





Introduction to Convolutional Neural Networks

Mehmet Kerem Turkcan


Vision Models

- **Fundamental Idea:**
 - There is a repeating structure in visual systems
 - So patches of the input to the retina are processed in the same way
 - How do we structure our thinking about neural networks to enable this?




Convolutional Neural Networks


Feedforward View of a Neural Network




Convolutional View of a Neural Network



Convolutional Neural Networks




Why do this? So that we can move the kernel around a larger input!



Channels/Depth

- **Fundamental Idea:**


- Images have *channels* (Red, Green, Blue)
- If we have K filters in a layer, then we will have K output *channels*
- We also call the number of channels in a representation its *depth*



This example has 3 input channels, and 1 output channel (since we only have 1 output unit)


Channels/Depth

- **Fundamental Idea:**
 - Images have *channels* (Red, Green, Blue)
 - If we have K filters in a layer, then we will have K output *channels*
 - We also call the number of channels in a representation its *depth*



This example has 1 input channel, and 2 output channels (since we have 2 output units)


Channels/Depth



Pooling

- **Fundamental Idea:**


- Modern images are pretty large! (3840x2160)
- We introduce “pooling” as a special form of a neural network “layer”, to reduce the size of the input arrays/matrices between one layer to another.




Flattening

- **Fundamental Idea:**

- When we classify a whole image, we no longer need the structure of the image
- Similar to the strawberry/blueberry example, we can flatten the output of a convolutional layer so that a neural network can end like a feedforward neural network



A whole Convolutional Neural Network



A whole Convolutional Neural Network

