

## Vision (just a little bit)

- Given a megapixels-big RGB image with one input neuron per pixel and a fully connected design there would be trillions of weights to work with
- And Locality matters
  - Neighbouring pixels contribute to real features, distant ones don't
  - But a neuron's position within a layer makes no difference
  - Might as well just present the pixels in random order – can't be right
- Also Spatial Invariance
  - A cat looks the same regardless of where it appears in an image
  - Cats should be recognized the same way regardless of position
  - So we expect there to be groups of neurons all with the same weights
- Convolutional Neural Networks
  - From studies of the Visual Cortex; the Receptive Field of a neuron
  - Groups of neurons all connected to the same neighbouring pixels
  - Those groups, Kernels, all have the same set of weights
  - Kernels of size  $k$ , with a Stride of  $s$
  - The operation of such a layer can be treated as a Matrix operation
  - GPUs are good for them

## Memory

- Recurrent neural networks
- There can be loops back in the connections
  - A delay is required at each step
- A neuron's new state can depend in some way on its previous state
  - And those of its neighbours too
- Just like building a flip-flop out of nand gates
- Can analyse sequential data