

Assignment 5

Convolutional Neural Networks

Image Recognition

Instructions. This assignment focuses on image classification using convolutional neural networks. We will use the MobileNetV2 CNN, which is a lightweight and efficient network developed by Google and is commonly used in smartphone apps. Explore the [Week 5 web app](#) and answer the questions below. Stick to 3 sentences maximum per question. Upload your responses to BruinLearn as a PDF by the end of the day (11:59pm PT).

Question 1. In the CNN Builder tab, train two models using 20 epochs and default settings (32 filters, 3×3 kernel, max pooling): one with 1 convolutional block and 500 samples, and one with 4 convolutional blocks and 100 samples. Report the validation accuracy for both. Which performs better?

Question 2. Now try to find the best model you can build from scratch using the CNN Builder. Experiment with the number of convolutional blocks, filters, and training samples to maximize validation accuracy. Report your best architecture and validation accuracy.

Question 3. Switch to the Transfer Learning tab. First train MobileNetV2 using the Frozen Backbone strategy with 500 samples, then train again using Fine-Tuning with 500 samples. Report the validation accuracy for both. Comparing all three approaches — your best scratch-built CNN from Question 2, Frozen Backbone, and Fine-Tuning — which performs best? Explain why briefly.

Question 4. A birding nonprofit is building a smartphone app to help volunteers identify three North American woodpecker species (Downy, Hairy, and Pileated) from field photos. Examples are shown below.



Downy Woodpecker



Hairy Woodpecker



Pileated Woodpecker

Based on your knowledge of computer vision, what advice would you give to the non-profit for building a reliable classifier for their app? Comment briefly on data, architecture, and your training approach. (Many correct answers).