

Study Exercises: Convolutional Neural Networks

1. Are fully-connected layers suitable for images? Discuss about it.
2. What is low-dimensional, and high dimensional data?
3. Describe the idea of a convolution.
4. Describe what is a kernel.
5. Considering a vector of size 1×10 , and a kernel of size 1×4 , what will be the size of the output vector after the 1D convolution?
6. What about the output size of a 2D convolution in which the input size is 6×6 and the kernel 3×3 ?
7. What about the output size of a 3D convolution in which the input size is $15 \times 15 \times 10$ and the kernel $3 \times 3 \times 3$?
8. What is the difference between convolution and matrix multiplication?
9. What is the primary purpose of using convolutional layers in a CNN?
10. To what structure of the multi-layer perceptron can we do a high level analogy with the kernel convolution in a CNN?
11. What should be the relationship between kernel depth and the input channels?
12. In a Convolutional Neural Network, why is it important to convolve the input with multiple kernels in different layers, and how does this contribute to the network's ability to understand and recognize features in the data?
13. What is a receptive field in CNNs?
14. How does the size of the receptive field change as you move deeper into a CNN architecture?
15. Explain the concepts of padding and stride
16. Explain the concept of weight sharing in CNNs.
17. What is the role of activation functions like ReLU in a convolutional layer?
18. How does the stride value affect the output size of a convolutional layer?
19. Explaining the concept of pooling.
20. In a CNN, what is the purpose of max-pooling, and how does it affect feature maps?
21. What is the purpose of the fully connected layers in a CNN?
22. How is backpropagation used to update the weights of a CNN during training?
23. Explain the term "vanishing gradient" and how it can be mitigated in CNNs.

24. What is data augmentation, and why is it beneficial when training CNNs?
25. How does dropout regularization work in CNNs, and when should it be used?
26. Define the term "transfer learning" and provide an example of when it can be advantageous.
27. In a convolutional layer, if the input image has a size of 128x128, and the filter size is 3x3 with a stride of 2, what is the size of the output feature map?
28. Explain the trade-off between computational complexity and feature richness in choosing the depth of a CNN.
29. What is the role of batch normalization in CNNs, and how does it help during training?
30. Define overfitting in the context of CNNs and discuss strategies for preventing it.
31. Calculate the number of parameters in a convolutional layer with a 4x4 input, eight 3x3 filters, and no padding.
32. What is the purpose of the Softmax function in the final layer of a CNN designed for image classification?
33. Why is flattening used in CNNs?
34. How does the choice of loss function impact the training of a CNN for a specific task, such as image segmentation or object detection?
35. Explain the concept of feature maps and their role in CNNs.
36. What are the advantages and disadvantages of using a larger batch size during CNN training?
37. Compare and contrast the architecture and use cases of popular CNN models like VGG, ResNet, and MobileNet.
38. Let's consider scenarios where labeled data is scarce, such as in agricultural drones for crop disease detection, how CNNs utilize transfer learning, and what benefits does this approach offer for accurate recognition?
39. Discuss about, in real-time applications like augmented reality, why might a shallower CNN architecture be preferred.
40. What is an ablation study?
41. How to prove that your CNN is doing a good job? What kind of experiments should be displayed in this proof?
42. How to achieve the best CNN architecture?
43. What is a GPU and why is it important for CNNs? Is there a significant difference in running the CNN model on a CPU?