# calc_intro_corrections

January 21, 2019

## 0.1 Multivariable Calculus Review

We will be covering some of the most relevant concepts from multivariable calculus for this course, and show how they can be extended to deal with matrices of training data.

### 0.1.1 Partial Derivatives

The derivative of a function of 2 variables $f(x, y)$ w.r.t. either one of its variables is defined as

$$\frac{\partial}{\partial x} f(x, y) := \lim_{h \to 0} \frac{f(x + h, y) - f(x, y)}{h}$$

$$\frac{\partial}{\partial y} f(x, y) := \lim_{h \to 0} \frac{f(x, y + h) - f(x, y)}{h}$$

This is a very obvious extension going from the derivative definition for functions of single variables. As an example, what is $\frac{\partial}{\partial x} f(x, y)$ where $f(x, y) = x^2 cos(xy)$?

$$\frac{\partial}{\partial x} f(x, y) = 2x cos(xy) - x^2 sin(xy) y$$

---

### 0.1.2 Directional Derivatives

Before we define the gradient, we discuss the closely related concept of a directional derivative.

The directional derivative of a function $f(x, y)$ at $(x_0, y_0)$ moving in the direction of a unit vector $u = [a, b]$ is

$$D_u f(x_0, y_0) := \lim_{h \to 0} \frac{f(x_0 + ha, y_0 + hb) - f(x_0, y_0)}{h}$$

This simply tells us how the scalar output of $f$ will change when we move in an arbitrary direction that is not necesarily axis-aligned. It can be written as

$$D_u f(x, y) = \left[ \frac{\partial}{\partial x} f(x, y), \frac{\partial}{\partial y} f(x, y) \right] \cdot u$$

---

### 0.1.3 Gradient

This is perhaps the term you will hear most often in the context of neural networks. The gradient is informally defined as the vector pointing in the direction of steepest increase for a given function. Thus, if we would like to minimize a loss function such as mean squared error (MSE), we move in the opposite direction of the gradient, i.e. in the direction of the negative gradient.

Now the gradient is simply defined as

$$\nabla f(x,y) := \left[\frac{\partial}{\partial x} f(x,y), \frac{\partial}{\partial y} f(x,y)\right]$$

Note that this is a vector-field (vector-valued function), and not a scalar-valued function!

What does $\nabla f(x,y)$ have to do with neural networks and minimizing loss functions? Let's see in which direction the directional derivative achieves the highest possible value. Using the gradient notation, we can rewrite the directional derivative as

$$\begin{aligned} D_u f(x,y) &= \nabla f(x,y) \cdot u \\ &= ||\nabla f(x,y)||_2 \cdot ||u||_2 \cdot cos(\theta) \\ &= ||\nabla f(x,y)||_2 \cdot cos(\theta) \end{aligned}$$

This quantity is maximized when $cos(\theta) = 1$, i.e. $\nabla f(x,y)$ and $u$ are parallel. Therefore, the gradient points in the direction of steepest increase for $f$.

So if $f$ is a neural network parameterized by some weights $w$, and $x \in R^D$ is some input to it, then taking steps in $\nabla_w f_w(x)$ will move us in the direction that maximizes the output of $f$. Thus, if $f$ is a binary classification network, say the probability of an image being of a car, then moving in this direction would maximimally increase the probability of $x$ being classified as a car.

---

## 0.2 Basic ML Setup

If you recall from CSC411, the basic supervised learning scenarios are regression and classification. Most buzz in computer vision has driven by success on classifcation problems, so we start with that first.

### 0.2.1 Classification

In $K$-class classification, we are given a dataset $\{\mathbf{x}^{(i)}, t^{(i)}\}_{i=1}^N$ of $N$ samples where $\mathbf{x}^{(i)} \in R^D$ and $t^{(i)} \in [1, ..., K]$.

We would like to fit a probabilistic model, in our case a neural network, that maximizes the probability

$$\prod_{i=1}^N p(t^{(i)}|\mathbf{x}^{(i)}, \mathbf{w})$$

we can vectorize these objects/terms by defining

$$\mathbf{X} = \begin{bmatrix} (\mathbf{x}^{(1)})^\top \\ \vdots \\ (\mathbf{x}^{(N)})^\top \end{bmatrix}$$

$$\mathbf{T} = \begin{bmatrix} (t^{(i)})^\top \\ \vdots \\ (t^{(N)})^\top \end{bmatrix}$$

where $\mathbf{X}$ is called the design matrix, and has dimensions $(N, D)$, where entry $\mathbf{X}_{ij}$ corresponds to the j-th feature for the i-th training sample. $\mathbf{T}$ is our matrix of labels. The i-th row of $\mathbf{T}$ is the one-hot encoded vector of the label for the i-th sample which is simply a $K$-dimensional vector of all 0s except for the a single 1 in the position of the label. If $K = 4$, and $t^{(i)} = 2$, the one-hot encoded version is

$$\begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix}$$

We assume that all vectors are column vectors.

Our goal becomes to maximize the probability

$$p(\mathbf{t}|\mathbf{X}, \mathbf{w}) = \prod_{i=1}^{N} p(t^{(i)}|\mathbf{x}^{(i)}, \mathbf{w})$$

$$= \prod_{i=1}^{N} \prod_{k=1}^{K} p(k|\mathbf{x}^{(i)}, \mathbf{w})^{\mathbf{T}_{ik}}$$

Note that just the term in the innermost product corresponding to the true label of the i-th can be less than 1. Since this expression would lead to numerical underflow given a large enough dataset, we can equivalently maximize a monotonically increasing function of this expression, i.e. take the log

$$\log \left( \prod_{i=1}^{N} \prod_{k=1}^{K} p(k|\mathbf{x}^{(i)}, \mathbf{w})^{\mathbf{T}_{ik}} \right) = \sum_{i=1}^{N} \sum_{k=1}^{K} \mathbf{T}_{ik} \log(p(k|\mathbf{x}^{(i)}, \mathbf{w}))$$

If we negate this term, we end up with the most popular loss function for classification known as cross-entropy

$$L_{CE} = - \sum_{i=1}^{N} \sum_{k=1}^{K} \mathbf{T}_{ik} \log(p(k|\mathbf{x}^{(i)}, \mathbf{w}))$$

Our goal is to **minimize** $L_{CE}$ by moving in the direction $-\nabla_{\mathbf{w}} L_{CE}$.

---

### 0.2.2 Linear Regression

**Multivariate input. Univariate output.** $R^M \to R$

**Forward pass** Suppose we have some features $x_1, x_2, \ldots, x_M$, and we predict a single scalar $y = w_1 x_1 + w_2 x_2 + \cdots + w_M x_M$. We can write this in matrix notation as a dot product

$$y = \sum_{i=1}^{M} w_i x_i = \mathbf{w}^T \mathbf{x} \quad \text{where} \quad \mathbf{w} = \begin{bmatrix} w_1 \\ w_2 \\ \vdots \\ w_M \end{bmatrix}, \mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_M \end{bmatrix}$$

We say $\mathbf{x} \in R^M$ to indicate that $\mathbf{x}$ is a $M$-dimensional vector.

**Backward pass** Suppose we have a target $t$ and a loss function $L = (y - t)^2$, what is the gradient $\nabla_{\mathbf{w}} L$?

Remember the definition of the gradient,

$$\nabla_w L = \begin{bmatrix} \frac{\partial L}{\partial w_1} \\ \frac{\partial L}{\partial w_2} \\ \vdots \\ \frac{\partial L}{\partial w_M} \end{bmatrix}$$

Each element is a partial derivative that can be written out using *chain rule*.

$$\frac{\partial L}{\partial w_i} = \frac{\partial L}{\partial y} \frac{\partial y}{\partial w_i} = 2(y - t) \, x_i$$

So the gradient is

$$\nabla_w L = \begin{bmatrix} 2(y - t) \, x_1 \\ \vdots \\ 2(y - t) \, x_M \end{bmatrix} = 2(y - t)\mathbf{x}$$

**Remark.** There are some matrix calculus formulas that are worth remembering. One that we could have used here is

$$\mathbf{w}^T \mathbf{x} \implies \nabla_{\mathbf{w}} \mathbf{w}^T \mathbf{x} = \mathbf{x}$$

So the math becomes much easier.

$$\nabla_{\mathbf{w}} L = \underbrace{(\nabla_y L)(\nabla_{\mathbf{w}} y)}_{\text{Chain rule holds for matrix calculus.}} = 2(y - t)\mathbf{x}$$

---

**Multivariate input. Multivariate output.** $R^D \rightarrow R^M$

**Forward pass** Suppose we have some features $x_1, x_2, \ldots, x_D$, and we predict a **vector** $\mathbf{y} \in \mathbf{R^M}$ where $y_j = w_{j1} x_1 + w_{j2} x_2 + \cdots + w_{jD} x_D$. We can write this in matrix notation as a **matrix-vector product**.

$$\mathbf{y} = \begin{bmatrix} w_{11} x_1, w_{12} x_2, \cdots w_{1D} x_D \\ w_{21} x_1, w_{22} x_2, \cdots w_{2D} x_D \\ \vdots \\ w_{M1} x_1, w_{M2} x_2, \cdots w_{MD} x_D \end{bmatrix} = \begin{bmatrix} w_1^T x \\ w_2^T x \\ \vdots \\ w_M^T x \end{bmatrix} = \mathbf{W}\mathbf{x} \quad \text{where} \quad W = \begin{bmatrix} w_1^T \\ w_2^T \\ \vdots \\ w_M^T \end{bmatrix} = \begin{bmatrix} w_{11} & w_{12} & \cdots & w_{1D} \\ w_{21} & w_{22} & \cdots & w_{2D} \\ \vdots \\ w_{M1} & w_{M2} & \cdots & w_{MD} \end{bmatrix}$$

We say $\mathbf{W} \in R^{M \times D}$ to indicate that $\mathbf{W}$ is a matrix of dimension $M$ by $D$.