

Back-propagation Learning in Multilayer Perceptron Artificial Neural Networks

K.M. Chugg, B.K. Jenkins
version 0.7

March 31, 2023

This document derives the back-propagation learning algorithm for Artificial Neural Networks (ANNs). In particular we focus on the case of Multilayer Perceptrons (MLPs) where each layer in the ANN is a fully-connected (or dense) layer. Back-propagation learning is really just the repeated application of the chain rule for derivatives from calculus. Specifically, each parameter w is the ANN is updated by $-\eta \frac{\partial J}{\partial w}$ at each iteration, where J is the criterion, or loss, function. Thus, the entire algorithm comprises the steps and recursions required to calculate this partial derivative for each trainable parameter in the ANN. It is very useful to express these steps/recursions in vector-matrix form so we start with a summary of the required vector calculus results.

1 Vector Calculus Preliminaries

Vector calculus is simply a set of book-keeping conventions for tracking many related partial derivatives in a compact notation. The use of vector calculus notation is entirely optional for deriving and implementing back-propagation learning, but it substantially simplifies both and therefore we adopt it herein.

We are familiar with the gradient of a function $f(\underline{x})$ that maps a vector valued argument to a scalar:

$$\nabla_{\underline{x}} f(\underline{x}) = \begin{bmatrix} \frac{\partial f(\underline{x})}{\partial x_1} \\ \frac{\partial f(\underline{x})}{\partial x_2} \\ \frac{\partial f(\underline{x})}{\partial x_3} \\ \vdots \\ \frac{\partial f(\underline{x})}{\partial x_M} \end{bmatrix} \quad (1)$$

Note that for $\underline{x} \in \mathbb{R}^M$, the gradient $\nabla_{\underline{x}} f(\underline{x}) \in \mathbb{R}^M$ is a column vector.

The gradient can be viewed as the derivative of a scalar valued function with respect to a vector argument. To derive backprop learning, we desire to extend this notion to the derivative of a vector-valued function with respect to a vector argument. Again, we emphasize that this is simply all partial derivatives involved, so it is more accurate to say that we seek a book-keeping convention for storing these partial derivatives. For example consider the vector-valued function

$\underline{f}(\underline{x})$ that maps \mathbb{R}^M into \mathbb{R}^K

$$\underline{f}(\underline{x}) = \begin{bmatrix} f_1(\underline{x}) \\ f_2(\underline{x}) \\ f_3(\underline{x}) \\ \vdots \\ f_K(\underline{x}) \end{bmatrix} \quad (2)$$

where each $f_k(\underline{x})$ maps \mathbb{R}^M into \mathbb{R} . We would like to compute all of the partial derivatives $\left\{ \frac{\partial f_k(\underline{x})}{\partial x_m} \right\}_{k,m}$. These are naturally stored in a matrix which can be thought of as $\frac{\partial \underline{f}(\underline{x})}{\partial \underline{x}}$. However, there are two conventions for doing so. In the *denominator convention* this is considered a $(M \times K)$ matrix with $(m, k)^{\text{th}}$ entry $\frac{\partial f_k(\underline{x})}{\partial x_m}$. In the *numerator convention* this is considered a $(K \times M)$ matrix with $(k, m)^{\text{th}}$ entry $\frac{\partial f_k(\underline{x})}{\partial x_m}$. Clearly, these two conventions are related by a simple transpose. While each of these conventions has desirable properties, **we adopt the denominator convention**.

For this denominator convention, for a scalar-valued function we have

$$\frac{\partial f(\underline{x})}{\partial \underline{x}} = \nabla_{\underline{x}} f(\underline{x}) \quad (3)$$

and for a vector-valued function we have

$$\frac{\partial \underline{f}(\underline{x})}{\partial \underline{x}} = \left[\nabla_{\underline{x}} f_1(\underline{x}) \quad \nabla_{\underline{x}} f_2(\underline{x}) \quad \cdots \quad \nabla_{\underline{x}} f_K(\underline{x}) \right] \quad (4)$$

$$= \begin{bmatrix} \frac{\partial f_1(\underline{x})}{\partial x_1} & \frac{\partial f_2(\underline{x})}{\partial x_1} & \cdots & \frac{\partial f_K(\underline{x})}{\partial x_1} \\ \frac{\partial f_1(\underline{x})}{\partial x_2} & \frac{\partial f_2(\underline{x})}{\partial x_2} & \cdots & \frac{\partial f_K(\underline{x})}{\partial x_2} \\ \frac{\partial f_1(\underline{x})}{\partial x_3} & \frac{\partial f_2(\underline{x})}{\partial x_3} & \cdots & \frac{\partial f_K(\underline{x})}{\partial x_3} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial f_1(\underline{x})}{\partial x_M} & \frac{\partial f_2(\underline{x})}{\partial x_M} & \cdots & \frac{\partial f_K(\underline{x})}{\partial x_M} \end{bmatrix} \quad (5)$$

1.1 Useful Relations

With the denominator convention adopted above, we have

$$\frac{\partial(\underline{A}\underline{x})}{\partial \underline{x}} = \underline{A}^T \quad (6)$$

with special case

$$\frac{\partial(\underline{b}^T \underline{x})}{\partial \underline{x}} = \underline{b} \quad (7)$$

Also, the key identity used in the derivation of backprop is the *right-to-left chain-rule*¹, that is associated with the denominator convention. Specifically, let $\underline{f} : \mathbb{R}^M \rightarrow \mathbb{R}^K$ and $\underline{g} : \mathbb{R}^K \rightarrow \mathbb{R}^L$,

¹In the numerator convention, the chain rule is left-to-right and all the vector derivatives are transposes of those obtained with the denominator convention.

Identities: vector-by-vector $\frac{\partial \mathbf{y}}{\partial \mathbf{x}}$

Condition	Expression	Numerator layout, i.e. by \mathbf{y} and \mathbf{x}^\top	Denominator layout, i.e. by \mathbf{y}^\top and \mathbf{x}
\mathbf{a} is not a function of \mathbf{x}	$\frac{\partial \mathbf{a}}{\partial \mathbf{x}} =$	$\mathbf{0}$	
	$\frac{\partial \mathbf{x}}{\partial \mathbf{x}} =$	\mathbf{I}	
\mathbf{A} is not a function of \mathbf{x}	$\frac{\partial \mathbf{A}\mathbf{x}}{\partial \mathbf{x}} =$	\mathbf{A}	\mathbf{A}^\top
\mathbf{A} is not a function of \mathbf{x}	$\frac{\partial \mathbf{x}^\top \mathbf{A}}{\partial \mathbf{x}} =$	\mathbf{A}^\top	\mathbf{A}
\mathbf{a} is not a function of \mathbf{x} , $\mathbf{u} = \mathbf{u}(\mathbf{x})$	$\frac{\partial \mathbf{a}\mathbf{u}}{\partial \mathbf{x}} =$	$\mathbf{a} \frac{\partial \mathbf{u}}{\partial \mathbf{x}}$	
$\mathbf{v} = \mathbf{v}(\mathbf{x})$, \mathbf{a} is not a function of \mathbf{x}	$\frac{\partial \mathbf{v}\mathbf{a}}{\partial \mathbf{x}} =$	$\mathbf{a} \frac{\partial \mathbf{v}}{\partial \mathbf{x}}$	$\frac{\partial \mathbf{v}}{\partial \mathbf{x}} \mathbf{a}^\top$
$\mathbf{v} = \mathbf{v}(\mathbf{x})$, $\mathbf{u} = \mathbf{u}(\mathbf{x})$	$\frac{\partial \mathbf{v}\mathbf{u}}{\partial \mathbf{x}} =$	$\mathbf{v} \frac{\partial \mathbf{u}}{\partial \mathbf{x}} + \mathbf{u} \frac{\partial \mathbf{v}}{\partial \mathbf{x}}$	$\mathbf{v} \frac{\partial \mathbf{u}}{\partial \mathbf{x}} + \frac{\partial \mathbf{v}}{\partial \mathbf{x}} \mathbf{u}^\top$
\mathbf{A} is not a function of \mathbf{x} , $\mathbf{u} = \mathbf{u}(\mathbf{x})$	$\frac{\partial \mathbf{A}\mathbf{u}}{\partial \mathbf{x}} =$	$\mathbf{A} \frac{\partial \mathbf{u}}{\partial \mathbf{x}}$	$\frac{\partial \mathbf{u}}{\partial \mathbf{x}} \mathbf{A}^\top$
$\mathbf{u} = \mathbf{u}(\mathbf{x})$, $\mathbf{v} = \mathbf{v}(\mathbf{x})$	$\frac{\partial (\mathbf{u} + \mathbf{v})}{\partial \mathbf{x}} =$	$\frac{\partial \mathbf{u}}{\partial \mathbf{x}} + \frac{\partial \mathbf{v}}{\partial \mathbf{x}}$	
$\mathbf{u} = \mathbf{u}(\mathbf{x})$	$\frac{\partial \mathbf{g}(\mathbf{u})}{\partial \mathbf{x}} =$	$\frac{\partial \mathbf{g}(\mathbf{u})}{\partial \mathbf{u}} \frac{\partial \mathbf{u}}{\partial \mathbf{x}}$	$\frac{\partial \mathbf{u}}{\partial \mathbf{x}} \frac{\partial \mathbf{g}(\mathbf{u})}{\partial \mathbf{u}}$
$\mathbf{u} = \mathbf{u}(\mathbf{x})$	$\frac{\partial \mathbf{f}(\mathbf{g}(\mathbf{u}))}{\partial \mathbf{x}} =$	$\frac{\partial \mathbf{f}(\mathbf{g})}{\partial \mathbf{g}} \frac{\partial \mathbf{g}(\mathbf{u})}{\partial \mathbf{u}} \frac{\partial \mathbf{u}}{\partial \mathbf{x}}$	$\frac{\partial \mathbf{u}}{\partial \mathbf{x}} \frac{\partial \mathbf{g}(\mathbf{u})}{\partial \mathbf{u}} \frac{\partial \mathbf{f}(\mathbf{g})}{\partial \mathbf{g}}$

Figure 1: Table of derivatives of vector quantities with respect to vector variables taken from [the Wikipedia page for matrix calculus](#). In this document we adopt the denominator layout convention. This Wikipedia page uses bold lower case for vectors and bold upper case for matrices. The chain rule is summarized in the last two entries.

then

$$\underbrace{\frac{\partial \mathbf{g}(\mathbf{f}(\mathbf{x}))}{\partial \mathbf{x}}}_{(M \times L)} = \underbrace{\frac{\partial \mathbf{f}}{\partial \mathbf{x}}}_{(M \times K)} \underbrace{\frac{\partial \mathbf{g}}{\partial \mathbf{f}}}_{(K \times L)} \tag{8}$$

These book-keeping conventions can be extended to define, for example, the derivative of a vector-valued function with respect to matrix variable. Surprisingly, one of the best references for this is [the Wikipedia page for matrix calculus](#). A table of relations for vector derivatives of vector-valued functions from this page is shown in Figure 1.

1.2 Vectorized Scalar Functions

One special case of the above is a vector-valued function that is a “vectorized scalar function.” This is a vector valued function that is obtained by applying the same scalar-valued function to

each component of a vector argument

$$\underline{f}(\underline{x}) = \mathbf{f}(\underline{x}) = \begin{bmatrix} f(x_1) \\ f(x_2) \\ f(x_3) \\ \vdots \\ f(x_M) \end{bmatrix} \quad (9)$$

where we have introduced the notation $\mathbf{f}(\underline{x})$ to denote this special case of a vector valued function. Note that this is the special case of (2) with $f_k(\underline{x}) = f(x_k)$. For the special case of a vectorized scalar function we have

$$\frac{\partial \mathbf{f}(\underline{x})}{\partial \underline{x}} = \underline{\underline{\text{diag}}} \left(\frac{d}{dx_1} f(x_1), \frac{d}{dx_2} f(x_2), \dots, \frac{d}{dx_M} f(x_M) \right) = \underline{\underline{\text{diag}}} \left(\dot{\mathbf{f}}(\underline{x}) \right) \quad (10)$$

where $\underline{\underline{\text{diag}}}(\cdot)$ indicates a diagonal matrix with the specified diagonal elements and $\dot{\mathbf{f}}(\underline{x})$ is the vectorization of the scalar function $\dot{f}(v) = \frac{df(v)}{dv}$.

As a special case of the chain rule in (8) when $\underline{y} = \underline{g}(\underline{f}(\underline{x})) = \underline{g}(\mathbf{f}(\underline{x}))$ – i.e., when $\underline{f}(\cdot) = \mathbf{f}(\cdot)$ is a vectorized scalar function is

$$\frac{\partial \underline{y}}{\partial \underline{x}} = \frac{\partial \mathbf{f}}{\partial \underline{x}} \frac{\partial \underline{g}}{\partial \mathbf{f}} = \underline{\underline{\text{diag}}} \left(\dot{\mathbf{f}}(\underline{x}) \right) \frac{\partial \underline{g}}{\partial \mathbf{f}} \quad (11)$$

2 Notation for MLP Processing

The processing in the l^{th} layer of a multilayer perceptron can be expressed in vector form as

$$\underline{v}^{(l)} = \mathbf{h} \left(\underline{\underline{W}}^{(l)} \underline{v}^{(l-1)} + \underline{w}_0^{(l)} \right) = \mathbf{h} \left(\underline{a}^{(l)} \right) \quad l = 1, 2, \dots, L \quad (12)$$

where $\underline{v}^{(l)}$ is the *activation* for layer l . Denoting the number of neurons (nodes) at layer l as M_l , the weight matrix $\underline{\underline{W}}^{(l)}$ is $(M_l \times M_{l-1})$, and the bias vector $\underline{w}_0^{(l)}$ is $(M_l \times 1)$. The activation function $h(\cdot)$ is vectorized to $\mathbf{h}(\cdot)$ as described in Section 1.2. This is shown in Figure 2. The *linear activation* or *pre-activation* for layer l is implicitly defined in (12) as

$$\underline{a}^{(l)} = \underline{\underline{W}}^{(l)} \underline{v}^{(l-1)} + \underline{w}_0^{(l)} \quad (13)$$

The processing in (12) is initialized with $\underline{v}^{(0)} = \underline{x}$ where \underline{x} is the network input (non-augmented). The output of the network is $\underline{v}^{(L)}$ for an L -layer network. Specifically, for a regression problem the final layer will typically have a linear activation and $\hat{\underline{y}}(\underline{x}) = \underline{v}^{(L)} = \underline{a}^{(L)}$. For a classification problem, the last layer typically contains a SoftMax operation and $\hat{\underline{p}} = \underline{v}^{(L)} = \text{SoftMax}(\underline{a}^{(L)})$ is the network's predictions of the class membership of the network input \underline{x} .²

²Note in the case of a SoftMax output layer, the activation is not a vectorized scalar function. Specifically, each coordinate of the SoftMax(\underline{a}) function is a function of every component of \underline{a} . This is discussed in detail in Section 3.3.2.

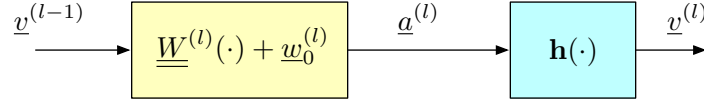


Figure 2: The processing for layer l where a vectorized scalar function has been assumed for the activation function.

3 Back-Propagation Algorithm

In Section 2 we have defined the notation for processing a single input vector \underline{x} to a single network output vector $\underline{v}^{(L)}$. In this section we will consider a single data-point gradient descent update for this model. In other words, we consider the input \underline{x}_n and compute $\frac{\partial J_n}{\partial \theta}$ for every trainable parameter θ in the network – *e.g.*, θ can be one of the elements of $\underline{W}^{(l)}$ or $\underline{w}_0^{(l)}$ for each layer $l \in \{1, 2, \dots, L\}$. Here, the total criterion function is $J = \sum_{n=1}^N J_n$ and $J_n = J_n(\underline{x}_n; \{\underline{W}^{(l)}, \underline{w}_0^{(l)}\}_{l=1}^L)$. For notational simplicity, let us consider in this section a single input $\underline{x} = \underline{x}_n$ and denote the associated loss contribution as $J = J_n$. We generalize this to batch-based backprop learning in the section that follows.

The back-propagation learning algorithm is a “backward” recursion on the gradient of the loss with respect to the layer pre-activations – *i.e.*, computing $\nabla_{\underline{a}^{(l-1)}} J$ from $\nabla_{\underline{a}^{(l)}} J$. Specifically, defining

$$\underline{\delta}^{(l)} \triangleq \nabla_{\underline{a}^{(l)}} J \quad (14)$$

we seek a recursion computing $\underline{\delta}^{(l-1)}$ from $\underline{\delta}^{(l)}$. With this recursion, we only need two other components for backprop learning: (i) an initialization for $\underline{\delta}^{(L)} = \nabla_{\underline{a}^{(L)}} J$ and (ii) a way to express the partial derivatives of all trainable parameters in terms of $\underline{\delta}^{(l)}$. Let us consider each of these three components of backprop in the following subsections.

3.1 Computing GD Updates from Deltas

Assume that we have available $\underline{\delta}^{(l)} = \nabla_{\underline{a}^{(l)}} J$ at each layer and we wish to compute the gradient descent updates for the bias vectors and weight matrices at each layer. To simplify the notation, let us drop the index l and focus on the mapping from activation \underline{v} to pre-activation \underline{a} as shown in Figure 3. Here the activation from the previous layer \underline{v} is mapped to the pre-activation for the current layer via

$$\underline{a} = \underline{W} \underline{v} + \underline{w}_0 \quad (15)$$

First, let us compute the gradient of the loss with respect to the bias vector. Using the right-to-left chain rule we have

$$\nabla_{\underline{w}_0} J = \frac{\partial J}{\partial \underline{w}_0} = \frac{\partial \underline{a}}{\partial \underline{w}_0} \frac{\partial J}{\partial \underline{a}} = \underline{I} \underline{\delta} = \underline{\delta} \quad (16)$$

where $\underline{\delta} = \nabla_{\underline{a}} J$.

To compute the GD updates for the matrix \underline{W} , we write out the expression in (15) for each component

$$a_m = \underline{w}_m^T \underline{v} + [w_0]_m \quad (17)$$

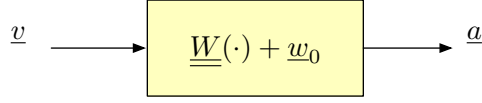


Figure 3: The mapping from activation to pre-activation for a given layer.

where $[w_0]_m$ is the m^{th} element of the vector \underline{w}_0 and \underline{w}_m^T is the m^{th} row of \underline{W}

$$\underline{W} = \begin{bmatrix} \underline{w}_1^T \\ \underline{w}_2^T \\ \underline{w}_3^T \\ \vdots \\ \underline{w}_M^T \end{bmatrix} \quad (18)$$

The gradient of J with respect to the \underline{w}_m is

$$\frac{\partial J}{\partial \underline{w}_m} = \frac{\partial a_m}{\partial \underline{w}_m} \frac{\partial J}{\partial a_m} = \underline{v} \delta_m \quad (19)$$

It follows that the GD update for the m^{th} row of \underline{W} is

$$\underline{w}_m^T(i+1) = \underline{w}_m^T(i) - \eta(i) \delta_m \underline{v}^T \quad (20)$$

Using this with (18) implies

$$\underline{W}(i+1) = \underline{W}(i) - \eta(i) \underline{\delta} \underline{v}^T \quad (21)$$

Using (16) and (21) and also identifying the layer indices from Figure 2 yields the method to use the $\underline{\delta}^{(l)}$ value to update the layer l weights and biases

$$\underline{W}^{(l)}(i+1) = \underline{W}^{(l)}(i) - \eta(i) \underline{\delta}^{(l)} \left(\underline{v}^{(l-1)} \right)^T \quad (22a)$$

$$\underline{w}_0^{(l)}(i+1) = \underline{w}_0^{(l)}(i) - \eta(i) \underline{\delta}^{(l)} \quad (22b)$$

3.2 The Delta Recursion

In order to get a recursion on $\underline{\delta}^{(l)} = \nabla_{\underline{a}^{(l)}} J$, consider the mapping from $\underline{a}^{(l-1)}$ to $\underline{a}^{(l)}$ as shown in Figure 4. The desired recursion follows from the right-to-left chain rule

$$\underline{\delta}^{(l-1)} = \frac{\partial J}{\partial \underline{a}^{(l-1)}} \quad (23)$$

$$= \frac{\partial \underline{v}^{(l-1)}}{\partial \underline{a}^{(l-1)}} \frac{\partial \underline{a}^{(l)}}{\partial \underline{v}^{(l-1)}} \frac{\partial J}{\partial \underline{a}^{(l)}} \quad (24)$$

$$= \underline{\text{diag}} \left(\dot{\mathbf{h}}(\underline{a}^{(l-1)}) \right) \left(\underline{W}^{(l)} \right)^T \underline{\delta}^{(l)} \quad (25)$$

$$= \dot{\mathbf{h}}(\underline{a}^{(l-1)}) \odot \left[\left(\underline{W}^{(l)} \right)^T \underline{\delta}^{(l)} \right] \quad (26)$$

$$= \dot{\mathbf{v}}^{(l-1)} \odot \left[\left(\underline{W}^{(l)} \right)^T \underline{\delta}^{(l)} \right] \quad (27)$$

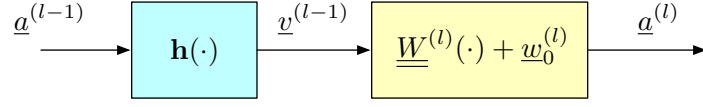


Figure 4: The mapping from pre-activation to pre-activation across a layer.

where we have used (6) and (11) and implicitly defined

$$\underline{v}^{(l)} \triangleq \dot{\mathbf{h}}(\underline{a}^{(l)}) \quad (28)$$

which are the *derivative activations* obtained by passing the pre-activations through the vectorized function $\dot{\mathbf{h}}(\cdot)$ – *i.e.*, the derivative of the activation function. Also implicitly defined in (27) is the Hadamard (or element-wise) product of two vectors

$$\underline{x} \odot \underline{y} \triangleq \begin{bmatrix} x_1 y_1 \\ x_2 y_2 \\ x_3 y_3 \\ \vdots \\ x_m y_m \end{bmatrix} \quad (29)$$

3.3 Initializing the Delta Recursion

In the previous two subsections, we have derived a backward recursion on $\underline{\delta}^{(l)}$ and shown how $\underline{\delta}^{(l)}$ can be mapped to the GD updates for the layer weight matrix and bias vector. All that remains is to determine how to initialize the delta recursion. Specifically, we need an expression for $\underline{\delta}^{(L)} = \nabla_{\underline{a}^{(L)}} J$. This follows again by the right-to-left chain rule to the processing shown in Figure 5. Specifically,

$$\underline{\delta}^{(L)} = \nabla_{\underline{a}^{(L)}} J \quad (30)$$

$$= \frac{\partial J}{\partial \underline{a}^{(L)}} \quad (31)$$

$$= \frac{\partial \underline{v}^{(L)}}{\partial \underline{a}^{(L)}} \frac{\partial J}{\partial \underline{v}^{(L)}} \quad (32)$$

$$= \underline{v}^{(L)} \odot \nabla_{\underline{v}^{(L)}} J \quad (33)$$

where it has been assumed that the final layer activation in a vectorized scalar function. Note that $\nabla_{\underline{v}^{(L)}} J$ is the standard gradient of the criterion function with respect to the model prediction – *i.e.*, this is determined by the functional form of the loss function.

3.3.1 Special Case: MSE Criterion

For the MSE criterion function, we have $J = \frac{1}{2} \|\underline{v}^{(L)} - \underline{y}\|^2$ where, for the typical case, \underline{y} is the target vector for a multi-dimensional target regression problem. It follows that

$$\nabla_{\underline{v}^{(L)}} J = (\underline{v}^{(L)} - \underline{y}) \quad (34)$$

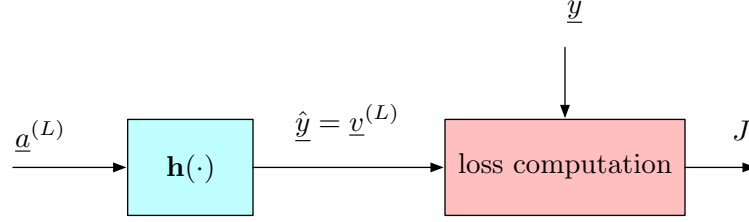


Figure 5: The computation of the loss function from the final layer output. A vectorized scalar activation function has been assumed for the final layer.

Thus, for MSE regression, we have

$$\underline{\delta}^{(L)} = \underline{\dot{v}}^{(L)} \odot (\underline{v}^{(L)} - \underline{y}) \quad (35)$$

and for MSE regression, typically the final layer activation is linear so that $\underline{\hat{y}} = \underline{a}^{(L)} = \underline{s}^{(L)}$ and $\dot{h}(v) = 1$. It follows that for MSE loss and a linear final layer activation

$$\underline{\delta}^{(L)} = (\underline{\hat{y}} - \underline{y}) \quad (36)$$

3.3.2 Special Case: SoftMax with Multiclass Cross-Entropy Loss

Another special case of interest is when the final layer has SoftMax activation and multiclass cross-entropy loss is used. Specifically, when

$$\underline{\hat{p}} = \underline{v}^{(L)} = \text{SoftMax}(\underline{a}^{(L)}) \quad (37)$$

and

$$J = \text{MCE}(\underline{p}, \underline{\hat{p}}) = - \sum_{m=1}^C p_m \ln(\hat{p}_m) \quad (38)$$

where p_m is the label probability for class m . In most cases $p_m = 1$ for the true class and $p_m = 0$ for all other classes. In the case of SoftMax output activation, the final activation is not a vectorized scalar function, as is illustrated in Figure 6. Again, applying the right-to-left chain rule, we obtain

$$\underline{\delta}^{(L)} = \nabla_{\underline{a}^{(L)}} J \quad (39)$$

$$= \frac{\partial J}{\partial \underline{a}^{(L)}} \quad (40)$$

$$= \frac{\partial \underline{v}^{(L)}}{\partial \underline{a}^{(L)}} \frac{\partial J}{\partial \underline{v}^{(L)}} \quad (41)$$

where now $\frac{\partial \underline{v}^{(L)}}{\partial \underline{a}^{(L)}}$ is a non-diagonal matrix. In the case of the SoftMax output activation, this is a $(C \times C)$ matrix, where C is the number of classes.

For the special case of SoftMax final activation and multiclass cross-entropy loss, it can be shown that

$$\underline{\delta}^{(L)} = \underline{v}^{(L)} - \underline{p} = \underline{\hat{p}} - \underline{p} \quad (42)$$

so that the delta recursion is initialized with the difference between the model prediction and the label/target vector.

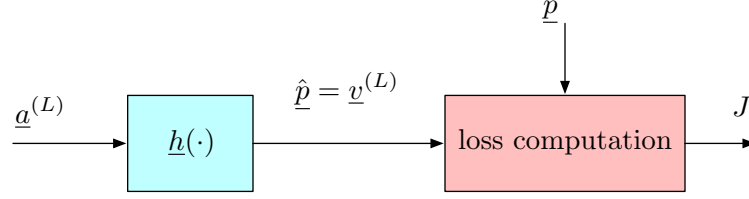


Figure 6: The computation of the loss function from the final layer output where the nonlinear activation of the final layer is assumed to be a general vector-valued function.

3.4 Single Data-point Back-Propagation Learning Summary

In summary, for a single input $\underline{v}^{(0)} = \underline{x}$ and output $\underline{v}^{(L)}$, the back-propagation learning algorithm is

$$\underline{v}^{(l)} = \mathbf{h} \left(\underline{\underline{W}}^{(l)} \underline{v}^{(l-1)} + \underline{w}_0^{(l)} \right) \quad \text{Layer Activation} \quad (43a)$$

$$\dot{\underline{v}}^{(l)} = \dot{\mathbf{h}} \left(\underline{\underline{W}}^{(l)} \underline{v}^{(l-1)} + \underline{w}_0^{(l)} \right) \quad \text{Layer Derivative Activation} \quad (43b)$$

$$\underline{\delta}^{(L)} = \frac{\partial \underline{v}^{(L)}}{\partial \underline{a}^{(L)}} \nabla_{\underline{v}^{(L)}} J \quad \text{Delta Recursion Initialization} \quad (43c)$$

$$\underline{\delta}^{(l-1)} = \dot{\underline{v}}^{(l-1)} \odot \left[\left(\underline{\underline{W}}^{(l)} \right)^T \underline{\delta}^{(l)} \right] \quad \text{Delta Recursion} \quad (43d)$$

$$\underline{\underline{W}}^{(l)}(i+1) = \underline{\underline{W}}^{(l)}(i) - \eta(i) \underline{\delta}^{(l)} \left(\underline{v}^{(l-1)} \right)^T \quad \text{Weight Matrix GD Update} \quad (43e)$$

$$\underline{w}_0^{(l)}(i+1) = \underline{w}_0^{(l)}(i) - \eta(i) \underline{\delta}^{(l)} \quad \text{Bias Vector GD Update} \quad (43f)$$

4 Extensions and Discussion

In the previous backprop development, we considered only a single data-point \underline{x} along with the associated target/label. Consider the case where mini-batches (or full-batch) are used. In this case, we compute (43a) - (43d) for each data point and then average the updates in (43e) and (43f) over the mini-batch. Specifically, let the data be $\{(\underline{x}_n, \underline{y}_n)\}_{n=1}^{N_{\text{train}}}$ and denote the result of using the the n^{th} data point by $(\cdot)_n$. Translating the above description to equations with this convention

yields³

$$\underline{v}_n^{(l)} = \mathbf{h} \left(\underline{\underline{W}}^{(l)} \underline{v}_n^{(l-1)} + \underline{w}_0^{(l)} \right) \quad \text{Layer Activations} \quad (44a)$$

$$\dot{\underline{v}}_n^{(l)} = \dot{\mathbf{h}} \left(\underline{\underline{W}}^{(l)} \underline{v}_n^{(l-1)} + \underline{w}_0^{(l)} \right) \quad \text{Layer Derivative Activations} \quad (44b)$$

$$\underline{\delta}_n^{(L)} = \frac{\partial v_n^{(L)}}{\partial \underline{a}_n^{(L)}} \nabla_{\underline{v}_n^{(L)}} J_n \quad \text{Delta Recursion Initializations} \quad (44c)$$

$$\underline{\delta}_n^{(l-1)} = \dot{\underline{v}}_n^{(l-1)} \odot \left[\left(\underline{\underline{W}}^{(l)} \right)^T \underline{\delta}_n^{(l)} \right] \quad \text{Delta Recursions} \quad (44d)$$

$$\underline{\underline{W}}^{(l)}(i+1) = \underline{\underline{W}}^{(l)}(i) - \eta(i) \left[\frac{1}{|\mathcal{B}_i|} \sum_{n \in \mathcal{B}_i} \underline{\delta}_n^{(l)} \left(\underline{v}_n^{(l-1)} \right)^T \right] \quad \text{Weight Matrix GD Update} \quad (44e)$$

$$\underline{w}_0^{(l)}(i+1) = \underline{w}_0^{(l)}(i) - \eta(i) \left[\frac{1}{|\mathcal{B}_i|} \sum_{n \in \mathcal{B}_i} \underline{\delta}_n^{(l)} \right] \quad \text{Bias Vector GD Update} \quad (44f)$$

where \mathcal{B}_i is the set of data indices for the i^{th} mini-batch and $|\mathcal{B}_i|$ is the number of data points in this mini-batch.

Another extension is the use of weight regularization. For example, if L2 regularization of the weights and biases is employed, with regularizer coefficient λ , then the weight updates in (43e) and bias updates in (43f) are modified to

$$\underline{\underline{W}}^{(l)}(i+1) = \underline{\underline{W}}^{(l)}(i) - \eta(i) \left(\underline{\delta}^{(l)} \left(\underline{v}^{(l-1)} \right)^T + 2\lambda \underline{\underline{W}}^{(l)}(i) \right) \quad (45a)$$

$$\underline{w}_0^{(l)}(i+1) = \underline{w}_0^{(l)}(i) - \eta(i) \left(\underline{\delta}^{(l)} + 2\lambda \underline{w}_0^{(l)}(i) \right) \quad (45b)$$

Finally, there are other types of layers used in modern ANNs. For example, convolutional layers, recurrent layers, and attention layers. the backprop equations for networks with these types of layers can be computed using the same principles employed herein. Specifically, the key step is to get a recursion on the pre-activations. In practice, most of the heavy lifting is done for us by auto-differentiation routines available in frameworks such as TensorFlow and PyTorch and one does not need to directly implement the backprop equations. However, for some applications (*e.g.*, imbedded programming, digital circuit acceleration) direct implementation of backprop using the equations derived herein is required.

³These updates can be vectorized over the variable n so that instead of looping over $|\mathcal{B}_i|$ matrix vector multiplications, one can do one matrix-matrix multiplication.