

## Homework 5 - CNN Models - Spring 2025

**Essay Question:** *Even though the basic convolution neural network (CNN) architecture appeared during the second wave of neural network research [Fukushima, IEEE SMC-Trans, 1983], there were several unique features in Krizhevsky's CNN model [Krizhevsky, NeurIPS, 2012] that were critical for the renewed interest in deep learning models over the past decade. Write a one paragraph essay that starts with the preceding prompt. The following sentences should support this assertion by identifying the unique features of Krizhevsky's work and discussing why these features were instrumental in driving the recent resurgence of neural network research. Your paragraph should close with your opinion regarding the relative importance of these features and how they may drive or hinder future neural network research.*

**Problem 1:** Consider a CNN model that takes an JPG image of shape (128,128,3) and passes it through 3 Conv2D layers with ReLU activations, a string of 2, and no zero padding. The  $\ell$ th Conv2D layer ( $\ell = 1, 2, 3$ ) outputs  $\ell \times 16$  channels. The output of the last Conv2D layer is flattened and mapped to 5 real-valued outputs through a Dense layer using a softmax activation function.

- Write a Keras script that instantiates the *Model* object for this CNN.
- Determine the shape of each layer's output.
- Determine how many trainable parameters are in the model.

**Problem 2:** Consider the 1D convolutional network declared by the following Keras script,

```
inputs = Keras.Input(shape=(6,1))
outputs = layers.Conv1D(1,3, activation = "relu", use_bias=False)(inputs)
model = keras.Model(inputs = inputs, outputs = outputs)
```

If all trainable parameters are equal to one, then determine the model's output when the input is a rank-1 tensor  $\begin{bmatrix} 0 & -1 & 2 & -1 & 1 & 0 \end{bmatrix}$ .

**Notebook Assignment 1 (brain tumor classification):** Gathering and labeling data is expensive and time consuming. It is therefore important to know how to train models when there is limited data available. One way of doing this is to *transfer* the knowledge contained in one model to another model. This notebook assignment asks to use a pretrained CNN model and use it for the classification analysis of brain MRI scan images. The classification task is to classify each MRI image as either having a "tumor" or "no-tumor".

The Brain MRI image dataset (zip file on course website) consists of 253 image samples of high-resolution brain MRI scans. The images are grayscale in nature and vary in size. The dataset has been pre-processed so the images are placed in two directories TRAIN and TEST. Each of these directories has two subdirectories NO or YES representing the true label of each image. The images are jpg files with shape (244,244,3).

1. Create a `train_ds` and `test_ds` dataset objects using the keras utility "image data set from directory". Your dataset should declare the image size as (244,244) and create batches of size 32. Use a 30% validation split to decompose the training dataset object into a p-training and validation dataset object. Print out the number of batches in the p-training, the validation, and the testing dataset objects.

2. Display 16 of the training images in a 4 by 4 array and have the title on each image indicates its true class label
3. From `keras.applications` import the pre-trained model ResNet50V2. Set the input shape to (244,244,3) and leave off the top layer of the model. Fix the model weights and then add a classification layer on top that consists of data augmentation that randomly flips, rotates, and zooms the input images, a GlobalAveragePooling2D, 64 node Dense layer, a Dropout layer with 0.3 dropout rate and a dense layer that has a single node as its output with a sigmoid activation function. Use the model summary routines to determine the total number of parameters in the model and the number of trainable weights. Be sure to rescale the inputs to the convolutional base so they are in the range from 0 to 1. Compile the model with a Adam optimizer using a binary crossentropy loss function. Print out the number of trainable parameters in your refined model. How many parameters are in the fixed convolutional base?
4. Checkpoint the model to save the model with the smallest validation loss. Train the model using p-training dataset for 20 epochs and that validates the models using the validation dataset. Plot the training curves for your model and determine the accuracy of the best model.
5. Let TP denote the number of true positive classifications, FP denotes the number of false positive classifications, TN being the number of true negative classifications, and FN being the number of false negative classifications. Accuracy is the (TP+TN) divided by total number of examples when the model output is greater than 0.5. It is a coarse measure of a classifier's performance. In practice, we look at the

$$\begin{aligned} \text{TPR} &= \text{True Positive Rate} = \frac{\text{TP}}{\text{TP} + \text{FN}} \\ \text{FPR} &= \text{False Positive Rate} = \frac{\text{FP}}{\text{TN} + \text{FP}} \end{aligned}$$

for a variety of threshold, not just 0.5. Compute the TPR and FPR for your classifier for the following thresholds

$$\text{thresholds} = \{0.025, 0.05, 0.075, .1, .2, .3, .4, .5, .6, .7, .8, .9, .925, .950, .975\}$$

and plot the TPR vs FPR to form the classifiers Receiver Operating Characteristic. Also plot the heatmap of the classifier's confusion matrix when the threshold is 0.5.