Runs full, general purpose operating system (OS)
- Can provide full command line or graphical user interface
- Requires more power






Microcontroller


- Less powerful
- Bare-metal (superloop) or real-time operating system (RTOS)
- Limited or no user interface
- Requires less power

Cortex-M series microcontrollers are becoming powerful

nRF52840 is built around the 32-bit ARM® Cortex™-M4 CPU with floating point unit running at 64 MHz.



ML in microcontrollers TinyML




Motion classification from accelerometer data




Acceleration (m/s²)

x	-1.4
y	0.4
z	9.6



class1
↓
class4


Motion classification from accelerometer data




Way too many inputs to the model

Problems with deep learning

1. Computational complexity
2. Requires lots of training data



Feature extraction




125 x 3


375 → 3 inputs

What is a feature


- Individual measurable property or characteristic of a phenomenon being observed




1 (x, y, z) accelerometer point
from “left-right” sample




Many (x, y, z) accelerometer
points from all classes




RMS of acceleration in each axis as the feature





Features don't need to be only in the time domain



Features in the frequency domain




Feature
extraction:
FFT




Feature
extraction:
PSD


Using both time and frequency domain features




Holdout Method for training, validation, and testing




Dataset balance




Data flow diagram






Inference: using the machine learning model to make predictions on unseen data in the wild


Final implementation




Final implementation




Final implementation




Final implementation




Final implementation




Confusion matrix in classification




Confusion matrix in classification




Confusion matrix in classification



Confusion matrix in classification



Confusion matrix in classification



Confusion matrix accuracy

Total accuracy: $\frac{\sum correct}{\sum all} = \frac{813}{1000} = 0.813$

Confusion Matrix

		Predicted Label			
Actual Label		Circle	Idle	Left-Right	Up-Down
	Circle	205	10	1	46
	Idle	6	199	0	32
	Left-Right	9	17	223	34
	Up-Down	21	8	3	186

Is this a good dataset?

Total accuracy: $\frac{\sum correct}{\sum all} = \frac{973}{1000} = 0.973$

Confusion Matrix

		Predicted Label			
		Circle	Idle	Left-Right	Up-Down
Actual Label	Circle	0	6	0	0
	Idle	0	973	0	0
	Left-Right	0	11	0	0
	Up-Down	0	10	0	0

Positives and negatives

Confusion Matrix


		Predicted Label	
		Positive	Negative
Actual Label	Positive	38	17
	Negative	3	42

- ✓ True Positive (TP): Predicted positive matches actual positive
- ✓ True Negative (TN): Predicted negative matches actual negative
- ✓ False Positive (FP) ("Type I Error"): Predicted positive does not match actual negative
- ✓ False Negative (FN) ("Type II Error"): Predicted negative does not match actual positive

False positives, false negatives

Confusion Matrix

Actual Label	Predicted Label			
	Circle	Idle	Left-Right	Up-Down
Circle	205	10	1	46
Idle	6	199	0	32
Left-Right	9	17	223	34
Up-Down	21	8	3	186



True Positive (TP): Predicted positive matches actual positive

True Negative (TN): Predicted negative matches actual negative

False Positive (FP) ("Type I Error"): Predicted positive does not match actual negative

False Negative (FN) ("Type II Error"): Predicted negative does not match actual positive

Accuracy for a single class

Confusion Matrix

Actual Label		Predicted Label			
		Circle	Idle	Left-Right	Up-Down
	Circle	205	10	1	46
	Idle	6	199	0	32
	Left-Right	9	17	223	34
Up-Down	21	8	3	186	

Accuracy

$$\begin{aligned}
 ACC &= \frac{TP + TN}{P + N} = \frac{TP + TN}{TP + TN + FP + FN} = \frac{\text{Green Box} + \text{Blue Box}}{\text{Total}} \\
 &= \frac{199 + 728}{1000} = 0.927
 \end{aligned}$$

True Positive Rate, Sensitivity

		Predicted Label				
		Circle	Idle	Left-Right	Up-Down	
Confusion Matrix	Actual Label	Circle	205	10	1	46
		Idle	6	199	0	32
		Left-Right	9	17	223	34
		Up-Down	21	8	3	186

True Positive Rate (TPR), Sensitivity, Recall, Hit Rate

$$\begin{aligned} TPR &= \frac{TP}{P} = \frac{TP}{TP + FN} = \frac{\text{Green Box}}{\text{Green Box} + \text{Purple Box}} \\ &= \frac{199}{199 + 38} = 0.840 \end{aligned}$$

True Negative Rate, Selectivity

		Predicted Label				
		Circle	Idle	Left-Right	Up-Down	
Confusion Matrix	Actual Label	Circle	205	10	1	46
		Idle	6	199	0	32
		Left-Right	9	17	223	34
		Up-Down	21	8	3	186

True Negative Rate (TNR), Specificity, Selectivity

$$\begin{aligned} TNR &= \frac{TN}{N} = \frac{TN}{TN + FP} = \frac{\boxed{\text{blue}}}{\boxed{\text{blue}} + \boxed{\text{orange}}} \\ &= \frac{728}{728 + 35} = 0.954 \end{aligned}$$

Precision (Positive Predictive Value)

Confusion Matrix

		Predicted Label			
		Circle	Idle	Left-Right	Up-Down
Actual Label	Circle	205	10	1	46
	Idle	6	199	0	32
	Left-Right	9	17	223	34
	Up-Down	21	8	3	186

Positive Predictive Value (PPV), Precision

$$\begin{aligned}
 PPV &= \frac{TP}{TP + FP} = \frac{\text{Green Box}}{\text{Green Box} + \text{Orange Box}} \\
 &= \frac{199}{199 + 35} = 0.850
 \end{aligned}$$

Predicted Label

Confusion Matrix	Actual Label		Circle	Idle	Left-Right	Up-Down
		Circle	205	10	1	46
		Idle	6	199	0	32
		Left-Right	9	17	223	34
		Up-Down	21	8	3	186

F1 Score

$$\begin{aligned}
 F_1 &= 2 \cdot \frac{PPV \cdot TPR}{PPV + TPR} = \frac{2TP}{2TP + FP + FN} \\
 &= 2 \cdot \frac{0.850 \cdot 0.840}{0.850 + 0.840} = 0.845
 \end{aligned}$$

		Predicted Label				
		Circle	Idle	Left-Right	Up-Down	
Confusion Matrix	Actual Label	Circle	205	10	1	46
	Idle	6	199	0	32	
	Left-Right	9	17	223	34	
	Up-Down	21	8	3	186	
Per-class accuracy		0.907	0.927	0.936	0.856	
F1 scores		0.815	0.845	0.875	0.721	

Total accuracy: 0.813


F1 average: 0.818

F1 score is indicative of dataset balance


Confusion Matrix		Predicted Label				
			Circle	Idle	Left-Right	Up-Down
		Circle	0	6	0	0
		Idle	0	973	0	0
		Left-Right	0	11	0	0
		Up-Down	0	10	0	0
Actual Label						
Per-class accuracy		0.994	0.973	0.989	0.99	
F1 scores		0.0	0.986	0.0	0.0	

Total accuracy: 0.973


F1 average: 0.247



Underfit: Model fails to capture trends in the data




Good fit: Model captures trends and can generalize to unseen data




Overfit: Model captures training data trends but fails on unseen data


Accuracy vs epoch to understand dataset



Underfit: Model performs poorly on training and validation data




Good fit: Model generalizes well from training to validation data




Overfit: Model predicts training data well but fails to generalize to validation data


Loss vs epoch to understand dataset



Underfit: Model performs poorly on training and validation data




Good fit: Model generalizes well from training to validation data



Overfit: Model predicts training data well but fails to generalize to validation data

Fixing underfit




Underfit: Model performs poorly on training and validation data

- Get more data
- Try different features or more features
- Train for longer
- Try a more complex model (more layers, more nodes, etc.)


Fixing overfit

- Get more data
- Early stopping
- Reduce model complexity
- Add regularization terms
- Add dropout layers (for neural networks)




Overfit: Model predicts training data well but fails to generalize to validation data


Model characterization in realtime



Output of the model



62.5 Hz sampling for 2 seconds
with 3 axes = 375 values



P(left-right) = 0.9143

P(up-down) = 0.0032

P(circle) = 0.0581

P(idle) = 0.0244

Action after classification

```
if (p_left_right > 0.5) {  
    // Do stuff  
}  
  
if (p_up_down > 0.5) {  
    // Do some things  
}  
  
if (p_circle > 0.8) {  
    // And now for something completely different  
}
```