Mathematical Foundations of Deep Neural Networks, M1407.001200
E. Ryu
Spring 2024

Homework 6
Due 5pm, Monday, April 22, 2024

Problem 1: *Dropout-ReLU=ReLU-Dropout.* Consider the following layer

```
class myLayer(nn.Module):
  def __init__(self, input_size, output_size):
    super(myLayer, self).__init__()
    self.linear = nn.Linear(input_size,output_size)
    self.sigma = nn.ReLU()
    # self.sigma = nn.Sigmoid()
    # self.sigma = nn.LeakyReLU()
    self.dropout= nn.Dropout(p=0.4)
  def forward(self, x):
    return dropout(sigma(linear))
    # return sigma(dropout(linear))  # Is this is equivalent?
```

In which of the three following cases are the operations linear-dropout-σ and linear-σ-dropout equivalent?

(a) `self.sigma = nn.ReLU()`

(b) `self.sigma = nn.Sigmoid()`

(c) `self.sigma = nn.LeakyReLU()`

Problem 2: *Default weight initialization.* Consider the multi-layer perceptron

$$y_L = A_L y_{L-1} + b_L$$
$$y_{L-1} = \sigma(A_{L-1} y_{L-2} + b_{L-1})$$
$$\vdots$$
$$y_2 = \sigma(A_2 y_1 + b_2)$$
$$y_1 = \sigma(A_1 x + b_1),$$

where $x \in \mathbb{R}^{n_0}$, $A_\ell \in \mathbb{R}^{n_\ell \times n_{\ell-1}}$, $b_\ell \in \mathbb{R}^{n_\ell}$, and $n_L = 1$. For the sake of simplicity, let

$$\sigma(z) = z.$$

Assume $x_1, \ldots, x_{n_0}$ are IID with zero-mean and unit variance. If this network is initialized with the default weight initialization of PyTorch, what will the mean and variance of y_L be?

Clarification. For this problem, you are being asked to read the PyTorch source code
`https://pytorch.org/docs/stable/_modules/torch/nn/modules/linear.html`
to identify the default initialization behavior and then to perform calculations.

Problem 3: *Backprop for MLP with residual connections.* Let $\sigma\colon \mathbb{R} \to \mathbb{R}$ be a differentiable activation function and consider the following MLP with residual connections

$$y_L = A_L y_{L-1} + b_L$$
$$y_{L-1} = \sigma(A_{L-1} y_{L-2} + b_{L-1}) + y_{L-2}$$
$$\vdots$$
$$y_3 = \sigma(A_3 y_2 + b_3) + y_2$$
$$y_2 = \sigma(A_2 y_1 + b_2) + y_1$$
$$y_1 = \sigma(A_1 x + b_1),$$

where $x \in \mathbb{R}^n$, $A_1 \in \mathbb{R}^{m \times n}$, $b_1 \in \mathbb{R}^m$, $A_\ell \in \mathbb{R}^{m \times m}$, $b_\ell \in \mathbb{R}^m$ for $\ell = 2, \ldots, L-1$, and $A_L \in \mathbb{R}^{1 \times m}$, $b_L \in \mathbb{R}^1$. (To clarify, σ is applied element-wise.) For notational convenience, define $y_0 = x$.

(i) Find formulae for
$$\frac{\partial y_\ell}{\partial y_{\ell-1}}$$
for $\ell = 2, \ldots, L$.

(ii) Find formulae for
$$\frac{\partial y_L}{\partial b_\ell}, \qquad \frac{\partial y_L}{\partial A_\ell}$$
for $\ell = 1, \ldots, L$.

(iii) The gradients
$$\frac{\partial y_L}{\partial b_i}, \qquad \frac{\partial y_L}{\partial A_i}$$
for $i = 1, \ldots, \ell$ need not vanish when $[A_j = 0$ for some $j \in \{\ell + 1, \ldots, L - 1\}]$ or $[\sigma'(A_j y_{j-1} + b_j) = 0$ for some $j \in \{\ell + 1, \ldots, L - 1\}]$. Explain why.

Problem 4: *Split-transform-merge convolutions.* Consider a series of 1×1, 3×3, 1×1 conv-ReLU operations with 256–128–128–256 channels:

```
class MyConvLayer(nn.Module):
  def __init__(self):
    super(MyConvLayer, self).__init__()
    self.conv1 = nn.Conv2d(256, 128, 1,)
    self.conv2 = nn.Conv2d(128, 128, 3, padding=1)
    self.conv3 = nn.Conv2d(128, 256, 1)
  def forward(self, x):
    out = torch.nn.functional.relu(self.conv1(x))
    out = torch.nn.functional.relu(self.conv2(out))
    out = torch.nn.functional.relu(self.conv3(out))
    return out
```

An issue with this construction, however, is that it has too many trainable parameters. To reduce the number of trainable parameters, we use the following *split-transform-merge* structure: [apply a series of 1×1, 3×3, 1×1 conv-ReLU operations with 256–4–4–256 channels] a total of 32 times and sum the 32 outputs. The following figure illustrates this construction.
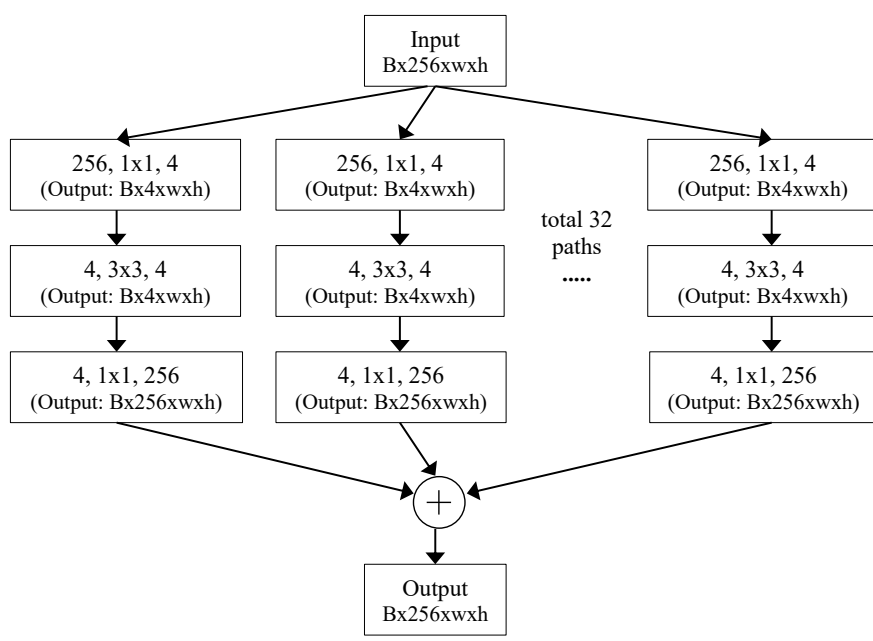
To clarify, all convolutions use biases and the strides are all equal to 1. ReLU is not applied after the sum operation.

(a) How many trainable parameters are present in both constructions?

(b) In the following page, implement this convolution with the split-transform-merge structure.

```python
class STMConvLayer(nn.Module):
  def __init__(self):
    super(STMConvLayer, self).__init__()
    #-----------------------------------
    # Fill in code here




    #-----------------------------------
  def forward(self, x):
    # [apply 1x1conv with 4 output channels
    #  apply 3x3conv with 4 output channels (with padding=1)
    #  apply 1x1conv with 256 output channels] X 32
    # Add all 32 outputs
    #-----------------------------------
    # Fill in code here





    #-----------------------------------

    return out
```

Problem 5: *Regularization can mitigate double descent.* Assume we have labels $Y_1, \ldots, Y_{N_{\text{train}}} \in \mathbb{R}$ generated IID as $X_i \sim \mathcal{N}(0, I_d)$ and $Y_i \sim X_i^\intercal \beta^\star + \mathcal{N}(0, \sigma^2)$ for $i = 1, \ldots, N_{\text{train}}$, where $\beta^\star \in \mathbb{R}^d$. Use $d = 35$ and $\sigma = 0.5$, and $N_{\text{train}} = 300$. Fit the data with a 2-layer ReLU network $f_{\theta,W}(x) = \theta^\intercal \text{ReLU}(Wx)$ with $\theta \in \mathbb{R}^p$ and $W \in \mathbb{R}^{p \times d}$. Assume $W_{ij} \sim \mathcal{N}(0, 1/p)$ IID. For simplicity, assume W is fixed (not trained) once initialized. Train θ via

$$\underset{\theta \in \mathbb{R}^p}{\text{minimize}} \quad \sum_{i=1}^{N_{\text{train}}} \frac{1}{2}(f_{\theta,W}(X_i) - Y_i)^2 + \frac{\lambda}{2}\|\theta\|^2$$

with $\lambda > 0$. Using the notation $\tilde{X}_i = \text{ReLU}(WX_i)$ for $i = 1, \ldots, N_{\text{train}}$ and

$$\tilde{X} = \begin{bmatrix} \tilde{X}_1^\intercal \\ \vdots \\ \tilde{X}_{N_{\text{train}}}^\intercal \end{bmatrix} \in \mathbb{R}^{N_{\text{train}} \times p}, \qquad Y = \begin{bmatrix} Y_1 \\ \vdots \\ Y_{N_{\text{train}}} \end{bmatrix} \in \mathbb{R}^{N_{\text{train}}},$$

we can equivalently express the optimization problem as

$$\underset{\theta \in \mathbb{R}^p}{\text{minimize}} \quad \frac{1}{2}\|\tilde{X}\theta - Y\|^2 + \frac{\lambda}{2}\|\theta\|^2.$$

Train (compute the global minimum) by using linear algebra to solve the least-squares problem. With the fixed regularization parameter $\lambda = 0.01$, we indeed observe the double descent phenomenon when we plot the test error against the number of parameters p. Show that the double descent phenomenon vanishes if λ is tuned. Specifically, use the training dataset to (precisely) compute θ and use the validation dataset of size $N_{\text{validation}} = 60$ to (roughly) tune for $\lambda \in [10^{-2}, 10^2]$. (You should separately tune λ for each p, as you would do in practice.) Then, use the test dataset of size $N_{\text{test}} = 30$ to plot the test error for each p and its corresponding optimal λ. Use the starter code `ddescent.py`.

Remark. This problem was inspired by [1].

Hint. The results should look something like:
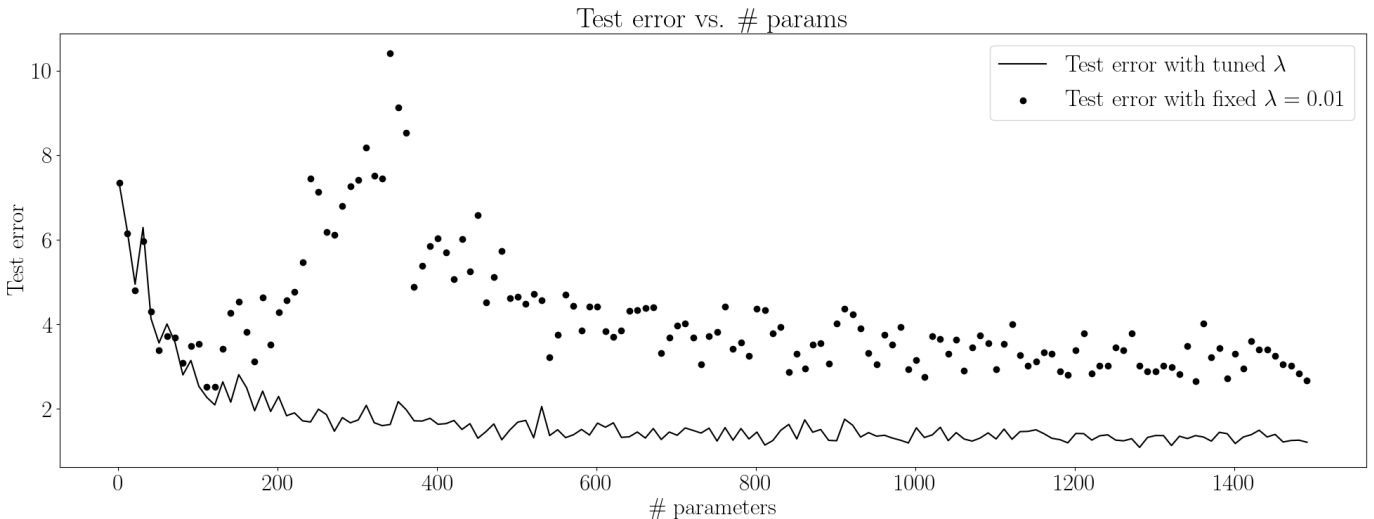
References

[1] P. Nakkiran, P. Venkat, S. Kakade, and T. Ma. Optimal regularization can mitigate double descent, *ICLR*, 2021.