

CS 4803 / 7643: Deep Learning

Topics:

- (Finish) Computing Gradients
 - Backprop in FC+ReLU NNs
 - Backprop in Conv Layers
- Forward mode vs Reverse mode AD
- Patterns in backprop

Zsolt Kira
Georgia Tech

Administrivia

- HW1 Reminder
 - Due: 02/15, 11:55pm

Matrix/Vector Derivatives Notation

	S	V	M	
S	$\frac{\partial y}{\partial x}$	$\frac{\partial y}{\partial \vec{x}}$	$\frac{\partial y}{\partial X}$	$x, y \in \mathbb{R}^1$
V	$\frac{\partial \vec{y}}{\partial x}$	$\frac{\partial \vec{y}}{\partial \vec{x}}$		$\vec{x} \in \mathbb{R}^d$
M	$\frac{\partial Y}{\partial x}$		Tens ²	$\vec{y} \in \mathbb{R}^k$

Convention: $\frac{\partial \vec{y}}{\partial x} = \begin{bmatrix} \frac{\partial y_1}{\partial x} \\ \vdots \\ \frac{\partial y_k}{\partial x} \end{bmatrix}$ \downarrow numerator = dim 1 = col-vector

[Gradient] $\frac{\partial y}{\partial \vec{x}} = \left[\frac{\partial y}{\partial x_1} \dots \frac{\partial y}{\partial x_d} \right]$ denominator = dim 2 = row-vector

→ [More concrete notation for scalars]
 $L'(x) = f'(g(x)) g'(x)$

→ [with intermediate variables]
 $y = g(x)$
 $z = f(y)$


$$\frac{\partial z}{\partial x} = \frac{\partial z}{\partial y} \cdot \frac{\partial y}{\partial x}$$

Example: $L_i(w) = -\log\left(\frac{1}{1+e^{-w \cdot x_i}}\right)$ [For $y_i = +1$]

$$= \left(\underbrace{-\log(\cdot)}_{\frac{\partial L}{\partial p}} \circ \underbrace{\frac{1}{1+e^{-a}}}_{\frac{\partial p}{\partial a}} \circ \underbrace{x^T(w)}_{\frac{\partial a}{\partial w}} \right) (w)$$

$$\frac{\partial L_i}{\partial w} = \begin{bmatrix} 1 \\ p \end{bmatrix} \cdot \begin{bmatrix} -1 & -e^{-a} \\ (1+e^{-a})^2 & p \cdot (1-p) \end{bmatrix} \cdot x^T = (1-p)x^T$$

← Path 1 : $\frac{\partial E}{\partial x_{i-1}}$ from $\frac{\partial E}{\partial x_i}$



w_i

An upward-pointing arrow originates from the label w_i and points to the bottom edge of the function box.

Path 2
 $\frac{\partial E}{\partial w_i}$ from $\frac{\partial E}{\partial x_i}$

A downward-pointing arrow originates from the text 'Path 2' and points to the label w_i .