(1) Monitoring the neural network learning process

(a) You are training three neural networks and you are observing the following loss functions on the training and validation sets. What can we conclude for each of the three networks? In which training epoch would you stop training each network?

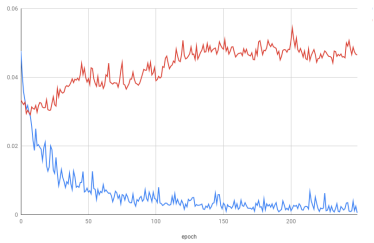

Network 1

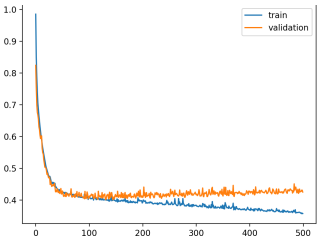

Network 2

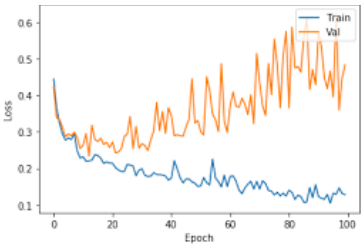

Network 3

(b) You have trained two neural networks and visualized various kernels learned in the first layer. Which of the two networks would you use and why?

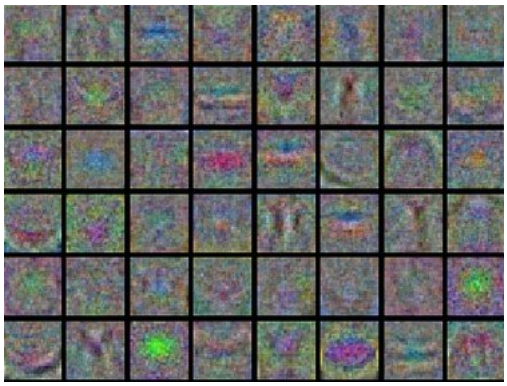

Network 1

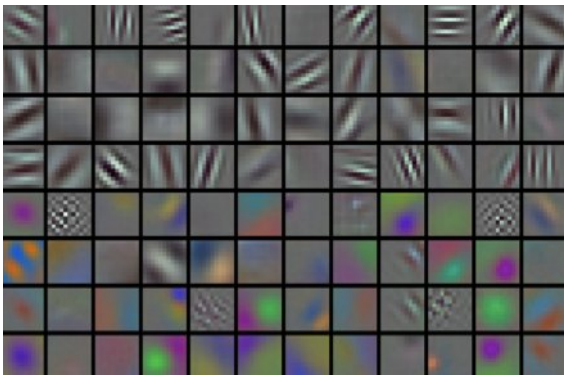

Network 2

(2) Convolution operations

(a) Consider a 1-dimensional signal

$$F = [1 \quad 2 \quad 1 \quad 3 \quad 2 \quad 3 \quad 1 \quad 2 \quad 3 \quad 8 \quad 7 \quad 8 \quad 9 \quad 9 \quad 7 \quad 8]$$

We will define the 1-dimensional convolution of the signal F with a filter F at point i as:

$$(F \star W)[i] = \sum_{-\infty}^{\infty} F[j]W[j]$$

(a.i) Calculate the convolution of F with filter $W_1 = [1 \quad 1 \quad 1]$, i.e., $F1 = F \star W_1$. What is the operation that filter W_1 performs? *Note:* You can perform zero-padding for the edge elements.

(a.ii) Calculate the convolution of F with filter $W_2 = [1 \quad 0 \quad -1]$, i.e., $F2 = F \star W_2$. What is the operation that filter W_2 performs? *Note:* You can perform zero-padding for the edge elements.

(b) Consider an image $I \in \mathbb{R}^{D \times D}$ and the following 2-d convolution filters.
What operations will the following filters apply to image I?

$$F_1 = \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}, F_2 = \begin{bmatrix} -1 & -2 & -1 \\ 0 & 0 & 0 \\ 1 & 2 & 1 \end{bmatrix}, F_3 = \begin{bmatrix} -1 & -1 & -1 \\ -1 & 9 & -1 \\ -1 & -1 & -1 \end{bmatrix}$$

(c.i) A gray color 10×10 image is the input to a CNN. The first layer of the CNN has 64 filters of dimensionality 3×3, stride size of 1, and zero-padding. What is the output of the first hidden layer and how many parameters are learned by this layer?

(c.ii) A color 10×10 image is the input to a CNN. The first layer of the CNN has 64 filters of dimensionality 3×3, stride size of 1, and zero-padding. What is the output of the first hidden layer and how many parameters are learned by this layer?

(3) Neural network hyper-parameters

(a) What are important factors that contribute to selecting the depth of a neural network?
(i) Type of neural network (eg. MLP, CNN etc)
(ii) Input data
(iii) Computation power, i.e. Hardware capabilities and software capabilities
(iv) Learning Rate

(b) You are fine-tuning a pre-trained neural network. What would be the best strategy for the learning rate?
(i) Keep the same learning rate as the one in the pre-trained network.
(ii) Have a larger learning rate compared to the one in the pre-trained network.
(iii) Have a smaller learning rate compared to the one in the pre-trained network.
(iv) The learning rate will not affect the fine-tuning result.

(4) Fine-tuning or not fine-tuning
You have pre-trained a CNN on a large source dataset of images with plants from the area of Alaska. You are also given a target dataset for which you are called to make decisions on: (i) whether to use and fine-tine the pre-trained CNN or train a CNN from scratch; and (ii) if fine-tuning is used, which part of the network to fine-tune (e.g., convolutional layers, fully-connected layers). Please provide your answer and justification for each of the following cases.

(a) The target dataset is small and includes plants from the area of Texas.

(b) The target dataset is large and includes plants from the area of Texas.

(c) The target dataset is small and includes images of various mechanical parts.

(d) The target dataset is large and includes images of various mechanical parts.