Deep Learning
Tutorial 1: Signals and Convolution
(unassessed)

January 5, 2021

Q1: The curse of dimensionality

1a) For a one-dimensional space of real numbers between 0 and 1, 100 observations are required to adequately cover this space such that histograms can be calculated and conclusions can be drawn. How many observations would be required to adequately cover the space for the following number of dimensions?

i) 10 dimensions

ii) 1000 dimensions

1b) Under what conditions is it still possible to perform Cluster analysis and outlier detection in high-dimensional spaces?

1c) Why is the curse of dimensionality a serious hurdle in machine learning problems?

1d) Why is using weight sharing common practice?

1	1	2	3	2	1	1
1	2	6	6	5	2	1
1	5	6	7	7	6	2
1	5	6	6	6	4	2
1	1	4	5	4	1	1
1	1	3	4	3	1	0
0	0	1	2	1	0	1

(a) Single channel image

0	1	0
0	-2	1
0	1	0

(b) 3x3 Filter

Figure 1: Perform a convolution

Q2: Convolutions

2a) Suppose you have the following single-channel image shown in Figure 1a:

i) Compute the result of max pooling of size 3x3 with stride=2.

ii) Compute the result of convolution of the input image with the 3x3 filter shown in Figure 1b using stride=2, no zero padding.

2bi) How would the filter shown in Figure 1b need to changed to obtain the correct definition of convolution rather than cross-correlation? Check your answer to **2aii)**.

2bii) Why do we need to perform this change?

2c) Which of the following properties hold for convolution? (True/False)

i) Non-Commutativity: $f * g \neq g * f$

ii) Associativity: $f * (g * h) = (f * g) * h$

iii) Non-Distributivity: $f * (g + h) \neq (f * g) + (f * h)$

iv) Associativity with scalar multiplication: $a(f * g) = (af) * g$

v) Derivative: $D(f * g) = (Df) * g = f * Dg$