Computing Neural Network Gradients

Kevin Clark

1   Introduction

The purpose of these notes is to demonstrate how to quickly compute neural network gradients. This will hopefully help you with question 3 of Assignment 2 (if you haven't already done it) and with the midterm (which will have at least one significant gradient computation question). It is **not** meant to provide an intuition for how backpropagation works – for that I recommend going over lecture 5[1] and the cs231 course notes[2] on backpropagation.

2   Vectorized Gradients

While it is a good exercise to compute the gradient of a neural network with respect to a single parameter (e.g., a single element in a weight matrix), in practice this tends to be quite slow. Instead, it is more efficient to keep everything in matrix/vector form. The basic building block of vectorized gradients is the *Jacobian Matrix*. Suppose we have a function $\boldsymbol{f} : \mathbb{R}^n \to \mathbb{R}^m$ that maps a vector of length n to a vector of length m: $\boldsymbol{f}(\boldsymbol{x}) = [f_1(x_1, ..., x_n), f_2(x_1, ..., x_n), ..., f_m(x_1, ..., x_n)]$. Then its Jacobian is.

$$\frac{\partial \boldsymbol{f}}{\partial \boldsymbol{x}} = \begin{bmatrix} \frac{\partial f_1}{\partial x_1} & \cdots & \frac{\partial f_1}{\partial x_n} \\ \vdots & \ddots & \vdots \\ \frac{\partial f_m}{\partial x_1} & \cdots & \frac{\partial f_m}{\partial x_n} \end{bmatrix}$$

That is, $(\frac{\partial \boldsymbol{f}}{\partial \boldsymbol{x}})_{ij} = \frac{\partial f_i}{\partial x_j}$ (which is just a standard non-vector derivative). The Jacobian matrix will be useful for us because we can apply the chain rule to a vector-valued function just by multiplying Jacobians.

As a little illustration of this, suppose we have a function $\boldsymbol{f}(x) = [f_1(x), f_2(x)]$ taking a scalar to a vector of size 2 and a function $\boldsymbol{g}(\boldsymbol{y}) = [g_1(y_1, y_2), g_2(y_1, y_2)]$ taking a vector of size two to a vector of size two. Now let's compose them to get $\boldsymbol{g}(x) = [g_1(f_1(x), f_2(x)), g_2(f_1(x), f_2(x))]$. Using the regular chain rule, we

---

[1] http://web.stanford.edu/class/cs224n/lectures/cs224n-2017-lecture5.pdf
[2] http://cs231n.github.io/optimization-2/

can compute the derivative of $\boldsymbol{g}$ as the Jacobian

$$\frac{\partial \boldsymbol{g}}{\partial x} = \begin{bmatrix} \frac{\partial}{\partial x} g_1(f_1(x), f_2(x)) \\ \frac{\partial}{\partial x} g_2(f_1(x), f_2(x)) \end{bmatrix} = \begin{bmatrix} \frac{\partial g_1}{\partial f_1} \frac{\partial f_1}{\partial x} + \frac{\partial g_1}{\partial f_2} \frac{\partial f_2}{\partial x} \\ \frac{\partial g_2}{\partial f_1} \frac{\partial f_1}{\partial x} + \frac{\partial g_2}{\partial f_2} \frac{\partial f_2}{\partial x} \end{bmatrix}$$

And we see this is the same as multiplying the two Jacobians:

$$\frac{\partial \boldsymbol{g}}{\partial x} = \frac{\partial \boldsymbol{g}}{\partial \boldsymbol{f}} \frac{\partial \boldsymbol{f}}{\partial x} = \begin{bmatrix} \frac{\partial g_1}{\partial f_1} & \frac{\partial g_1}{\partial f_2} \\ \frac{\partial g_2}{\partial f_1} & \frac{\partial g_2}{\partial f_2} \end{bmatrix} \begin{bmatrix} \frac{\partial f_1}{\partial x} \\ \frac{\partial f_2}{\partial x} \end{bmatrix}$$

3 Useful Identities

This section will now go over how to compute the Jacobian for several simple functions. It will provide some useful identities you can apply when taking neural network gradients.

(1) **Matrix times column vector with respect to the column vector**
($\boldsymbol{z} = \boldsymbol{W}\boldsymbol{x}$, what is $\frac{\partial \boldsymbol{z}}{\partial \boldsymbol{x}}$?)

Suppose $\boldsymbol{W} \in \mathbb{R}^{n \times m}$. Then we can think of $\boldsymbol{z}$ as a function of $\boldsymbol{x}$ taking an m-dimensional vector to an n-dimensional vector. So its Jacobian will be $n \times m$. Note that

$$z_i = \sum_{k=1}^{m} W_{ik} x_k$$

So an entry $(\frac{\partial \boldsymbol{z}}{\partial \boldsymbol{x}})_{ij}$ of the Jacobian will be

$$(\frac{\partial \boldsymbol{z}}{\partial \boldsymbol{x}})_{ij} = \frac{\partial z_i}{\partial x_j} = \frac{\partial}{\partial x_j} \sum_{k=1}^{m} W_{ik} x_k = \sum_{k=1}^{m} W_{ik} \frac{\partial}{\partial x_j} x_k = W_{ij}$$

because $\frac{\partial}{\partial x_j} x_k = 1$ if $k = j$ and 0 if otherwise. So we see that $\boxed{\dfrac{\partial \boldsymbol{z}}{\partial \boldsymbol{x}} = \boldsymbol{W}}$

(2) **Row vector times matrix with respect to the row vector**
($\boldsymbol{z} = \boldsymbol{x}\boldsymbol{W}$, what is $\frac{\partial \boldsymbol{z}}{\partial \boldsymbol{x}}$?)

A computation similar to (1) shows that $\boxed{\dfrac{\partial \boldsymbol{z}}{\partial \boldsymbol{x}} = \boldsymbol{W}^T}$.

(3) **A vector with itself**
($\boldsymbol{z} = \boldsymbol{x}$, what is $\frac{\partial \boldsymbol{z}}{\partial \boldsymbol{x}}$?)
We have $z_i = x_i$. So

$$(\frac{\partial \boldsymbol{z}}{\partial \boldsymbol{x}})_{ij} = \frac{\partial z_i}{\partial x_j} = \frac{\partial}{\partial x_j} x_i = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{if otherwise} \end{cases}$$

So we see that the Jacobian $\frac{\partial \boldsymbol{z}}{\partial \boldsymbol{x}}$ is a diagonal matrix where the entry at (i, i) is 1. This is just the identity matrix: $\boxed{\dfrac{\partial \boldsymbol{z}}{\partial \boldsymbol{x}} = \boldsymbol{I}}$. When applying the chain rule, this term will disappear because a matrix multiplied by the identity matrix does not change.

(4) **An elementwise function applied a vector**
($\boldsymbol{z} = f(\boldsymbol{x})$, what is $\frac{\partial \boldsymbol{z}}{\partial \boldsymbol{x}}$?)
If f is being applied elementwise, we have $z_i = f(x_i)$. So

$$(\frac{\partial \boldsymbol{z}}{\partial \boldsymbol{x}})_{ij} = \frac{\partial z_i}{\partial x_j} = \frac{\partial}{\partial x_j} f(x_i) = \begin{cases} f'(x_i) & \text{if } i = j \\ 0 & \text{if otherwise} \end{cases}$$

So we see that the Jacobian $\frac{\partial \boldsymbol{z}}{\partial \boldsymbol{x}}$ is a diagonal matrix where the entry at (i, i) is the derivative of f applied to x_i. We can write this as $\boxed{\dfrac{\partial \boldsymbol{z}}{\partial \boldsymbol{x}} = \text{diag}(f'(\boldsymbol{x}))}$. Since multiplication by a diagonal matrix is the same as doing elementwise multiplication by the diagonal, we could also write $\boxed{\circ f'(\boldsymbol{x})}$ when applying the chain rule.

(5) **Matrix times column vector with respect to the matrix**
($\boldsymbol{z} = \boldsymbol{W}\boldsymbol{x}$, $\boldsymbol{\delta} = \frac{\partial J}{\partial \boldsymbol{z}}$ what is $\frac{\partial J}{\partial \boldsymbol{W}} = \frac{\partial J}{\partial \boldsymbol{z}} \frac{\partial \boldsymbol{z}}{\partial \boldsymbol{W}} = \boldsymbol{\delta} \frac{\partial \boldsymbol{z}}{\partial \boldsymbol{W}}$?)

This is a bit more complicated than the other identities. The reason for including $\frac{\partial J}{\partial \boldsymbol{z}}$ in the above problem formulation will become clear in a moment.

First suppose we have a loss function J (a scalar) and are computing its gradient with respect to a matrix $\boldsymbol{W} \in \mathbb{R}^{n \times m}$. Then we could think of J as a function of $\boldsymbol{W}$ taking nm inputs (the entries of $\boldsymbol{W}$) to a single output (J). This means the Jacobian $\frac{\partial J}{\partial \boldsymbol{W}}$ would be a $1 \times nm$ vector. But in practice this is not a very useful way of arranging the gradient. It would be much nicer if the derivatives were in a $n \times m$ matrix like this:

$$\frac{\partial J}{\partial \boldsymbol{W}} = \begin{bmatrix} \frac{\partial J}{\partial W_{11}} & \cdots & \frac{\partial J}{\partial W_{1m}} \\ \vdots & \ddots & \vdots \\ \frac{\partial J}{\partial W_{n1}} & \cdots & \frac{\partial J}{\partial W_{nm}} \end{bmatrix}$$

Since this matrix has the same shape as $\boldsymbol{W}$, we could just subtract it (times the learning rate) from $\boldsymbol{W}$ when doing gradient descent. So (in a slight abuse of notation) let's find this matrix as $\frac{\partial J}{\partial \boldsymbol{W}}$ instead.

This way of arranging the gradients becomes complicated when computing $\frac{\partial \boldsymbol{z}}{\partial \boldsymbol{W}}$. Unlike J, $\boldsymbol{z}$ is a vector. So if we are trying to rearrange the gradients like with $\frac{\partial J}{\partial \boldsymbol{W}}$, $\frac{\partial \boldsymbol{z}}{\partial \boldsymbol{W}}$ would be an $n \times m \times n$ tensor! Luckily, we can avoid the issue by taking the gradient with respect to a single weight W_{ij} instead.

$\frac{\partial z}{\partial W_{ij}}$ is just a vector, which is much easier to deal with. We have

$$z_k = \sum_{l=1}^{m} W_{kl} x_l$$

$$\frac{\partial z_k}{\partial W_{ij}} = \sum_{l=1}^{m} x_l \frac{\partial}{\partial W_{ij}} W_{kl}$$

Note that $\frac{\partial}{\partial W_{ij}} W_{kl} = 1$ if $i = k$ and $j = l$ and 0 if otherwise. So if $k \neq i$ everything in the sum is zero and the gradient is zero. Otherwise, the only nonzero element of the sum is when $l = j$, so we just get x_j. Thus we find $\frac{\partial z_k}{\partial W_{ij}} = x_j$ if $k = i$ and 0 if otherwise. Another way of writing this is

$$\frac{\partial z}{\partial W_{ij}} = \begin{bmatrix} 0 \\ \vdots \\ 0 \\ x_j \\ 0 \\ \vdots \\ 0 \end{bmatrix} \leftarrow i\text{th element}$$

Now let's compute $\frac{\partial J}{\partial W_{ij}}$

$$\frac{\partial J}{\partial W_{ij}} = \frac{\partial J}{\partial z} \frac{\partial z}{\partial W_{ij}} = \boldsymbol{\delta} \frac{\partial z}{\partial W_{ij}} = \sum_{k=1}^{m} \delta_k \frac{\partial z_k}{\partial W_{ij}} = \delta_i x_j$$

(the only nonzero term in the sum is $\delta_i \frac{\partial z_i}{\partial W_{ij}}$). To get $\frac{\partial J}{\partial W}$ we want a matrix where entry (i, j) is $\delta_i x_j$. This matrix is equal to the outer product

$$\boxed{\frac{\partial J}{\partial W} = \boldsymbol{\delta}^T x}$$

(6) **Row vector time matrix with respect to the matrix**
$(z = xW, \boldsymbol{\delta} = \frac{\partial J}{\partial z}$ what is $\frac{\partial J}{\partial W} = \boldsymbol{\delta} \frac{\partial z}{\partial W}$?)

A similar computation to (5) shows that $\boxed{\dfrac{\partial J}{\partial W} = x^T \boldsymbol{\delta}}$.

(7) **Cross-entropy loss with respect to logits** $(\hat{y} = \text{softmax}(\boldsymbol{\theta}), J = CE(\boldsymbol{y}, \hat{\boldsymbol{y}})$, what is $\frac{\partial J}{\partial \boldsymbol{\theta}}$?)

You computed this in Assignment 1! The gradient is $\boxed{\dfrac{\partial J}{\partial \boldsymbol{\theta}} = \hat{\boldsymbol{y}} - \boldsymbol{y}}$

(or $(\hat{\boldsymbol{y}} - \boldsymbol{y})^T$ if $\boldsymbol{y}$ is a column vector).

These identities will be enough to let you quickly compute the gradients for many neural networks. However, it's important to know how to compute Jacobians for other functions as well in case they show up. Some examples if you want practice: dot product of two vectors, elementwise product of two vectors, 2-norm of a vector. Feel free to use these identities on the midterm and assignments.

4 Example: 1-Layer Neural Network with Embeddings

This section provides an example of computing the gradients of a full neural network. In particular we are going to compute the gradients of the dependency parser you are building in Assignment 2. First let's write out the forward pass of the model.

$$\boldsymbol{x} = [\boldsymbol{L}_{w_0}, \boldsymbol{L}_{w_1}, ..., \boldsymbol{L}_{w_{m-1}}]$$
$$\boldsymbol{z} = \boldsymbol{x}\boldsymbol{W} + \boldsymbol{b}_1$$
$$\boldsymbol{h} = \text{ReLU}(\boldsymbol{z})$$
$$\boldsymbol{\theta} = \boldsymbol{h}\boldsymbol{U} + \boldsymbol{b}_2$$
$$\boldsymbol{\hat{y}} = \text{softmax}(\boldsymbol{\theta})$$
$$J = CE(\boldsymbol{y}, \boldsymbol{\hat{y}})$$

It helps to break up the model into the simplest parts possible, so note that we defined $\boldsymbol{z}$ and $\boldsymbol{\theta}$ to split up the activation functions from the linear transformations in the network's layers. The dimensions of the model's parameters are

$$\boldsymbol{L} \in \mathbb{R}^{|V| \times d} \qquad \boldsymbol{b}_1 \in \mathbb{R}^{1 \times D_h} \qquad \boldsymbol{W} \in \mathbb{R}^{md \times D_h} \qquad \boldsymbol{b}_2 \in \mathbb{R}^{1 \times N_c} \qquad \boldsymbol{U} \in \mathbb{R}^{D_h \times N_c}$$

where $|V|$ is the vocabulary size, d is the size of our word vectors, m is the number of features, D_h is the size of our hidden layer, and N_c is the number of classes.

In this example, we will compute all the gradients:

$$\frac{\partial J}{\partial \boldsymbol{U}} \qquad \frac{\partial J}{\partial \boldsymbol{b}_2} \qquad \frac{\partial J}{\partial \boldsymbol{W}} \qquad \frac{\partial J}{\partial \boldsymbol{b}_1} \qquad \frac{\partial J}{\partial \boldsymbol{L}_{w_i}}$$

To start with, recall that $\text{ReLU}(x) = \max(x, 0)$. This means

$$\text{ReLU}'(x) = \begin{cases} 1 & \text{if } x > 0 \\ 0 & \text{if otherwise} \end{cases} = \text{sgn}(\text{ReLU}(x))$$

where sgn is the signum function. Note that as you did in Assignment 1 with sigmoid, we are able to write the derivative of the activation in terms of the activation itself.

Now let's write out the chain rule for $\frac{\partial J}{\partial \boldsymbol{U}}$ and $\frac{\partial J}{\partial \boldsymbol{b}_2}$:

$$\frac{\partial J}{\partial \boldsymbol{U}} = \frac{\partial J}{\partial \hat{\boldsymbol{y}}} \frac{\partial \hat{\boldsymbol{y}}}{\partial \boldsymbol{\theta}} \frac{\partial \boldsymbol{\theta}}{\partial \boldsymbol{U}}$$

$$\frac{\partial J}{\partial \boldsymbol{b}_2} = \frac{\partial J}{\partial \hat{\boldsymbol{y}}} \frac{\partial \hat{\boldsymbol{y}}}{\partial \boldsymbol{\theta}} \frac{\partial \boldsymbol{\theta}}{\partial \boldsymbol{b}_2}$$

Notice that $\frac{\partial J}{\partial \hat{\boldsymbol{y}}} \frac{\partial \hat{\boldsymbol{y}}}{\partial \boldsymbol{\theta}} = \frac{\partial J}{\partial \boldsymbol{\theta}}$ is present in both gradients. This makes the math a bit cumbersome. Even worse, if we're implementing the model without automatic differentiation, computing $\frac{\partial J}{\partial \boldsymbol{\theta}}$ twice will be inefficient. So it will help us to define some variables to represent the intermediate derivatives:

$$\boldsymbol{\delta}_1 = \frac{\partial J}{\partial \boldsymbol{\theta}} \qquad \boldsymbol{\delta}_2 = \frac{\partial J}{\partial \boldsymbol{z}}$$

These can be thought as the error signals passed down to $\boldsymbol{\theta}$ and $\boldsymbol{z}$ when doing backpropagation. We can compute them as follows:

$$\boldsymbol{\delta}_1 = \frac{\partial J}{\partial \boldsymbol{\theta}} = \hat{\boldsymbol{y}} - \boldsymbol{y} \qquad\qquad \text{this is just identity (7)}$$

$$\boldsymbol{\delta}_2 = \frac{\partial J}{\partial \boldsymbol{z}} = \frac{\partial J}{\partial \boldsymbol{\theta}} \frac{\partial \boldsymbol{\theta}}{\partial \boldsymbol{h}} \frac{\partial \boldsymbol{h}}{\partial \boldsymbol{z}} \qquad\qquad \text{using the chain rule}$$

$$= \boldsymbol{\delta}_1 \frac{\partial \boldsymbol{\theta}}{\partial \boldsymbol{h}} \frac{\partial \boldsymbol{h}}{\partial \boldsymbol{z}} \qquad\qquad \text{substituting in } \boldsymbol{\delta}_1$$

$$= \boldsymbol{\delta}_1 \, \boldsymbol{U}^T \frac{\partial \boldsymbol{h}}{\partial \boldsymbol{z}} \qquad\qquad \text{using identity (2)}$$

$$= \boldsymbol{\delta}_1 \, \boldsymbol{U}^T \circ \text{ReLU}'(\boldsymbol{z}) \qquad\qquad \text{using identity (4)}$$

$$= \boldsymbol{\delta}_1 \, \boldsymbol{U}^T \circ \text{sgn}(\boldsymbol{h}) \qquad\qquad \text{we computed this earlier}$$

A good way of checking our work is by looking at the dimensions of the terms in the derivative:

$$\begin{array}{ccccc} \dfrac{\partial J}{\partial \boldsymbol{z}} & = & \boldsymbol{\delta}_1 & \boldsymbol{U}^T & \circ & \text{sgn}(\boldsymbol{h}) \\[2mm] (1 \times D_h) & & (1 \times N_c) & (N_c \times D_h) & & (D_h) \end{array}$$

We see that the dimensions of all the terms in the gradient match up (i.e., the number of columns in a term equals the number of rows in the next term). This will always be the case if we computed our gradients correctly.

Now we can use the error terms to compute our gradients:

$$\frac{\partial J}{\partial \boldsymbol{U}} = \frac{\partial J}{\partial \boldsymbol{\theta}}\frac{\partial \boldsymbol{\theta}}{\partial \boldsymbol{U}} = \boldsymbol{\delta}_1 \frac{\partial \boldsymbol{\theta}}{\partial \boldsymbol{U}} = \boldsymbol{h}^T \boldsymbol{\delta}_1 \qquad \text{using identity (6)}$$

$$\frac{\partial J}{\partial \boldsymbol{b}_2} = \frac{\partial J}{\partial \boldsymbol{\theta}}\frac{\partial \boldsymbol{\theta}}{\partial \boldsymbol{b}_2} = \boldsymbol{\delta}_1 \frac{\partial \boldsymbol{\theta}}{\partial \boldsymbol{b}_2} = \boldsymbol{\delta}_1 \qquad \text{using identity (3)}$$

$$\frac{\partial J}{\partial \boldsymbol{W}} = \frac{\partial J}{\partial \boldsymbol{\theta}}\frac{\partial \boldsymbol{z}}{\partial \boldsymbol{W}} = \boldsymbol{\delta}_2 \frac{\partial \boldsymbol{z}}{\partial \boldsymbol{W}} = \boldsymbol{x}^T \boldsymbol{\delta}_2 \qquad \text{using identity (6)}$$

$$\frac{\partial J}{\partial \boldsymbol{b}_1} = \frac{\partial J}{\partial \boldsymbol{\theta}}\frac{\partial \boldsymbol{z}}{\partial \boldsymbol{b}_1} = \boldsymbol{\delta}_2 \frac{\partial \boldsymbol{z}}{\partial \boldsymbol{b}_1} = \boldsymbol{\delta}_2 \qquad \text{using identity (3)}$$

$$\frac{\partial J}{\partial \boldsymbol{L}_{w_i}} = \frac{\partial J}{\partial \boldsymbol{z}}\frac{\partial \boldsymbol{z}}{\partial \boldsymbol{L}_{w_i}} = \boldsymbol{\delta}_2 \frac{\partial \boldsymbol{z}}{\partial \boldsymbol{L}_{w_i}}$$

All that's left is to compute $\frac{\partial \boldsymbol{z}}{\partial \boldsymbol{L}_{w_i}}$. It helps to split up $\boldsymbol{W}$ by rows like this:

$$\boldsymbol{x}\boldsymbol{W} = [\boldsymbol{L}_{w_0}, \boldsymbol{L}_{w_1}, ..., \boldsymbol{L}_{w_{m-1}}]\boldsymbol{W} = [\boldsymbol{L}_{w_0}, \boldsymbol{L}_{w_1}, ..., \boldsymbol{L}_{w_{m-1}}]\begin{bmatrix} \boldsymbol{W}_{0:d} \\ \boldsymbol{W}_{d:2d} \\ \vdots \\ \boldsymbol{W}_{(m-1)d:md} \end{bmatrix}$$

$$= \boldsymbol{L}_{w_0}\boldsymbol{W}_{0:d} + \boldsymbol{L}_{w_1}\boldsymbol{W}_{d:2d} + \cdots + \boldsymbol{L}_{w_{m-1}}\boldsymbol{W}_{(m-1)d:md} = \sum_{j=0}^{m-1} \boldsymbol{L}_{w_j}\boldsymbol{W}_{dj:d(j+1)}$$

When we compute $\frac{\partial \boldsymbol{z}}{\partial \boldsymbol{L}_{w_i}}$, only the ith term in this sum is nonzero, so we get

$$\frac{\partial \boldsymbol{z}}{\partial \boldsymbol{L}_{w_i}} = \frac{\partial}{\partial \boldsymbol{L}_{w_i}} = \boldsymbol{L}_{w_i}\boldsymbol{W}_{di:d(i+1)} = (\boldsymbol{W}_{di:d(i+1)})^T$$

using identity (2).