DEEP LEARNING WEEK 10

1. Which of the following 3x3 kernels works best for the edge detection task in an image?

   a) $\begin{bmatrix} 0 & -1 & 4 \\ -1 & 7 & -1 \\ 0 & -1 & 0 \end{bmatrix}$

   b) $\begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$

   c) $\begin{bmatrix} 1 & 1 & 1 \\ 1 & -8 & 1 \\ 1 & 1 & 1 \end{bmatrix}$

   d) $\begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}$

   Answer: c)
   Solution: The result of convolution with the third kernel will give 0 when all neighbors of a pixel are the same and non zero number when the neighbors of a pixel are not the same i.e. in the edges.

2. What is the motivation behind using multiple filters in one Convolution layer?
   a) Reduced complexity of the network
   b) Reduced size of the convolved image
   c) Each filter captures some feature of the image separately
   d) Insufficient information

   Answer: c)

   Solution: Increasing the number of filters at each layer creates more trainable parameters and increases the image's dimensions. However, we believe that each filter learns to capture some important image aspects. This is analogous to feature engineering in classical machine learning.

3. We have the following two images. We observe that a neuron is in the 3rd convolutional layer of a CNN. Which of the following statements is always true?

Image 1
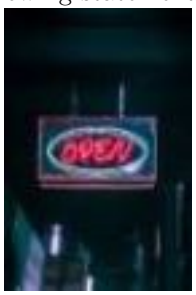

Image 2

   a)The neuron will fire again because both images contain the word 'OPEN'
   b)The neuron will not fire again because the color of the word 'OPEN' is different in both images
   c)The neuron will fire again since the background for images is dark.

d)None of these.

Answer: d)
Solution: Since we are at a very deep convolutional layer we can't say anything directly about what abstract representations of the image this neuron has learned.

4. What is the main innovation proposed in Resnet architecture?
   a)Use of 1x1 convolutions
   b)Residual connections
   c)Use of different-sized filters at a single layer
   d)Increasing the width of the network.

   Answer:b)

   Solution: Residual connections were the main innovation proposed in Resnet architecture to tackle the problem of vanishing/exploding gradients in deep CNN.

5. Consider a convolution operation with an input image of size 256x256x3 and 63 filters of size 11x11x3, using a stride of 4 and a padding of 2. What is the depth (h in w x b x h) of the output image?
   a) 63
   b) 64
   c) 40
   d) 3

   Answer: a)

   Solution: The height of the image is equal to the number of filters.

6. What is the result of convolving a 3x3 matrix with a 2x2 filter with stride 1 and no padding?
   a)2×2 matrix
   b)1×1 matrix
   c)4×4 matrix
   d)3×3 matrix

   Answer: a) A 2x2 matrix

   Solution: Size of output matrix is (W-H+2P)/S+1=3-2+1=2

7. What is the purpose of using a stride larger than 1 in convolutional neural networks?
   a) To increase the number of convolution operation in a layer
   b) To reduce the size of the output feature map
   c) To increase the number of trainable parameters
   d) To increase the receptive field of the convolutional layer

   Answer: b) To reduce the size of the output feature map

   Solution: Using a stride larger than 1 in convolutional neural networks reduces the size of the output feature map. This is because the stride determines the distance between each application of the filter to the input matrix. A larger stride means that the filter is applied less frequently, resulting in a smaller output matrix.

8. Which of the following layers doesn't contain learnable parameters in CNN?
   a) Max-pooling layers
   b) Convolutional layers
   c) Feedforward Network's hidden layers

d) Feedforward Network's output layer

Answer: a)

Solution: Max pool layers don't contain any learnable parameters. They simply produce the max value in their kernel window.

9. Which statement is true about the size of filters in CNNs?
   a) Larger filters capture more global features.
   b) Smaller filters capture more local features.
   c) The size of the filter does not affect the features it captures.
   d) The size of the filter only affects the computation time.

   Answer: a),b)

   Solution: Smaller filters capture more local features because they can detect smaller patterns in the input image. Larger filters, on the other hand, capture more global features because they can detect larger patterns in the input image

10. What is the purpose of guided backpropagation in CNNs?
    a) To visualize which pixels in an image are most important for a particular class prediction.
    b) To train the CNN to improve its accuracy on a given task.
    c) To reduce the size of the input images in order to speed up computation.
    d) None of the above.

    Answer: a)

    Explanation: Guided backpropagation is a technique used to visualize the parts of an input image that are most important for a particular class prediction. It achieves this by backpropagating the gradients of the output class concerning the input image but only allowing positive gradients to flow through the network