Return to Edit Assessment

Name: ________________________
Score: ______ / ______

Worksheet 05 - autoencoder and convnet

Part 1

Please state the names of all the students you worked with on this assignment:

Answer Point Value: 0.0 points

Model Short Answer: --------

Part 2: Autoencoder/PCA

Multiple applications of PCA

Consider a dataset $\mathbf{X} \in \mathbb{R}^{T\times n}$. We perform a PCA on the dataset $\mathbf{X}$ by computing

$$\begin{aligned} \mathbf{C}^{(1)} &= \frac{\mathbf{X}^\top \mathbf{X}}{T} \\ \mathbf{C}^{(1)} \mathbf{w}_i^{(1)} &= \left(\sigma_i^{(1)}\right)^2 \mathbf{w}_i^{(1)} \end{aligned}$$

and projecting onto the $m < n$ principal components $\mathbf{W}_{1\ldots m}^{(1)} = [\mathbf{w}_1, \ldots, \mathbf{w}_m]$ with largest eigenvalues:

$$\mathbf{Y} = \mathbf{X}\mathbf{W}_{1\ldots m}^{(1)}.$$

We perform PCA a second time by computing:

$$\begin{aligned} \mathbf{C}^{(2)} &= \frac{\mathbf{Y}^\top \mathbf{Y}}{T} \\ \mathbf{C}^{(2)} \mathbf{w}_i^{(2)} &= \left(\sigma_i^{(2)}\right)^2 \mathbf{w}_i^{(2)} \end{aligned}$$

projecting onto the single principal component $\mathbf{w}_1^{(2)}$ with largest eigenvalue
\projecting onto the single principal component $\mathbf{w}_1^{(2)}$ with largest eigenvalue

$\left(\sigma_1^{(2)}\right)^2$, resulting in $\mathbf{z} \in \mathbb{R}^T$.

Which of the following is true?

- A. $\sigma_1^{(2)} < \sigma_2^{(2)}$
- B. $\mathbf{w}_1^{(1)} = \mathbf{w}_2^{(2)}$
- C. $\mathbf{Y}_{*,1} = \mathbf{z}$
- D. $\left\| \mathbf{w}_1^{(1)} - \mathbf{w}_1^{(2)} \right\| = 1$

Answer Point Value: 1.0 points

Answer Key: C

Autoencoder

Consider grayscale image data with 128×128 pixels. We consider two autoencoder architectures with Rectified Linear Units as nonlinearities:

Dense:
Input layer
Dense hidden layer with 256 neurons
Dense output layer.

Convolutional:
Input layer
Strided convolutional layer with 64 filters, stride 8, "valid" padding (no zeros added) and a filter size of 8×8.
Max-Pooling across the 64 channels.
Strided transposed convolutional layer with 64 filters, stride 8 that recover(s) the input (image) dimensions as output.
Sum-Pooling across the 64 channels.

Ignoring the bias parameters, how many parameters do we have?

- A. $32,768$ **(dense) and** $1,048,576$ **(convolutional)**
- B. $4,194,304$ **(dense) and** $4,096$ **(convolutional)**
- C. $4,194,304$ **(dense) and** $8,192$ **(convolutional)**
- D. $8,388,608$ **(dense) and** $8,192$ **(convolutional)**

- E. $8,388,608$ **(dense) and** $16,384$ **(convolutional)**
- F. $16,777,216$ **(dense) and** $8,192$ **(convolutional)**
- G. $16,777,216$ **(dense) and** $16,384$ **(convolutional)**

Answer Point Value: 1.0 points

Answer Key: D

PCA

Given a real-valued dataset $\mathbf{X}$ and the covariance matrix $\mathbf{C} = \mathbf{X}^\top \mathbf{X}$.

Which of the following statements is <u>not</u> true?

- A. **The PCA eigenvalues can come in complex conjugate pairs.**
- B. $\mathbf{C} = \mathbf{C}^\top$
- C. **The PCA eigenvalues are always non-negative.**

Answer Point Value: 1.0 points

Answer Key: A

Part 3: Convnet

Convolutional Neural Network

Given a dataset with 130×130 px images and 3 channels, consider the following convolutional neural network:

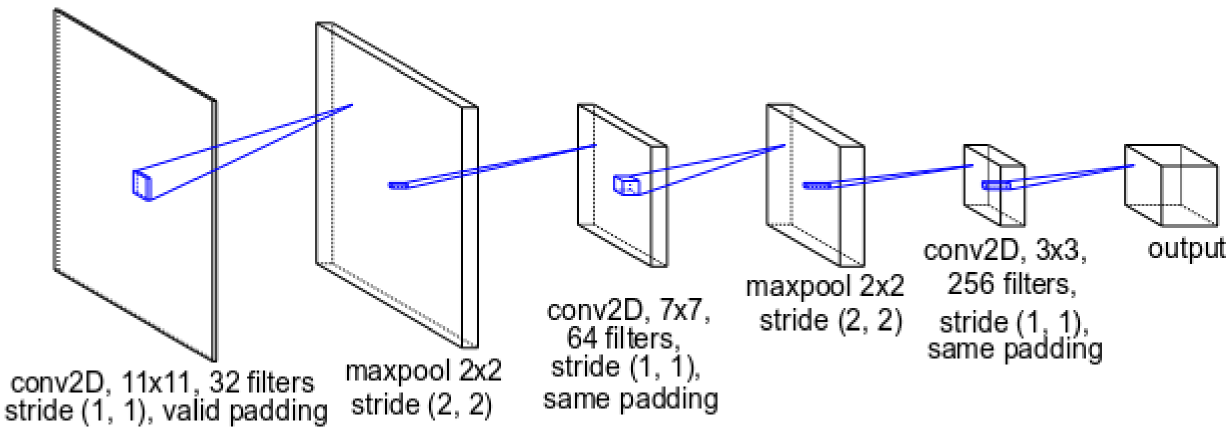

What is the output dimension?

- A. $30 \times 30 \times 256$
- B. $30 \times 30 \times 524288$
- C. $60 \times 60 \times 256$
- D. $60 \times 60 \times 524288$

Answer Point Value: 1.0 points

Answer Key: A

Kernels

We have a Gaussian function in two dimensions given by

$$G(x, y) = \frac{1}{2\pi\sigma^2} e^{-\frac{x^2+y^2}{2\sigma^2}}.$$

One can use such a Gaussian to parameterize a convolution kernel K by, e.g., spanning an odd-sized equidistant grid centered around $(0, 0)^\top$, evaluating G on the grid points, and normalizing the resulting matrix with respect to the sum of its values.

Can K be written as outer product $K = \mathbf{u}\mathbf{v}^\top$?

- A. **Yes, in this case.**
- B. **No**
- C. **Yes, this can be done for any kernel.**

Answer Point Value: 1.0 points

Answer Key: A

Output dimensions

Consider a convolution kernel of size $n \times n$, a stride $s > 0$, and an input image of size $w \times h$.

Assuming w, h, and n being divisible by the stride s, what is the output dimension when applying valid padding?

- A. $[(w-n)/s+1] \times [(h-n)/s+1]$
- B. $[(w-2n)/s+1] \times [(h-2n)/s+1]$
- C. $[wh/s+1] \times [wh/s+1]$
- D. $[(w-n)s+1] \times [(h-n)s+1]$
- E. $[(w-n)/s-1] \times [(h-n)/s-1]$

Answer Point Value: 1.0 points

Answer Key: A

Sparsity in convolutional neural networks

Given images of dimension 8×10 and a network of depth d that contains the same convolutional layer d times with kernel size 5×5, stride 1, and same padding.

At which minimum depth does every feature pixel in the last layer depend on all input pixels?

- A. 4
- B. 5
- C. 6
- D. **Such a convolutional layer never depends on all pixels of the input image.**

Answer Point Value: 1.0 points

Answer Key: B

Valid padding

Given a concatenation of d identical convolutional layers with valid padding, kernel size 4×4, stride 2 and input image size 20×20.

At which depth d of the network do the kernel dimensions exceed the dimensions of the layer's input?

- A. $d = 2$
- B. $d = 3$
- C. $d = 4$
- D. $d = 5$

Answer Point Value: 1.0 points

Answer Key: B