

Perceptron, Neural Nets and CNN

Your Name *

Sihan Wei

Which of the following are true about a typical MLP architecture? (Select all that apply)

1 point

- It contains one or more hidden layers between input and output.
- Each neuron performs a weighted sum followed by an activation function.
- Without non-linear activation functions, an MLP reduces to a linear model.
- All layers in an MLP must have the same number of neurons.
- MLPs can only be used for classification tasks, not regression.

Which of the following statements about the Perceptron algorithm are correct? (Select all that apply)

1 point

- It computes a weighted sum of inputs and applies a threshold to produce the output.
- It updates weights only when it misclassifies a sample.
- It uses gradient descent to minimize a loss function.
- It can classify non-linearly separable data like XOR.
- It is suitable for binary classification tasks.

Which of the following models can represent non-linear decision boundaries? (Select all that apply)

1 point

- Multi-layer perceptron with ReLU activation
- CNN with convolutional and ReLU layers
- Single-layer Perceptron
- Logistic regression
- Decision tree

In a CNN, which of the following affect the number of parameters in a convolutional layer? (Select all that apply)

1 point

- Kernel size
- Number of input channels
- Stride
- Number of filters
- Padding

Suppose your input is a 300 by 300 color (RGB) image, and you use a convolutional layer with 100 filters that are each 3x3. How many parameters does this hidden layer have (including the bias parameters)? 1 point

- 2700
- 2800
- 900
- 1000

Your input has shape 100 x 100 x 32, and you apply a convolutional layer with 64 filters of shape 5 x 5, stride 1, and padding 2. What is the shape of the output volume? 1 point

- 98 x 98 x 64
- 100 x 100 x 64
- 96 x 96 x 64
- 100 x 100 x 32

You apply a max-pooling layer with 2x2 window and stride 2 on a volume of size 64 x 64 x 32. What is the shape of the output volume? 1 point

- 32 x 32 x 32
- 63 x 63 x 32
- 31 x 31 x 32
- 32 x 32 x 16

What are the roles of activation functions in neural networks? (Select all that apply)

1 point

- Introduce non-linearity into the model
- Help networks learn complex functions
- Make the model interpretable
- Prevent overfitting by default

Which of the following are hyperparameters of the pooling layers? (Choose all that apply)

1 point

- Filter size
- Average weights
- Whether it is Max or Average
- Number of filters

You're working with an input volume of size $64 \times 64 \times 32$ (height \times width \times channels), and you're applying 40 filters of size 9×9 , with stride 1. If you want to use "same" padding (so that the output has the same spatial size as the input), how much padding should you use?

1 point

- 0
- 2
- 4
- 6
- 8

This content is neither created nor endorsed by Google.

Google Forms