

ML with PyTorch

Weeks 1 & 2

Intro

- Course uses python (hence Py) several versions avail:
 - Py: 2.7, 3.5, 3.6, 3.7
 - C++: 11 (according to an unofficial source)
- Course caveats:
 - No GPU support
 - Simulated (a.k.a. clean) data

Week 1

- Mostly just reviews how to use tensors
- Some things to note:
 - You may import numpy arrays as PyTorch tensors (and vice versa)
 - You may add constants to tensors:
 - Equivalent to adding constant to each element of tensor
 - Multiplication of tensors by default is element-wise (!)
 - Specify matrix multiplication using `torch.mm(a, b)`
 - Built-in functions to get mean, max, etc

Week 1 Lab

- Review image manipulation with PyTorch:
 - Import
 - Scale
 - Flip

Week 2

- Linear Regression
 - Concepts of cost, loss, epochs, etc
- Concepts *not* covered (courtesy of Wojtek):
 - Early stopping and patience
 - Early stopping: when you stop training... early
 - Patience: how many epochs must pass before stopping/adjusting
- How to train and validate using PyTorch

Jupyter

- Allows for access in web browser
- Allows for code to be partitioned and loaded
 - Memory isn't deallocated
 - Can load function and 'toy' around with results
- Used in labs (and Wojtek, and industry, apparantly)
 - Love/hate