

Data Science Toolbox Question Sheet

09.1 Neural Networks

Daniel Lawson

Block 9

1. What is an activation function? Could $f(x) = x$ be used and why?
2. Given a simple neuron which fires when $f(x > 0) = 1$ and is zero otherwise, can the following logical functions be achieved in one layer? If so, give an example. AND, OR, NOT, XOR, ...
3. A colleague claims that because even single layer neural networks are universal function approximators, the architecture isn't really important. Critique this claim.
4. A colleague claims that overfitting in neural networks is no worse a problem than in a linear model. Critique this claim.
5. Define penalisation using L1 and L2 norms on the parameters.
6. Consider the softmax equation $p(c_i) = \exp(z_i) / \sum_j \exp(z_j)$, where c_i is the output of the final layer and z_i is the input. In what sense are these probabilities for Bayesians, frequentists and in practice?
7. Explain feed-forward, convolutional, recurrent and recursive neural networks at a high level.
8. What is an auto-encoder and how are they useful? When are they a good idea to use?