2025

4Geeks Academy: data science cohort 12

DAY 28: DEEP LEARNING CONT'D

TODO

DEEP LEARNING

Convolutional and recurrent neural networks, where to get free GPU access

IMAGE CLASSIFICATION PROJECT

Work on image-classifier-project-tutorial (Intro to Deep Learning module), plan to finish by next Wednesday

TOPICS

01 RECURRENT NEURAL NETWORKS

02 CONVOLUTIONAL NEURAL NETWORKS

03 KAGGLE: GPU COMPUTE

RECURRENT NEURAL NETWORKS

WHAT Neural network that sends its own output back as one of the inputs

WHY
- Gives the neuron a 'memory' for the 1D order of data
- Modeling sequence data: time-series, text, etc.

HOW The output from the last set of inputs is treated like an input for the next set of inputs

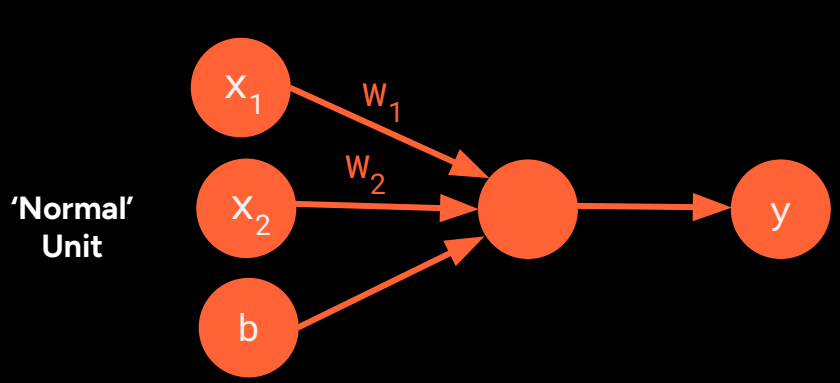

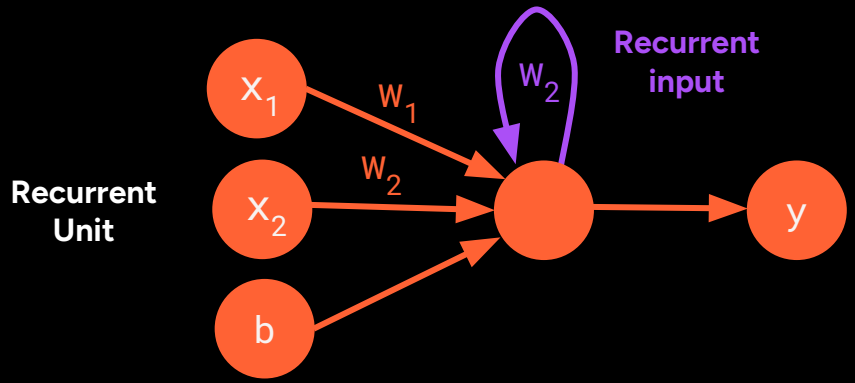

CONVOLUTIONAL NEURAL NETWORKS

WHAT Neural network that has its 'weights' arranged in 2 dimensions

WHY
- Gives the network 'awareness' of the 2D relationships in data
- Modeling spatial data: images, maps, etc.

HOW Uses convolutional filters (kernels) as 2D matrices of weights

Hyperparameters:

- The number of filters
- The size of the filters
- The filtering 'stride'

Pooling by averaging or taking the max value over an area is often applied between layer too.

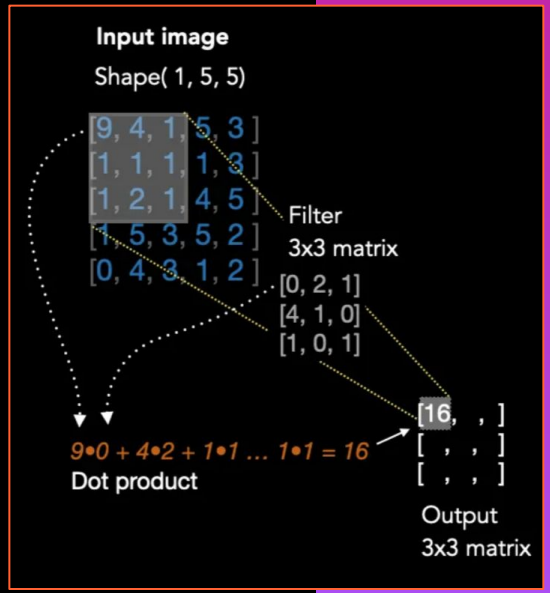

KAGGLE: GPU COMPUTE

WHAT

- Kaggle is a data science and machine learning platform
- They host competitions, datasets and offer a notebook compute environment
- Free tier account gets 30 hours of GPU compute per week

WHY

- GPU hardware is specifically designed for the kind of matrix math used in neural networks
- Kaggle notebooks also have more CPU/memory resources than our codespaces

DOWN SIDE

- Less flexibility than codespace
- No access to underlying OS
- Multi-notebook, multi-module projects a pain