

Training Neural Networks

CMPUT 366: Intelligent Systems

GBC §6.5

Lecture Outline

1. Recap & Logistics
2. Gradient Descent for Neural Networks
3. Automatic Differentiation
4. Back-Propagation

Assignment #2

- Assignment #2 was due on Monday
 - Late submission deadline **TODAY at 11:59pm**
 - Submit via eClass

Recap: Nonlinear Features

$$y = f(\mathbf{x}; \mathbf{w}) = g(\mathbf{w}^T \mathbf{x}) = g\left(\sum_{i=1}^n w_i x_i\right)$$

Generalized linear model: **Activation function** g of linear combination of inputs


Extension: Learn a generalized linear model on **richer inputs**

1. Define a **feature mapping** $\phi(\mathbf{x})$ that returns **functions** of the original inputs
2. Learn a linear model of the **features** instead of the **inputs**

$$y = f(\mathbf{x}; \mathbf{w}) = g(\mathbf{w}^T \phi(\mathbf{x})) = g\left(\sum_{i=1}^n w_i [\phi(\mathbf{x})]_i\right)$$

Recap:

Feedforward Neural Network



$$h_1(\mathbf{x}; \mathbf{w}^{(1)}, b^{(1)}) = g \left(b^{(1)} + \sum_{i=1}^n w_i^{(1)} x_i \right)$$

$$y(\mathbf{x}; \mathbf{w}, \mathbf{b}) = g \left(b^{(y)} + \sum_{i=1}^n w_i^{(y)} h_i(\mathbf{x}_i; \mathbf{w}^{(i)}, b^{(i)}) \right)$$

- A **neural network** is many **units composed** together
- **Feedforward neural network:** Units arranged into **layers**
 - Each layer takes outputs of **previous layer** as its **inputs**

Recap: Chain Rule of Calculus

$$\frac{dz}{dx} = \frac{dz}{dy} \frac{dy}{dx}$$

i.e.,

$$h(x) = f(g(x)) \implies h'(x) = f'(g(x))g'(x)$$


If we know formulas for the derivatives of **components** of a function, then we can build up the derivative of their composition mechanically

Recap: Training Neural Networks

- Specify a **loss** L and a set of **training examples**:

$$E = (\mathbf{x}^{(1)}, y^{(1)}), \dots, (\mathbf{x}^{(n)}, y^{(n)})$$

- Training by **gradient descent**:

1. Compute **loss** on training data: $L(\mathbf{W}, \mathbf{b}) = \sum_i \ell$ 

2. Compute **gradient** of loss: $\nabla L(\mathbf{W}, \mathbf{b})$

3. **Update parameters** to make loss smaller:

$$\begin{bmatrix} \mathbf{W}^{new} \\ \mathbf{b}^{new} \end{bmatrix} = \begin{bmatrix} \mathbf{W}^{old} \\ \mathbf{b}^{old} \end{bmatrix} - \eta \nabla L(\mathbf{W}^{old}, \mathbf{b}^{old})$$


Three Representations

A function $f(x, y)$ can be represented in multiple ways:

1. As a **formula**:

$$f(x, y) = (xy + \sin x + 4)(3y^2 + 6)$$

2. As a **computational graph**:



3. As a **finite numerical algorithm**


$$\begin{aligned} s_1 &= x \\ s_2 &= y \\ s_3 &= s_1 \times s_2 \\ s_4 &= \sin(s_1) \\ s_5 &= s_3 + s_4 \\ s_6 &= s_5 + 4 \\ s_7 &= \text{sqr}(s_2) \\ s_8 &= 3 \times s_7 \\ s_9 &= s_8 + 6 \\ s_{10} &= s_6 \times s_9 \end{aligned}$$

Symbolic Differentiation

$$\begin{aligned}z &= f(y) \\y &= f(x) \\x &= f(w)\end{aligned}$$

$$z = f(f(f(w)))$$

$$\begin{aligned}\frac{\partial z}{\partial w} &= \frac{\partial z}{\partial y} \frac{\partial y}{\partial x} \frac{\partial x}{\partial w} \\&= f'(f(f(w)))f'(f(w))f'(w)\end{aligned}$$



$$f(w) = \begin{cases} w & \text{if } w > 0 \\ 0 & \text{otherwise.} \end{cases}$$

- We can differentiate a nested **formula** by recursively applying the **chain rule** to derive a **new formula** for the gradient
- **Problem:** This can result in a lot of repeated subexpressions
- **Question:** What happens if the nested function is defined **piecewise**?

Automatic Differentiation: Forward Mode


- The forward mode converts a **finite numerical algorithm** for computing a function into an **augmented** finite numerical algorithm for computing the **function's derivative**
- For each step, a new step is constructed representing the derivative of the corresponding step in the original program:

$$\begin{array}{l} s_1 = x \\ s_2 = y \\ s_3 = s_1 + s_2 \\ s_4 = s_1 \times s_2 \\ \vdots \end{array} \quad \Longrightarrow \quad \begin{array}{l} s'_1 = 1 \\ s'_2 = 0 \\ s'_3 = s'_1 + s'_2 \\ s'_4 = s_1 \times s'_2 + s'_1 \times s_2 \\ \vdots \end{array}$$

- To compute the partial derivative $\frac{\partial s_n}{\partial s_1}$, set $s'_1 = 1$ and $s'_2 = 0$ and run augmented algorithm
- This takes roughly twice as long to run as the original algorithm (**why?**)

Forward Mode Example

Let's compute $\left. \frac{\partial f}{\partial y} \right|_{x=2, y=8}$ using forward mode:



$$\begin{aligned}
 s_1 &= x &= 2 \\
 s_2 &= y &= 8 \\
 s_3 &= s_1 \times s_2 &= 16 \\
 s_4 &= \sin(s_1) &\approx 0.034 \\
 s_5 &= s_3 + s_4 &= 16.034 \\
 s_6 &= s_5 + 4 &= 20.034 \\
 s_7 &= \text{sqr}(s_2) &= 64 \\
 s_8 &= 3 \times s_7 &= 192 \\
 s_9 &= s_8 + 6 &= 198 \\
 s_{10} &= s_6 \times s_9 &= 3966.732
 \end{aligned}$$

$$\begin{aligned}
 s'_1 &= 0 \\
 s'_2 &= 1 \\
 s'_3 &= s_1 \times s'_2 + s'_1 \times s_2 = 2 \\
 s'_4 &= \cos(s_1) \times s'_1 = 0 \\
 s'_5 &= s'_3 + s'_4 = 2 \\
 s'_6 &= s'_5 = 2 \\
 s'_7 &= s'_2 \times 2 \times s_2 = 16 \\
 s'_8 &= 3 \times s'_7 = 48 \\
 s'_9 &= s'_8 = 48 \\
 s'_{10} &= s_6 \times s'_9 + s'_6 \times s_9 = 1357.632
 \end{aligned}$$

Question: What is the problem with this approach for **neural networks**?

Forward Mode Performance

- To compute the full gradient of a function of m inputs requires computing m partial derivatives
- In forward mode, this requires m forward passes
- For our toy examples, that means running the forward pass twice
- Neural networks can easily have **thousands** of parameters
- We don't want to run the network *thousands of times* for each gradient update!

Automatic Differentiation: Backward Mode

- **Forward mode** sweeps through the graph:

- For each s_i , computes $s'_i = \frac{\partial s_i}{\partial s_1}$ for each s_i

- The **numerator varies**, and the **denominator is fixed**

- **Backward mode** does the opposite:

- For each s_i , computes the **local gradient** $\bar{s}_i = \frac{\partial s_n}{\partial s_i}$

- The **numerator is fixed**, and the **denominator varies**


- At the end, we have computed $\bar{x}_i = \frac{\partial s_n}{\partial x_i}$ for each input x_i

$$\begin{array}{l}
 s_1 = x \\
 s_2 = y \\
 s_3 = s_1 \times s_2 \\
 s_4 = \sin(s_1) \\
 s_5 = s_3 + s_4 \\
 s_6 = s_5 + 4 \\
 s_7 = \text{sqr}(s_2) \\
 s_8 = 3 \times s_7 \\
 s_9 = s_8 + 6 \\
 s_{10} = s_6 \times s_9
 \end{array}
 \left. \vphantom{\begin{array}{l} s_1 \\ s_2 \\ s_3 \\ s_4 \\ s_5 \\ s_6 \\ s_7 \\ s_8 \\ s_9 \\ s_{10} \end{array}} \right\} \frac{\partial s_3}{\partial s_1} \left. \vphantom{\begin{array}{l} \frac{\partial s_3}{\partial s_1} \\ s_3 \\ s_4 \\ s_5 \\ s_6 \\ s_7 \\ s_8 \\ s_9 \\ s_{10} \end{array}} \right\} \frac{\partial s_4}{\partial s_1}$$

$$\frac{\partial s_{10}}{\partial s_7} \left\{ \begin{array}{l} \frac{\partial s_{10}}{\partial s_8} \left\{ \begin{array}{l} s_8 = 3 \times s_7 \\ s_9 = s_8 + 6 \\ s_{10} = s_6 \times s_9 \end{array} \right. \\ s_7 = \text{sqr}(s_2) \end{array} \right.$$

Backward Mode Example

Let's compute $\left. \frac{\partial f}{\partial x} \right|_{x=2, y=8}$ and $\left. \frac{\partial f}{\partial y} \right|_{x=2, y=8}$ using backward mode:




$$\begin{aligned}
 s_1 &= x &&= 2 \\
 s_2 &= y &&= 8 \\
 s_3 &= s_1 \times s_2 &&= 16 \\
 s_4 &= \sin(s_1) &&\approx 0.034 \\
 s_5 &= s_3 + s_4 &&= 16.034 \\
 s_6 &= s_5 + 4 &&= 20.034 \\
 s_7 &= \text{sqr}(s_2) &&= 64 \\
 s_8 &= 3 \times s_7 &&= 192 \\
 s_9 &= s_8 + 6 &&= 198 \\
 s_{10} &= s_6 \times s_9 &&= 3966.732
 \end{aligned}$$

$$\begin{aligned}
 \bar{s}_{10} &= \frac{\partial s_{10}}{\partial s_{10}} = 1 \\
 \bar{s}_9 &= \frac{\partial s_{10}}{\partial s_9} = \frac{\partial s_{10}}{\partial s_{10}} \frac{\partial s_{10}}{\partial s_9} = \bar{s}_{10} s_6 = 20.034 \\
 \bar{s}_8 &= \frac{\partial s_{10}}{\partial s_8} = \frac{\partial s_{10}}{\partial s_9} \frac{\partial s_9}{\partial s_8} = \bar{s}_9 1 = 20.034 \\
 \bar{s}_7 &= \frac{\partial s_{10}}{\partial s_7} = \frac{\partial s_{10}}{\partial s_8} \frac{\partial s_8}{\partial s_7} = \bar{s}_8 3 = 60.102 \\
 \bar{s}_6 &= \frac{\partial s_{10}}{\partial s_6} = s_9 = 198 \\
 &\vdots
 \end{aligned}$$

Backward Mode Example (2)

Let's compute $\left. \frac{\partial f}{\partial x} \right|_{x=2, y=8}$ and $\left. \frac{\partial f}{\partial y} \right|_{x=2, y=8}$ using backward mode:



$$\begin{aligned}
 s_1 &= x &&= 2 \\
 s_2 &= y &&= 8 \\
 s_3 &= s_1 \times s_2 &&= 16 \\
 s_4 &= \sin(s_1) &&\approx 0.034 \\
 s_5 &= s_3 + s_4 &&= 16.034 \\
 s_6 &= s_5 + 4 &&= 20.034 \\
 s_7 &= \text{sqr}(s_2) &&= 64 \\
 s_8 &= 3 \times s_7 &&= 192 \\
 s_9 &= s_8 + 6 &&= 198 \\
 s_{10} &= s_6 \times s_9 &&= 3966.732
 \end{aligned}$$

$$\begin{aligned}
 \bar{s}_6 &= \frac{\partial s_{10}}{\partial s_6} = s_9 = 198 \\
 \bar{s}_5 &= \frac{\partial s_{10}}{\partial s_5} = \frac{\partial s_{10}}{\partial s_6} \frac{\partial s_6}{\partial s_5} = \bar{s}_6 \cdot 1 = 198 \\
 \bar{s}_4 &= \frac{\partial s_{10}}{\partial s_4} = \frac{\partial s_{10}}{\partial s_5} \frac{\partial s_5}{\partial s_4} = \bar{s}_5 \cdot 1 = 198 \\
 \bar{s}_3 &= \frac{\partial s_{10}}{\partial s_3} = \frac{\partial s_{10}}{\partial s_5} \frac{\partial s_5}{\partial s_3} = \bar{s}_5 \cdot 1 = 198 \\
 \bar{s}_2 &= \frac{\partial s_{10}}{\partial s_2} = \frac{\partial s_{10}}{\partial s_3} \frac{\partial s_3}{\partial s_2} + \frac{\partial s_{10}}{\partial s_7} \frac{\partial s_7}{\partial s_2} = \bar{s}_3 s_1 + \bar{s}_7 2s_2 \approx 1357.632 \\
 \bar{s}_1 &= \frac{\partial s_{10}}{\partial s_1} = \frac{\partial s_{10}}{\partial s_3} \frac{\partial s_3}{\partial s_1} + \frac{\partial s_{10}}{\partial s_4} \frac{\partial s_4}{\partial s_1} = \bar{s}_3 s_2 + \bar{s}_4 \cos s_1 \approx 1781.9
 \end{aligned}$$

Back-Propagation

$$L(\mathbf{W}, \mathbf{b}) = \sum_i \ell (f(\mathbf{x}^{(i)}; \mathbf{W}, \mathbf{b}), y^{(i)})$$

Back-propagation is simply automatic differentiation in **backward mode**, used to compute the gradient $\nabla_{\mathbf{W}, \mathbf{b}} L$ of the **loss function** with respect to its **parameters** \mathbf{W}, \mathbf{b} :

1. At each layer, compute the **local gradients** of the layer's computations
2. These local gradients will be used as inputs to the **next layer's** local gradient computations
3. At the end, we have a partial derivative for each of the parameters, which we can use to take a **gradient step**

Summary

- The loss function of a **deep feedforward networks** is simply a very nested function of the **parameters** of the model
- **Automatic differentiation** can compute these gradients more efficiently than symbolic differentiation or finite-differences numeric computations
 - Symbolic differentiation is **interleaved** with numeric computation
 - In **forward mode**, m sweeps are required for a function of m parameters
 - In **backward mode**, only a single sweep is required
- **Back-propagation** is simply automatic differentiation **applied to neural networks** in backward mode