

RECITATION 5

NEURAL NETWORKS

10-301/10-601: INTRODUCTION TO MACHINE LEARNING

2024-02-28

1 Matrix Calculus

Consider $\mathbf{x} \in \mathbb{R}^p$, $\mathbf{y} \in \mathbb{R}^r$, $\mathbf{z} \in \mathbb{R}^n$ where $\mathbf{z} = g(\mathbf{y})$, and $\mathbf{y} = f(\mathbf{x})$. We want to derive $d\mathbf{z}/d\mathbf{x}$ (a vector form of the scalar chain rule).

1. If \mathbf{x} , \mathbf{y} , and \mathbf{z} were all scalars, what would dz/dx be?

Shape matching:

2. Fill in the following shapes:

$$\frac{d\mathbf{y}}{d\mathbf{x}} :$$

$$\frac{d\mathbf{z}}{d\mathbf{y}} :$$

$$\frac{d\mathbf{z}}{d\mathbf{x}} :$$

3. Therefore, the correct derivative is

$$\frac{d\mathbf{z}}{d\mathbf{x}} =$$

Generalizing a single element: In order to ensure your derivatives are correct, we recommend you use matrix calculus rules whenever possible. When you're not sure how to apply a rule or if one applies, use the method of generalizing a single element.

4. Example: Suppose $\mathbf{A} \in \mathbb{R}^{n \times n}$, $\mathbf{x} \in \mathbb{R}^n$, and we want to compute $\frac{d\mathbf{Ax}}{d\mathbf{x}}$. Note that our numerator and denominator are both length- n vectors, so our derivative has shape $\mathbb{R}^{n \times n}$.
5. Is $\left(\frac{d\mathbf{Ax}}{d\mathbf{x}}\right)_{ij}$ equal to $\frac{d\mathbf{Ax}_i}{dx_j}$ or $\frac{d\mathbf{Ax}_j}{dx_i}$? Why?

6. Compute $\left(\frac{d\mathbf{Ax}}{d\mathbf{x}}\right)_{ij}$:

Applying Matrix Calculus For example, suppose we are finding the closed-form solution to linear regression: given $\mathbf{X} \in \mathbb{R}^{n \times d}$, $\mathbf{y} \in \mathbb{R}^n$, we wish to find $\boldsymbol{\theta} \in \mathbb{R}^d$ that minimizes $\|\mathbf{y} - \mathbf{X}\boldsymbol{\theta}\|_2^2$.

7. What is the shape of $\mathbf{X}^\top \mathbf{y}$? Is this our solution?
8. What is the closed-form solution? Use matrix calculus to derive the solution.

2 Forward Propagation




Figure 1: Neural Network For Example Questions

Forward Propagation is the process of calculating the value of your loss function, given data, weights and activation functions. Given the input data \mathbf{x} , we can transform it by the given weights, $\boldsymbol{\alpha}$, then apply the corresponding activation function to it and finally pass the result to the next layer. Forward propagation does not involve taking derivatives and proceeds from the input layer to the output layer.

Network Overview Consider the neural network with one hidden layer shown in Figure 2. The input layer consists of 2 features $\mathbf{x} = [x_1, x_2]^T$, the hidden layer has 2 nodes with output $\mathbf{z} = [z_1, z_2]^T$, and the output layer is a scalar \hat{y} . We also add an intercept to the input, $x_0 = 1$ and the output of the hidden layer $z_0 = 1$, both of which are fixed to 1.

$\boldsymbol{\alpha}$ is the matrix of weights from the inputs to the hidden layer and $\boldsymbol{\beta}$ is the matrix of weights from the hidden layer to the output layer. $\alpha_{j,i}$ represents the weight going to the node z_j in the hidden layer from the node x_i in the input layer (e.g. $\alpha_{1,2}$ is the weight from x_2 to z_1), and $\boldsymbol{\beta}$ is defined similarly. We will use a **ReLU** activation function for the hidden layer and no activation for the output layer.

Network Details Equivalently, we define each of the following.

The input:

$$\mathbf{x} = [x_0, x_1, x_2]^T \quad (1)$$

Linear combination at the first (hidden) layer:

$$a_j = \sum_{i=0}^2 \alpha_{j,i} \cdot x_i, \quad \forall j \in \{1, \dots, 2\} \quad (2)$$

Activation at the first (hidden) layer:

$$z_j = \text{ReLU}(a_j) = \max(0, a_j), \quad \forall j \in \{1, \dots, 2\} \quad (3)$$

$$\mathbf{z} = [z_0, z_1, z_2]^T \quad (4)$$

Linear combination at the second (output) layer:

$$\hat{y} = \sum_{j=0}^2 \beta_j \cdot z_j, \quad (5)$$

Here we fold in the intercept term $\alpha_{j,0}$ by thinking of $x_0 = 1$, and fold in β_0 by thinking of $z_0 = 1$.

Loss We will use Squared error loss, $\ell(\hat{y}, y)$:

$$\ell(\hat{y}, y) = \frac{1}{2}(\hat{y} - y)^2 \quad (6)$$

1. Why and how do we include an intercept term in the input and in the hidden-layer?

2. Why do we need to use nonlinear activation functions in our neural net?

We initialize the network weights as:

$$\boldsymbol{\alpha} = \begin{bmatrix} 2 & 1 & 0 \\ 0 & 2 & 0 \end{bmatrix}$$

$$\boldsymbol{\beta} = [0 \quad 1 \quad 2]$$

For the following questions, we use $y = 3$.

1. **Scalar Form:**

- Given $x_1 = 1$, $x_2 = 2$, What are the values of a ?

$$a_2 = \sum_{i=0}^2 \alpha_{2,i} x_i =$$

- Given $z_1 = 0$, $z_2 = 1$ calculate \hat{y}, l

$$\hat{y} = \sum_{i=0}^2 \beta_i \cdot z_i =$$


2. **Vector Form:** Find the vector form of forward computation, given \mathbf{x} is a column vector.

3 Computation Diagrams

- For the following function f , create the computation graph using the conventions defined in lecture.

$$f(x, y, z) = \frac{\sin(xy)}{z} + \frac{e^{xz}}{xy}$$

- For the following neural network, draw the corresponding computation graph. Assume that all hidden units use the ReLU function as the activation function and that the loss is mean squared error. Provide the shape of all parameters defined in the computation graph. Assume the weights for the first layer and second layers are respectively the matrices α and β .



4 Backward Propagation




Figure 2: Neural Network For Example Questions

Given a Neural Network and a corresponding loss function $J(\theta)$, backpropagation gives us the gradient of the loss function with respect to the weights of the neural network. The method is called *backward* propagation because we calculate the gradients of the final layer of weights first, then proceed backward to the first layer. In a simple neural network with one hidden layer, the partial derivatives that we need for learning are $\frac{\partial \ell}{\partial \alpha_{ji}}$ and $\frac{\partial \ell}{\partial \beta_{kj}}$, and we need to apply chain rule recursively to obtain these. Note that in implementation, it is easier to use matrix/vector forms to conduct computations.

1. Many gradients are calculated in back propagation. Which of these gradients are used to update the weights? Do not include intermediate value(s) used to calculate these gradient(s).

2. **Scalar Form:** Given

- $x_1 = 1, x_2 = 2$
- $a_1 = 3, a_2 = 2$
- $z_1 = 3, z_2 = 2$
- $\alpha = \begin{bmatrix} 2 & 1 & 0 \\ 0 & 2 & 0 \end{bmatrix}$
- $\beta = [0 \ 1 \ 2]$
- $y = 3,$

what are the values of $\frac{\partial \ell}{\partial \beta_1}, \frac{\partial \ell}{\partial \alpha_{1,1}}$?

Hint: Derive expressions for $\frac{\partial \ell}{\partial \beta_i}$ and $\frac{\partial \ell}{\partial \alpha_{j,i}}$ first, then substitute in values.

For convenience, the computation graph for the neural network is displayed below:

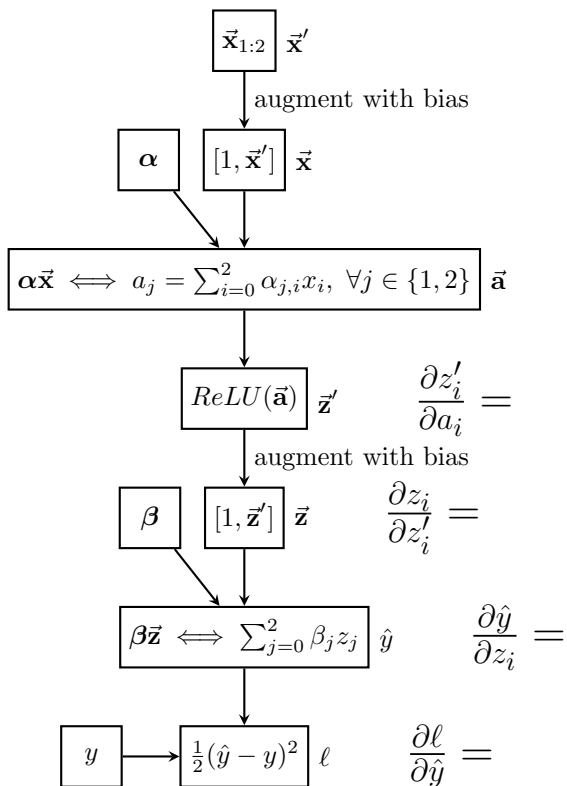


Figure 3: Computation graph for the neural network in Figure 2

Hint: $\frac{\partial ReLU(x)}{\partial x} = 1$ if $x > 0$, 0 if $x \leq 0$

$$\frac{\partial \ell}{\partial \beta_i} = \frac{\partial \ell}{\partial \hat{y}} \frac{\partial \hat{y}}{\partial \beta_i} \qquad \frac{\partial \ell}{\partial \beta_1} =$$

$$\frac{\partial \ell}{\partial \alpha_{j,i}} = \frac{\partial \ell}{\partial \hat{y}} \frac{\partial \hat{y}}{\partial z_j} \frac{\partial z_j}{\partial a_j} \frac{\partial a_j}{\partial \alpha_{j,i}}$$

$$\frac{\partial \ell}{\partial \alpha_{1,1}} =$$

3. **Vector Form:** What are the values of $\frac{\partial \ell}{\partial \beta}$, $\frac{\partial \ell}{\partial \alpha}$?