

# Named Tensor Notation

**David Chiang**

*University of Notre Dame*

**Alexander M. Rush**

*Cornell University*

**Boaz Barak**

*Harvard University*

**Reviewed on OpenReview:** <https://openreview.net/forum?id=hVT7SHlilx>

## Abstract

We propose a notation for tensors with named axes, which relieves the author, reader, and future implementers of machine learning models from the burden of keeping track of the order of axes and the purpose of each. The notation makes it easy to lift operations on low-order tensors to higher order ones, for example, from images to minibatches of images, or from an attention mechanism to multiple attention heads.

After a brief overview and formal definition of the notation, we illustrate it through several examples from modern machine learning, from building blocks like attention and convolution to full models like Transformers and LeNet. We then discuss differential calculus in our notation and compare with some alternative notations. Our proposals build on ideas from many previous papers and software libraries. We hope that our notation will encourage more authors to use named tensors, resulting in clearer papers and more precise implementations.

## 1 Introduction

Formal descriptions of neural networks primarily adopt the notation of vectors and matrices from applied linear algebra (Goodfellow et al., 2016). When used to describe vector spaces, this notation is both concise and unambiguous. However, when applied to neural networks, these properties are lost. Consider the equation for attention as notated in the Transformer paper (Vaswani et al., 2017):

$$\text{Attention}(Q, K, V) = \left( \text{softmax} \frac{QK^\top}{\sqrt{d_k}} \right) V.$$

The equation relates  $Q$ ,  $K$ , and  $V$  (for query, key, and value, respectively) as sequences of feature vectors, packed into possibly identically-sized matrices. While concise, this equation is ambiguous. Does the product  $QK^\top$  sum over the sequence, or over the features? We know that it sums over columns, but there is not enough information to know what the columns represent. Is the softmax taken over the query sequence or the key sequence? The usual notation does not offer an answer. Perniciously, the implementation of an incorrect interpretation might still run without errors. With the addition of more axes, like multiple attention heads or multiple sentences in a minibatch, the notation becomes even more cumbersome.

We propose an alternative mathematical notation for tensors with *named axes*.<sup>1</sup> The notation has a formal underpinning, but is hopefully intuitive enough that machine learning researchers can understand it without much effort. In named tensor notation, the above equation becomes

$$\text{Attention}: \mathbb{R}^{\text{key}} \times \mathbb{R}^{\text{seq} \times \text{key}} \times \mathbb{R}^{\text{seq} \times \text{val}} \rightarrow \mathbb{R}^{\text{val}}$$

<sup>1</sup>We follow NumPy in using the term *axis*. Other possible terms would be *index*, *dimension*, *way*, or *mode* (Tucker, 1964), but we felt that *axis* had the least potential for confusion.

$$\text{Attention}(Q, K, V) = \left( \underset{\text{seq}}{\text{softmax}} \frac{Q \underset{\text{key}}{\odot} K}{\sqrt{|\text{key}|}} \right) \underset{\text{seq}}{\odot} V.$$

The type signature introduces three named axes: the `key` axis is for features of queries and keys, the `val` axis is for features of values, and the `seq` axis is for tokens in a sequence. (Please see Section 2.2 for an explanation of our naming convention.) This notation makes the types of each input tensor explicit. Tensor  $Q$  is a query vector that is compared with key vectors, so it has a `key` axis. Tensor  $K$  is a sequence of key vectors, so it has `seq` and `key` axes. Tensor  $V$  is a sequence of value vectors, so it has `seq` and `val` axes. Unlike with matrix notation, the reader is not required to remember whether `seq` corresponds to rows or columns in either of these tensors.

The function itself uses the named axes to precisely apply operations. The expression  $Q \underset{\text{key}}{\odot} K$  is a dot product over the `key` axis shared between  $K$  and  $Q$ ; there is no ambiguity about rows or columns. Similarly, the `softmax` function is annotated with the axis along which it is applied, removing any ambiguity or reliance on convention.

Furthermore, named tensor notation naturally extends to *lifting* (also known as vectorizing and/or broadcasting) a function to tensors with more axes. For example, if instead of being a tensor with the single axis `key`,  $Q$  has three axes `key`, `seq` and `batch` (corresponding to tokens of a sequence and examples in a minibatch, respectively) then the Attention function works as written, acting on each example in a minibatch in parallel. Similarly, we can also add a `heads` axis to the inputs to get multiple attention heads. These additional axes are often elided in neural network papers, possibly avoiding notational complexity, but possibly also hiding critical model details.

**Our contributions.** This work proposes a *mathematical notation* for named tensors and a fully specified *semantic interpretation* for the notation. Through examples, we demonstrate that this notation enables specifying machine learning models and operations in a succinct yet precise manner. The need for named tensors has been recognized by several software packages, including `xarray` (Hoyer & Hamman, 2017), `Nexus` (Chen, 2017), `tsalib` (Sinha, 2018), `axisarrays` (Bauman, 2018), `NamedTensor` (Rush, 2019), `PyTorch` (Torch Contributors, 2019), `Dex` (Paszke et al., 2021), `JAX` (JAX authors, 2021), `einops` (Rogozhnikov, 2022), and `torchdim` (DeVito, 2023). While our notation is inspired by these efforts, our focus is on mathematical notation to be used in papers, whereas previous efforts have focused on code. Our hope is that our notation will be adopted by authors, leading to clearer, more replicable papers, and that this, in turn, will encourage more implementers to adopt named tensor libraries, leading to clearer, more correct code.

## 2 Named Tensors

In standard notation, a vector, matrix, or tensor is indexed by an integer or sequence of integers; if it has dimensions  $n_1, \dots, n_r$ , it can be thought of as a map that takes as input  $(i_1, \dots, i_r) \in [n_1] \times \dots \times [n_r]$  and outputs a real number (or an element of a different field). For example, if  $A \in \mathbb{R}^{3 \times 3}$ , then the order of the two axes matters:  $A_{1,3}$  and  $A_{3,1}$  are not the same element. It is up to the reader to remember what each axis of each tensor stands for. This problem is exacerbated in modern machine learning, where tensors have multiple axes with different meanings (batches, channels, etc.), and different operations act on different axes.

In contrast, we propose *named tensors*, in which each axis has a *name* that describes it and ensures there is no confusion between axes. We write  $\text{ax}[n]$  for an axis with name `ax` and size  $n$ , and we write  $\text{ax}(i)$  to index the  $i$ -th element along axis `ax`. So if a tensor has axes  $\text{ax}_1[n_1], \dots, \text{ax}_r[n_r]$  (with  $\text{ax}_1, \dots, \text{ax}_r$  being distinct names), it can be thought of as a map that takes as input a *record*  $\{\text{ax}_1(i_1), \dots, \text{ax}_r(i_r)\}$ , with  $i_1 \in [n_1], \dots, i_r \in [n_r]$ , and outputs a field element.

In summary the key difference is that, while a tensor in standard notation takes as input an ordered tuple of indices, a named tensor takes as input a record, which is an unordered set of named indices. We illustrate with some examples below, then give formal definitions.

## 2.1 By example

For example, if  $A$  represents a  $3 \times 3$  grayscale image, we can make it a named tensor like so (writing it two equivalent ways to show that the order of axes does not matter):

$$A \in \mathbb{R}^{\text{height}[3] \times \text{width}[3]} = \mathbb{R}^{\text{width}[3] \times \text{height}[3]}$$

$$A = \text{height} \begin{bmatrix} \text{width} & & \\ 3 & 1 & 4 \\ 1 & 5 & 9 \\ 2 & 6 & 5 \end{bmatrix} = \text{width} \begin{bmatrix} \text{height} & & \\ 3 & 1 & 2 \\ 1 & 5 & 6 \\ 4 & 9 & 5 \end{bmatrix}. \quad (1)$$

We access elements of  $A$  using named indices, whose order again does not matter:  $A_{\text{height}(1), \text{width}(3)} = A_{\text{width}(3), \text{height}(1)} = 4$ . We also allow partial indexing:

$$A_{\text{height}(1)} = \begin{bmatrix} \text{width} \\ 3 & 1 & 4 \end{bmatrix} \quad A_{\text{width}(3)} = \begin{bmatrix} \text{height} \\ 4 & 9 & 5 \end{bmatrix}.$$

It does not matter if we write  $A_{\text{height}(1)}$  or  $A_{\text{width}(3)}$  as row and column vectors. In many contexts, an axis name is used with only one size. If so, we can simply write **height** for the unique axis with name **height**, as in  $\mathbb{R}^{\text{height} \times \text{width}}$ . We can leave the size of an axis unspecified at first, and specify its size later (e.g., deferring it to an appendix on experimental details). For example, we can specify  $|\text{height}| = |\text{width}| = 28$  if we want to prescribe the precise size of an image, or just write  $|\text{height}| = |\text{width}|$  to specify that it's a square image.

## 2.2 What's in a name?

Although users of this notation are free to choose any names for axes, we offer the following recommendations. First, we recommend *words* instead of single letters, to communicate better the meaning of each axis.

More subtly, we recommend words that describe a *whole* rather than its parts. For example, to represent a minibatch of examples, we would name the axis **batch**; to represent a sequence of tokens, we would name the axis **seq**. One reason for this choice is that there are cases, like **height** and **width**, where there is a name for the whole, but no unambiguous name for the part. By contrast, in cases where there is a name for the part but not the whole, it's always possible to use the plural form of the name of the part. For example, if we wanted  $A$  to have red, green, and blue channels, we would name the axis **chans**.

Section 4 contains many more examples of axis names.

## 2.3 Formal definition

We now define formally the notation we use.

**Definition 1** (Names, indices, and axes). An *axis* is a pair, written  $\text{ax}[I]$ , where

- $\text{ax}$  is the *name* of the axis, which is simply a string of letters. We write both names and variables ranging over names using sans-serif font.
- $I$  is a set of *indices*. In this paper,  $I$  is always of the form  $\{1, \dots, n\}$  for some  $n$ , so we abbreviate  $\text{ax}\{\{1, \dots, n\}\}$  as  $\text{ax}[n]$ .

In many contexts, there is only one axis with name  $\text{ax}$ , and so we refer to the axis simply as  $\text{ax}$ . The context always makes it clear whether  $\text{ax}$  is a name or an axis. If  $\text{ax}$  is an axis, we write  $\text{ind}(\text{ax})$  for its index set, and we write  $|\text{ax}|$  as shorthand for  $|\text{ind}(\text{ax})|$ .

**Definition 2** (Named indices and records). If  $\text{ax}[I]$  is an axis and  $i \in I$ , then a *named index* is a pair, written  $\text{ax}(i)$ . A *record* is a set of named indices  $\{\text{ax}_1(i_1), \dots, \text{ax}_r(i_r)\}$ , where  $\text{ax}_1, \dots, \text{ax}_r$  are pairwise distinct names.

**Definition 3** (Shapes). A *shape* is a set of axes, written  $\mathbf{ax}_1[I_1] \times \cdots \times \mathbf{ax}_r[I_r]$ , where  $\mathbf{ax}_1, \dots, \mathbf{ax}_r$  are pairwise distinct names. We write  $\emptyset$  for the empty shape. A shape defines a set of records:

$$\text{rec}(\mathbf{ax}_1[I_1] \times \cdots \times \mathbf{ax}_r[I_r]) = \{\{\mathbf{ax}_1(i_1), \dots, \mathbf{ax}_r(i_r)\} \mid i_1 \in I_1, \dots, i_r \in I_r\}.$$

We say two shapes  $\mathcal{S}$  and  $\mathcal{T}$  are *compatible* if whenever  $\mathbf{ax}[I] \in \mathcal{S}$  and  $\mathbf{ax}[J] \in \mathcal{T}$ , then  $I = J$ . We say that  $\mathcal{S}$  and  $\mathcal{T}$  are *orthogonal* if there is no  $\mathbf{ax}$  such that  $\mathbf{ax}[I] \in \mathcal{S}$  and  $\mathbf{ax}[J] \in \mathcal{T}$  for any  $I, J$ . If  $t \in \text{rec } \mathcal{T}$  and  $\mathcal{S} \subseteq \mathcal{T}$ , then we write  $t|_{\mathcal{S}}$  for the unique record in  $\text{rec } \mathcal{S}$  such that  $t|_{\mathcal{S}} \subseteq t$ .

**Definition 4** (Named tensors). Let  $F$  be a field and let  $\mathcal{S}$  be a shape. Then a *named tensor over  $F$  with shape  $\mathcal{S}$*  is a mapping from  $\text{rec } \mathcal{S}$  to  $F$ . If  $X$  has shape  $\mathcal{S}$  then we write  $\text{shp } X = \mathcal{S}$ . We write the set of all named tensors with shape  $\mathcal{S}$  as  $F^{\mathcal{S}}$ .

We don't make any distinction between a scalar (an element of  $F$ ) and a named tensor with empty shape (an element of  $F^{\emptyset}$ ).

If  $X \in F^{\mathcal{S}}$ , then we access an element of  $X$  by applying it to a record  $s \in \text{rec } \mathcal{S}$ ; but we write this using the usual subscript notation:  $X_s$  rather than  $X(s)$ . To avoid clutter, in place of  $X_{\{\mathbf{ax}_1(i_1), \dots, \mathbf{ax}_r(i_r)\}}$ , we usually write  $X_{\mathbf{ax}_1(i_1), \dots, \mathbf{ax}_r(i_r)}$ . When a named tensor is an expression like  $(X + Y)$ , we index it by surrounding it with square brackets like this:  $[X + Y]_{\mathbf{ax}_1(i_1), \dots, \mathbf{ax}_r(i_r)}$ .

We also allow partial indexing. If  $X$  is a tensor with shape  $\mathcal{T}$  and  $s \in \text{rec } \mathcal{S}$  where  $\mathcal{S} \subseteq \mathcal{T}$ , then we define  $X_s$  to be the named tensor with shape  $\mathcal{T} \setminus \mathcal{S}$  such that, for any  $t \in \text{rec}(\mathcal{T} \setminus \mathcal{S})$ ,

$$[X_s]_t = X_{s \cup t}.$$

(For the edge case  $\mathcal{T} = \emptyset$ , our definitions for indexing and partial indexing coincide: one gives a scalar and the other gives a tensor with empty shape, but we don't distinguish between the two.)

### 3 Operations

A significant benefit of named tensor notation is that it allows one to unambiguously specify *operations* that map tensors to tensors, and defines precisely how operations can be *lifted* when an operation is applied to tensors with more axes than are present in its signature and how *broadcasting* happens when different arguments add different axes.

We start with the formal definition of named tensor operations and lifting, then show how this definition leads to many common operations.

#### 3.1 Formal definition

By (*named*) *tensor function* or (*named*) *tensor operation*, we mean not only functions from tensors to tensors, but also operators like negation ( $-$ ), addition ( $+$ ), and so on. We will extend the standard function/operator notation by allowing tensor operations to be *lifted* to higher-order tensors.

**Definition 5** (lifting, unary). Let  $f: F^{\mathcal{S}} \rightarrow G^{\mathcal{T}}$  be a function from tensors to tensors. For any shape  $\mathcal{S}'$  orthogonal to both  $\mathcal{S}$  and  $\mathcal{T}$ , we can define the *lift*  $f^{\mathcal{S}'}$  of  $f$  with the shape  $\mathcal{S}'$  to be the map

$$f^{\mathcal{S}'}: F^{\mathcal{S} \cup \mathcal{S}'} \rightarrow G^{\mathcal{T} \cup \mathcal{S}'}$$

$$\left[ f^{\mathcal{S}'}(X) \right]_{s'} = f(X_{s'}) \quad \text{for all } X \in F^{\mathcal{S} \cup \mathcal{S}'} \text{ and } s' \in \text{rec } \mathcal{S}'.$$

Usually, we simply write  $f$  instead of  $f^{\mathcal{S}'}$ . That is, for every tensor  $X$  with shape  $\mathcal{R} \supseteq \mathcal{S}$ , we let  $f(X) = f^{\mathcal{R} \setminus \mathcal{S}}(X)$ .

If  $f$  is a multary function, we can lift each of its arguments to larger shapes, and we don't have to add the same axes to all the arguments; an axis present in one argument but not another is *broadcast* from the former to the latter. We consider just the case of two arguments; three or more arguments are analogous.

**Definition 6** (lifting, binary). Let  $f: F^{\mathcal{S}} \times G^{\mathcal{T}} \rightarrow H^{\mathcal{U}}$  be a binary function from tensors to tensors. For any shapes  $\mathcal{S}'$  and  $\mathcal{T}'$  that are compatible with each other and orthogonal to  $\mathcal{S}$  and  $\mathcal{T}$ , respectively, and such that  $\mathcal{S}' \cup \mathcal{T}'$  is orthogonal to  $\mathcal{U}$ , we can lift  $f$  to:

$$f^{\mathcal{S}' \cup \mathcal{T}'}: F^{\mathcal{S} \cup \mathcal{S}'} \times G^{\mathcal{T} \cup \mathcal{T}'} \rightarrow H^{\mathcal{U} \cup \mathcal{S}' \cup \mathcal{T}'}$$

$$\left[ f^{\mathcal{S}' \cup \mathcal{T}'}(X, Y) \right]_{s'} = f(X_{s'|_{\mathcal{S}'}}, Y_{s'|_{\mathcal{T}'}}) \quad \text{for all } X \in F^{\mathcal{S} \cup \mathcal{S}'}, Y \in G^{\mathcal{T} \cup \mathcal{T}'}, s' \in \text{rec}(\mathcal{S}' \cup \mathcal{T}').$$

Again, we usually write  $f$  instead of  $f^{\mathcal{S}' \cup \mathcal{T}'}$ .

In the following sections, we present some consequences of the above lifting rules. In particular, we show how they allow one to lift some common operations from operating on scalars, vectors, or matrices to operating on tensors with more axes, and how they correspond to vectorizing and broadcasting (as implemented by NumPy and related packages).

### 3.2 Elementwise operations and broadcasting

Any function from a scalar to a scalar corresponds to a tensor function with signature  $F^{\emptyset} \rightarrow F^{\emptyset}$ . Hence lifting it to any tensor shape, by Definition 5, corresponds to elementwise application. For example, given the logistic sigmoid function,

$$\sigma: \mathbb{R} \rightarrow \mathbb{R}$$

$$\sigma(x) = \frac{1}{1 + \exp(-x)}$$

we can lift it to tensors. If  $A \in \mathbb{R}^{\text{height}[3] \times \text{width}[3]}$  is the example tensor (1), then

$$\sigma(A) = \text{height} \begin{bmatrix} \frac{1}{1+\exp(-3)} & \frac{1}{1+\exp(-1)} & \frac{1}{1+\exp(-4)} \\ \frac{1}{1+\exp(-1)} & \frac{1}{1+\exp(-5)} & \frac{1}{1+\exp(-9)} \\ \frac{1}{1+\exp(-2)} & \frac{1}{1+\exp(-6)} & \frac{1}{1+\exp(-5)} \end{bmatrix}.$$

Similarly for rectified linear units ( $\text{relu}(x) = \max(0, x)$ ), negation, and so on.

Any function with signature  $\mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ , including binary operators like addition (+), can be applied to two named tensors with the same shape. But if we apply a binary function or operator to tensors with different shapes, then, by Definition 6, broadcasting applies. For example, let

$$x \in \mathbb{R}^{\text{height}[3]} \qquad y \in \mathbb{R}^{\text{width}[3]}$$

$$x = \text{height} \begin{bmatrix} 2 \\ 7 \\ 1 \end{bmatrix} \qquad y = \begin{bmatrix} \text{width} \\ 1 & 4 & 1 \end{bmatrix}.$$

(We write  $x$  as a column just to make the broadcasting easier to visualize.) Then, to evaluate  $A + x$ , we effectively replace  $x$  with a new tensor with a copy of  $x$  for every index of axis `width`. Likewise for  $A + y$ :

$$A + x = \text{height} \begin{bmatrix} \text{width} \\ 3+2 & 1+2 & 4+2 \\ 1+7 & 5+7 & 9+7 \\ 2+1 & 6+1 & 5+1 \end{bmatrix} \qquad A + y = \text{height} \begin{bmatrix} \text{width} \\ 3+1 & 1+4 & 4+1 \\ 1+1 & 5+4 & 9+1 \\ 2+1 & 6+4 & 5+1 \end{bmatrix}.$$

Similarly for other operations. We write elementwise multiplication (Hadamard product) as  $\odot$ .

### 3.3 Reductions

The same rules apply to functions from vectors to scalars, called *reductions*. We specify which axis a reduction applies to using a subscript (equivalent to the `axis` argument in NumPy and `dim` in PyTorch), so that even after lifting, we know which axis to reduce. For example, we can define summation:

$$\sum_{\text{ax}} -: \mathbb{R}^{\text{ax}[I]} \rightarrow \mathbb{R}$$

$$\sum_{\text{ax}} X = \sum_{i \in I} X_{\text{ax}(i)}$$

and use it on  $A$  from Eq. (1):

$$\sum_{\text{height}} A = \sum_i A_{\text{height}(i)} = \begin{bmatrix} 3 + 1 + 2 & 1 + 5 + 6 & 4 + 9 + 5 \end{bmatrix}$$

$$\sum_{\text{width}} A = \sum_j A_{\text{width}(j)} = \begin{bmatrix} 3 + 1 + 4 & 1 + 5 + 9 & 2 + 6 + 5 \end{bmatrix}.$$

We can also write multiple names to sum over multiple axes:

$$\sum_{\substack{\text{height} \\ \text{width}}} A = \sum_i \sum_j A_{\text{height}(i), \text{width}(j)} = 3 + 1 + 4 + 1 + 5 + 9 + 2 + 6 + 5.$$

But a summation with an index variable (like  $i$  or  $j$  above) is a standard summation over values of that variable, and a summation with no subscript is a standard summation over a set.

Other examples of reductions include:

$$\begin{aligned} \text{norm}_{\text{ax}} X &= \sqrt{\sum_{\text{ax}} X^2} & \text{norm}_p X &= \left( \sum_{\text{ax}} X^p \right)^{\frac{1}{p}} \\ \min_{\text{ax}} X &= \min\{X_{\text{ax}(i)} \mid i \in I\} & \max_{\text{ax}} X &= \max\{X_{\text{ax}(i)} \mid i \in I\} \\ \text{mean}_{\text{ax}} X &= \frac{1}{|\text{ax}|} \sum_{\text{ax}} X & \text{var}_{\text{ax}} X &= \frac{1}{|\text{ax}|} \sum_{\text{ax}} (X - \text{mean}_{\text{ax}} X)^2 \end{aligned}$$

### 3.4 Contraction

The vector dot-product is a function from *two* vectors to a scalar. We write it as follows:

$$- \odot -: \mathbb{R}^{\text{ax}[I]} \times \mathbb{R}^{\text{ax}[I]} \rightarrow \mathbb{R}$$

$$X \odot_{\text{ax}} Y = \sum_{i \in I} X_{\text{ax}(i)} Y_{\text{ax}(i)}$$

When lifted to higher-order tensors, the dot-product generalizes to the ubiquitous *contraction* operator, which can also be thought of as elementwise multiplication followed by summation over one axis, that is,

$$X \odot_{\text{ax}} Y = \sum_{\text{ax}} X \odot Y. \quad (2)$$

For example,

$$A \odot_{\text{height}} x = \sum_{\text{height}} A \odot x = \begin{bmatrix} 3 \cdot 2 + 1 \cdot 7 + 2 \cdot 1 & 1 \cdot 2 + 5 \cdot 7 + 6 \cdot 1 & 4 \cdot 2 + 9 \cdot 7 + 5 \cdot 1 \end{bmatrix}$$

$$A \underset{\text{width}}{\odot} y = \sum_{\text{width}} A \odot y = \text{height} \begin{bmatrix} 3 \cdot 1 + 1 \cdot 4 + 4 \cdot 1 \\ 1 \cdot 1 + 5 \cdot 4 + 9 \cdot 1 \\ 2 \cdot 1 + 6 \cdot 4 + 5 \cdot 1 \end{bmatrix}.$$

Again, we can write multiple names to contract multiple axes at once:

$$A \underset{\text{height}}{\underset{\text{width}}{\odot}} A = \sum_{\text{height}} \sum_{\text{width}} A \odot A = 3 \cdot 3 + 1 \cdot 1 + 4 \cdot 4 + 1 \cdot 1 + 5 \cdot 5 + 9 \cdot 9 + 2 \cdot 2 + 6 \cdot 6 + 5 \cdot 5$$

$A \odot$  with no axis name under it contracts zero axes and is equivalent to elementwise multiplication, which is the reason we use the same symbol  $\odot$  for elementwise multiplication and contraction. The contraction operator can be used for many multiplication-like operations:

$$\begin{aligned} x \underset{\text{height}}{\odot} x &= \sum_{\text{height}} x \odot x && \text{inner product} \\ x \odot y &= \text{height} \begin{bmatrix} \text{width} \\ 2 \cdot 1 & 2 \cdot 4 & 2 \cdot 1 \\ 7 \cdot 1 & 7 \cdot 4 & 7 \cdot 1 \\ 1 \cdot 1 & 1 \cdot 4 & 1 \cdot 1 \end{bmatrix} && \text{outer product} \\ A \underset{\text{width}}{\odot} y &= \sum_{\text{width}} A \odot y && \text{matrix-vector product} \\ x \underset{\text{height}}{\odot} A &= \sum_{\text{height}} x \odot A && \text{vector-matrix product} \\ A \underset{\text{width}}{\odot} B &= \sum_{\text{width}} A \odot B && \text{matrix-matrix product } (B \in \mathbb{R}^{\text{width} \times \text{width}'}) \end{aligned}$$

A contraction of three more tensors can be written as a sum. For example, the three-way inner product of vectors  $x, y, z \in \mathbb{R}^{\text{width}}$  can be written as  $\sum_{\text{width}} x \odot y \odot z$ .

Like the dot-product from which it is lifted, but unlike matrix multiplication, the contraction operator is commutative, but not associative. However, contraction does obey the following associative-like law.

$$X \underset{\mathcal{S} \cup \mathcal{T}}{\odot} (Y \underset{\mathcal{U}}{\odot} Z) = (X \underset{\mathcal{S}}{\odot} Y) \underset{\mathcal{T} \cup \mathcal{U}}{\odot} Z \quad \text{if } \mathcal{S} \cap \text{shp } Z = \mathcal{U} \cap \text{shp } X = \emptyset. \quad (3)$$

The special case

$$X \underset{\mathcal{S}}{\odot} (Y \underset{\mathcal{U}}{\odot} Z) = (X \underset{\mathcal{S}}{\odot} Y) \underset{\mathcal{U}}{\odot} Z \quad \text{if } \mathcal{S} \cap \text{shp } Z = \mathcal{U} \cap \text{shp } X = \emptyset \quad (4)$$

will be useful in Section 5 for moving  $Z$  from inside one or more sets of parentheses to the outside.

### 3.5 Vectors to vectors and beyond

Functions from vectors to vectors ( $\mathbb{R}^{\text{ax}[I]} \rightarrow \mathbb{R}^{\text{ax}[I]}$ ) lift to functions on tensors that operate along one axis, but leave the tensor shape unchanged. Such functions are particularly problematic in standard notation, which does not provide any way (to our knowledge) of specifying which axis the operation should be performed over. Such functions include:

$$\text{softmax}_{\text{ax}} X = \frac{\exp X}{\sum_{\text{ax}} \exp X} \quad (5a)$$

$$\text{argmax}_{\text{ax}} X = \lim_{\alpha \rightarrow \infty} \text{softmax}_{\text{ax}} \alpha X \quad (5b)$$

$$\operatorname{argmin}_{\mathbf{ax}} X = \lim_{\alpha \rightarrow -\infty} \operatorname{softmax}_{\mathbf{ax}} \alpha X \tag{5c}$$

For example, we can clearly distinguish between two ways of performing a softmax on  $A$ :

$$\begin{aligned} \operatorname{softmax}_{\text{height}} A &= \operatorname{height} \left[ \begin{array}{c|c|c} \text{width} & & \\ \hline \frac{\text{exp 3}}{\text{exp 3+exp 1+exp 2}} & \frac{\text{exp 1}}{\text{exp 1+exp 5+exp 6}} & \frac{\text{exp 4}}{\text{exp 4+exp 9+exp 5}} \\ \frac{\text{exp 1}}{\text{exp 3+exp 1+exp 2}} & \frac{\text{exp 5}}{\text{exp 1+exp 5+exp 6}} & \frac{\text{exp 9}}{\text{exp 4+exp 9+exp 5}} \\ \frac{\text{exp 2}}{\text{exp 3+exp 1+exp 2}} & \frac{\text{exp 6}}{\text{exp 1+exp 5+exp 6}} & \frac{\text{exp 5}}{\text{exp 4+exp 9+exp 5}} \end{array} \right] \\ \operatorname{softmax}_{\text{width}} A &= \operatorname{width} \left[ \begin{array}{c|c|c} \text{width} & & \\ \hline \frac{\text{exp 3}}{\text{exp 3+exp 1+exp 4}} & \frac{\text{exp 1}}{\text{exp 3+exp 1+exp 4}} & \frac{\text{exp 4}}{\text{exp 3+exp 1+exp 4}} \\ \frac{\text{exp 1}}{\text{exp 1+exp 5+exp 9}} & \frac{\text{exp 5}}{\text{exp 1+exp 5+exp 9}} & \frac{\text{exp 9}}{\text{exp 1+exp 5+exp 9}} \\ \frac{\text{exp 2}}{\text{exp 2+exp 6+exp 5}} & \frac{\text{exp 6}}{\text{exp 2+exp 6+exp 5}} & \frac{\text{exp 5}}{\text{exp 2+exp 6+exp 5}} \end{array} \right] \end{aligned}$$

### 3.6 Renaming and reshaping

It's often useful to rename an axis (analogous to a transpose operation in standard notation). We can think of this as the lift of a function from vectors to vectors, but with different input and output axes:

$$\begin{aligned} [-]_{\mathbf{ax} \rightarrow \mathbf{ax}'} : \mathbb{R}^{\mathbf{ax}[I]} &\rightarrow \mathbb{R}^{\mathbf{ax}'[I]} \\ [X_{\mathbf{ax} \rightarrow \mathbf{ax}'}]_{\mathbf{ax}'(i)} &= X_{\mathbf{ax}(i)} \end{aligned}$$

For example,

$$A_{\text{height} \rightarrow \text{height}'} = \operatorname{height}' \begin{array}{c|c|c} \text{width} & & \\ \hline 3 & 1 & 4 \\ 1 & 5 & 9 \\ 2 & 6 & 5 \end{array}.$$

We can also define notation for reshaping two or more axes into one axis:

$$A_{(\text{height}, \text{width}) \rightarrow \text{layer}} = \begin{array}{c|c|c|c|c|c|c|c|c} \text{layer} & & & & & & & & \\ \hline 3 & 1 & 4 & 1 & 5 & 9 & 2 & 6 & 5 \end{array}$$

The order of elements in the new axis is undefined. Authors who need a particular ordering may write a more specific definition.

### 3.7 Indexing<sup>2</sup>

NumPy and its derivatives provide various ways to recombine elements of a tensor to form a new tensor: integer array indexing, and functions like `numpy.take`, `numpy.take_along_axis`, `torch.index_select`, and `torch.gather`. Using named tensors, we can write nearly all of these operations as lifts of a single function:

$$\begin{aligned} \operatorname{index}_{\mathbf{ax}} : \mathbb{R}^{\mathbf{ax}[n]} \times [n] &\rightarrow \mathbb{R} \\ \operatorname{index}_{\mathbf{ax}}(X, i) &= X_{\mathbf{ax}(i)}. \end{aligned}$$

For example, suppose we have

$$\begin{aligned} E &\in \mathbb{R}^{\text{vocab}[n] \times \text{emb}} \\ i &\in [n] \\ I &\in [n]^{\text{seq}} \\ P &\in \mathbb{R}^{\text{seq} \times \text{vocab}[n]} \end{aligned}$$

<sup>2</sup>We are grateful to Tongfei Chen and Chu-Cheng Lin for contributing the original idea behind this section, as well as the example.



Tensor  $E$  contains word embeddings for all the words in the vocabulary. Integer  $i$  is the numeric identifier of a word, while tensor  $I$  is a sequence of numeric identifiers of words. Tensor  $P$  contains a sequence of probability distributions over the vocabulary (e.g., the predictions of a language model). Then:

- $\text{index}(E, i)$  broadcasts the `emb` axis from  $E$  to  $i$ , giving the word embedding of word  $i$ . This is the same as partial indexing ( $E_{\text{vocab}(i)}$ ).
- $\text{index}(E, I)$  also broadcasts the `seq` axis from  $I$  to  $E$ , giving a sequence of word embeddings. This is the same as integer array indexing ( $E[I]$ ), `numpy.take(E, I, 0)`, or `torch.index_select(E, 0, I)`.
- $\text{index}(P, I)$  aligns  $P$ 's and  $I$ 's `seq` axes, giving a sequence of probabilities. This is the same as `numpy.take_along_axis(P, I, 0)` or `torch.gather(P, 0, I)`.
- If  $P$  and  $I$  additionally had a `batch` axis (before the other axes), then  $\text{index}(P, I)$  would be the same as `tensorflow.gather(P, I, axis=1, batch_dims=1)`.

In NumPy, indexing using two or more integer arrays requires a special definition with some surprising special cases. With named tensors, we simply apply the indexing function twice. For example, if we wanted to get probabilities of words  $J$  at a subset  $I$  of positions, we could let:

$$\begin{aligned} |\text{seq}| &= m \\ I &\in [m]^{\text{subseq}} && \text{positions} \\ J &\in [n]^{\text{subseq}} && \text{numeric identifiers} \\ S &= \underset{\text{vocab}}{\text{index}}(\underset{\text{seq}}{\text{index}}(P, I), J) \in \mathbb{R}^{\text{subseq}} \end{aligned}$$

so that

$$S_{\text{subseq}(k)} = P_{\text{seq}(I_{\text{subseq}(k)}), \text{vocab}(I_{\text{subseq}(k)})}$$

## 4 Worked Examples: Neural Networks

In this section we give a series of worked examples illustrating how standard neural network model components can be written using named tensors. Appendix A builds some of these components into complete specifications of the Transformer and LeNet.

### 4.1 Feedforward neural networks

A multi-layer, feedforward neural network with different-sized layers can be written as:

$$\begin{aligned} X^0 &\in \mathbb{R}^{\text{input}} \\ X^1 &= \sigma(W^1 \underset{\text{input}}{\odot} X^0 + b^1) && W^1 \in \mathbb{R}^{\text{hidden}_1 \times \text{input}} && b^1 \in \mathbb{R}^{\text{hidden}_1} \\ X^2 &= \sigma(W^2 \underset{\text{hidden}_1}{\odot} X^1 + b^2) && W^2 \in \mathbb{R}^{\text{hidden}_2 \times \text{hidden}_1} && b^2 \in \mathbb{R}^{\text{hidden}_2} \\ X^3 &= \sigma(W^3 \underset{\text{hidden}_2}{\odot} X^2 + b^3) && W^3 \in \mathbb{R}^{\text{out} \times \text{hidden}_2} && b^3 \in \mathbb{R}^{\text{out}} \end{aligned}$$

The layer sizes can be specified by writing  $|\text{hidden}_1| = n_1$ , etc. As noted above,  $\sigma$  is applied elementwise and does not require additional annotation.

Alternatively, the layer equation can be abstracted by defining:

$$\text{FullConn}^l(x) = \sigma \left( W^l \underset{\text{layer}}{\odot} x + b^l \right)_{\text{layer}' \rightarrow \text{layer}}$$

where

$$\begin{aligned} W^l &\in \mathbb{R}^{\text{layer}'[n_l] \times \text{layer}[n_{l-1}]} \\ b^l &\in \mathbb{R}^{\text{layer}'[n_l]}. \end{aligned}$$

The function  $\text{FullConn}^l$  encapsulates both the equation for layer  $l$  as well as its parameters  $W^l, b^l$  (analogous to what TensorFlow and PyTorch call *modules*). Since we chose to use the same axis name `layer` for all the layers (with different sizes  $n_l$ ),  $\text{FullConn}^l$  temporarily computes its output over axis `layer'`, then renames it back to `layer`. The network can be defined like this:

$$\begin{aligned} X^0 &\in \mathbb{R}^{\text{layer}[n_0]} \\ X^1 &= \text{FullConn}^1(X^0) \\ X^2 &= \text{FullConn}^2(X^1) \\ X^3 &= \text{FullConn}^3(X^2). \end{aligned}$$

## 4.2 Recurrent neural networks

As a second example, we consider a simple (Elman) RNN. This model is similar to the feedforward network, except that the number of timesteps is variable and parameters are shared over time. At each time step, it produces a tensor with a new axis `hidden'` which is then renamed `hidden` for the next step in the recursion.

$$\begin{aligned} x^t &\in \mathbb{R}^{\text{input}} & t &= 1, \dots, n \\ W^h &\in \mathbb{R}^{\text{hidden} \times \text{hidden}'} & |\text{hidden}| &= |\text{hidden}'| \\ W^i &\in \mathbb{R}^{\text{input} \times \text{hidden}'} \\ b &\in \mathbb{R}^{\text{hidden}'} \\ h^0 &\in \mathbb{R}^{\text{hidden}} \\ h^t &= \sigma \left( W^h \underset{\text{hidden}}{\odot} h^{t-1} + W^i \underset{\text{input}}{\odot} x^t + b \right)_{\text{hidden}' \rightarrow \text{hidden}} & t &= 1, \dots, n \end{aligned}$$

## 4.3 Attention

In the introduction (§1), we described difficulties in interpreting the equation for attention as used with Transformers (Vaswani et al., 2017). In our notation, it looks like this:

$$\text{Attention}: \mathbb{R}^{\text{key}} \times \mathbb{R}^{\text{seq} \times \text{key}} \times \mathbb{R}^{\text{seq} \times \text{val}} \rightarrow \mathbb{R}^{\text{val}} \quad (6)$$

$$\text{Attention}(Q, K, V) = \left( \underset{\text{seq}}{\text{softmax}} \frac{Q \odot K}{\sqrt{|\text{key}|}} \right) \underset{\text{seq}}{\odot} V. \quad (7)$$

This definition takes a single query  $Q$  vector and returns a single result vector (and actually could be further reduced to a scalar values as `val` is not strictly necessary). To apply to a sequence, we can give  $Q$  a `seq'` axis, and the function will compute an output sequence. Providing  $Q, K,$  and  $V$  with a `heads` axis lifts the function to compute multiple attention heads.

For Transformers we often need to apply a mask to ensure attention is only applied to valid keys (e.g. for causal language models). We can modify the equation to include this mask:

$$\text{Attention}: \mathbb{R}^{\text{key}} \times \mathbb{R}^{\text{seq} \times \text{key}} \times \mathbb{R}^{\text{seq} \times \text{val}} \times \mathbb{R}^{\text{seq}} \rightarrow \mathbb{R}^{\text{val}}$$

$$\text{Attention}(Q, K, V, M) = \left( \underset{\text{seq}}{\text{softmax}} \left( \frac{Q \odot_{\text{key}} K}{\sqrt{|\text{key}|}} + M \right) \right) \odot_{\text{seq}} V.$$

Appendix A.1 includes a full specification of the complete Transformer model using the named tensor notation.

#### 4.4 Convolution

Standard neural network convolutions can be specified by “unrolling” a vector and then applying a standard dot product. We define an axis-parameterized unrolling function that converts a one-axis tensor to a sequence of kernel sized vectors:

$$\begin{aligned} \underset{\text{seq}}{\text{unroll}}: \mathbb{R}^{\text{seq}[n]} &\rightarrow \mathbb{R}^{\text{seq}[n-|\text{kernel}|+1], \text{kernel}} \\ \underset{\text{seq}}{\text{unroll}} X &= Y, \text{ where} \\ Y_{\text{seq}(i), \text{kernel}(j)} &= X_{\text{seq}(i+j-1)}. \end{aligned}$$

A 1d convolution with input channels  $\text{chans}$  and output channels  $\text{chans}'$  consists of unrolling along the  $\text{seq}$  axis and then taking a dot product:

$$\begin{aligned} \text{Conv1d}: \mathbb{R}^{\text{chans} \times \text{seq}[n]} &\rightarrow \mathbb{R}^{\text{chans}' \times \text{seq}[n']} \\ \text{Conv1d}(X; W, b) &= W \odot_{\substack{\text{chans} \\ \text{kernel}}} \underset{\text{seq}}{\text{unroll}} X + b \end{aligned}$$

where

$$\begin{aligned} W &\in \mathbb{R}^{\text{chans}' \times \text{chans} \times \text{kernel}} \\ b &\in \mathbb{R}^{\text{chans}'} \end{aligned}$$

Unrolling easily generalizes to higher-dimensional convolutions:

$$\begin{aligned} \text{Conv2d}: \mathbb{R}^{\text{chans} \times \text{height}[h] \times \text{width}[w]} &\rightarrow \mathbb{R}^{\text{chans}' \times \text{height}[h'] \times \text{width}[w']} \\ \text{Conv2d}(X; W, b) &= W \odot_{\substack{\text{chans} & \text{height} & \text{width} \\ \text{kh}, \text{kw} & \text{kh} & \text{kw}}} \underset{\text{seq}}{\text{unroll}} \underset{\text{seq}}{\text{unroll}} X + b \end{aligned}$$

where

$$\begin{aligned} W &\in \mathbb{R}^{\text{chans}' \times \text{chans} \times \text{kh} \times \text{kw}} \\ b &\in \mathbb{R}^{\text{chans}'} \end{aligned}$$

#### 4.5 Pooling

Pooling is similar to convolutions. We first define a function to partition a tensor into windows.

$$\begin{aligned} \underset{\text{seq}, \text{kernel}}{\text{pool}}: \mathbb{R}^{\text{seq}[n]} &\rightarrow \mathbb{R}^{\text{seq}[n/|\text{kernel}|], \text{kernel}} \\ \underset{\text{seq}, \text{kernel}}{\text{pool}} X &= Y, \text{ where} \\ Y_{\text{seq}(i), \text{kernel}(j)} &= X_{\text{seq}((i-1) \cdot |\text{kernel}| + j)}. \end{aligned}$$

Then we can define aggregations over kernel. We define max-pooling as:

$$\begin{aligned} \text{MaxPool1d}_k &: \mathbb{R}^{\text{seq}[n]} \rightarrow \mathbb{R}^{\text{seq}[n/k]} \\ \text{MaxPool1d}_k(X) &= \max_{\text{kernel seq, kernel}} \text{pool } X \\ &|\text{kernel}| = k \\ \text{MaxPool2d}_{kh, kw} &: \mathbb{R}^{\text{height}[h] \times \text{width}[w]} \rightarrow \mathbb{R}^{\text{height}[h/kh] \times \text{width}[w/kw]} \\ \text{MaxPool2d}_{kh, kw}(X) &= \max_{kh, kw \text{ height, kh width, kw}} \text{pool } \text{pool } X \\ &|\text{kh}| = kh \\ &|\text{kw}| = kw. \end{aligned}$$

#### 4.6 Normalization layers

Normalization layers are used in all large-scale deep learning models, with different architectures requiring different types of normalization. However, despite their importance, the differences between them are often not clearly communicated. For example, the PyTorch documentation (PyTorch Contributors, 2022) describes all of them using the same equation (where  $\epsilon > 0$  is a small constant for numerical stability):

$$Y = \frac{X - \text{mean}(X)}{\sqrt{\text{var}(X) + \epsilon}} \odot \gamma + \beta$$

Wu & He (2018) give essentially the same equation and explain the differences using a combination of equations, words, and pictures. But they do not capture differences in  $\gamma$  and  $\beta$  among different normalization layers.

Critically, the layers do differ by which axes are *standardized* as well as their parameters. We define a single named standardization function as:

$$\begin{aligned} \text{standardize}_{\text{ax}} &: \mathbb{R}^{\text{ax}} \rightarrow \mathbb{R}^{\text{ax}} \\ \text{standardize}_{\text{ax}}(X) &= \frac{X - \text{mean}_{\text{ax}}(X)}{\sqrt{\text{var}_{\text{ax}}(X) + \epsilon}} \end{aligned}$$

Then, we can define the three kinds of normalization layers, all with type  $\mathbb{R}^{\text{batch} \times \text{chans} \times \text{layer}} \rightarrow \mathbb{R}^{\text{batch} \times \text{chans} \times \text{layer}}$ . While superficially similar, these functions differ in their standardized axes and their parameter shape.

$$\begin{aligned} \text{BatchNorm}(X; \gamma, \beta) &= \text{standardize}_{\text{batch, layer}}(X) \odot \gamma + \beta & \gamma, \beta \in \mathbb{R}^{\text{chans}} \\ \text{InstanceNorm}(X; \gamma, \beta) &= \text{standardize}_{\text{layer}}(X) \odot \gamma + \beta & \gamma, \beta \in \mathbb{R}^{\text{chans}} \\ \text{LayerNorm}(X; \gamma, \beta) &= \text{standardize}_{\text{layer, chans}}(X) \odot \gamma + \beta & \gamma, \beta \in \mathbb{R}^{\text{chans, layer}} \end{aligned}$$

## 5 Differential Calculus

In many machine learning applications, we need to compute derivatives of functions from tensors to tensors. In standard vector/matrix notation, this can become complicated. For example, if  $f$  maps from vectors to vectors, then the partial derivatives of  $f$  form a matrix (the Jacobian). It has an “input” axis for the directions in which  $X$  could change, and an “output” axis for the directions in which  $f(X)$  could change. But there are conflicting conventions about whether the first axis is the input axis (“denominator layout”) or the output axis (“numerator layout”). The derivative of a function from vectors to matrices or matrices to vectors cannot be represented as a matrix at all, so one must resort to flattening the matrices into vectors.

With non-named tensor index notation, taking derivatives is not difficult (Laue et al., 2018), but again a convention must be adopted that the input axes come after the output axes, separated by a comma.

With named tensors, axes are not ordered, so we don't need to remember whether the input or output axes come first. But we do need to ensure that the input and output axes have different names.

## 5.1 Definition

**Definition 7.** Let  $\mathcal{S} = \mathbf{ax}_1 \times \cdots \times \mathbf{ax}_r$  be a shape. Then we write  $\mathcal{S}^* = \mathbf{ax}_1^* \times \cdots \times \mathbf{ax}_r^*$ , and if  $s = \{\mathbf{ax}_1(i_1), \dots, \mathbf{ax}_r(i_r)\} \in \text{rec } \mathcal{S}$ , then we write  $s^* = \{\mathbf{ax}_1^*(i_1), \dots, \mathbf{ax}_r^*(i_r)\}$ . Furthermore, if  $X \in \mathbb{R}^{\mathcal{S}}$  then we write  $X^* = X_{\mathcal{S} \rightarrow \mathcal{S}^*}$ .

**Definition 8.** Let  $f: \mathbb{R}^{\mathcal{S}} \rightarrow \mathbb{R}^{\mathcal{T}}$ . The *derivative* of  $f(X)$  with respect to  $X^*$  is the tensor such that

$$\frac{\partial f(X)}{\partial X^*} \in \mathbb{R}^{\mathcal{S}^* \cup \mathcal{T}}$$

$$\left[ \frac{\partial f(X)}{\partial X^*} \right]_{s^*, t} = \frac{\partial f(X)_t}{\partial X_s}$$

for all  $s \in \text{rec } \mathcal{S}$  and  $t \in \text{rec } \mathcal{T}$ .

The above definition assumes that  $\mathcal{S}^*$  and  $\mathcal{T}$  are orthogonal; if not, the axes in  $\mathcal{S}$  should be renamed to something else. For example, the second derivative (the Hessian) could be

$$\frac{\partial^2 f(X)}{\partial X^* \partial X^\dagger} \in \mathbb{R}^{\mathcal{S}^* \cup \mathcal{S}^\dagger \cup \mathcal{T}}$$

$$\left[ \frac{\partial^2 f(X)}{\partial X^* \partial X^\dagger} \right]_{r^*, s^\dagger, t} = \frac{\partial^2 f(X)_t}{\partial X_r \partial X_s}$$

for all  $r, s \in \text{rec } \mathcal{S}$  and  $t \in \text{rec } \mathcal{T}$ .

## 5.2 Differentials

We could derive rules like the chain rule and the sum and product rules, and use them to compute derivatives; however, ensuring that input and output shapes are orthogonal is inconvenient. Instead, we recommend the method of differentials (Magnus & Neudecker, 1985), which reduces renaming to a minimum.

The first-order Taylor approximation of  $f$  around  $X \in \mathbb{R}^{\mathcal{S}}$  is

$$f(X + H) \approx f(X) + \frac{\partial f(X)}{\partial X^*} \underset{\mathcal{S}^*}{\odot} H^* \quad H \in \mathbb{R}^{\mathcal{S}}.$$

The *differential* of  $f(X)$  with increment  $H$ , written  $\partial f(X; H)$ , is the second term of this approximation; we defer a formal definition to Appendix B.

For example,

- If  $\text{id}$  is the identity function, then

$$\begin{aligned} \text{id}(X + H) &= X + H \\ \partial \text{id}(X; H) &= H. \end{aligned} \tag{8a}$$

- If  $f(X) = X \underset{\text{ax}}{\odot} X$  where  $X \in \mathbb{R}^{\text{ax}}$ , then

$$\begin{aligned} f(X + H) &= (X + H) \underset{\text{ax}}{\odot} (X + H) \\ &= X \underset{\text{ax}}{\odot} X + 2X \underset{\text{ax}}{\odot} H + H \underset{\text{ax}}{\odot} H \\ \partial f(X; H) &= 2X \underset{\text{ax}}{\odot} H. \end{aligned} \tag{8b}$$

It's often more convenient to work directly with the expression  $X \odot_{\text{ax}} X$  instead of  $f(X)$ , and to write  $\partial(X \odot_{\text{ax}} X)$  for  $\partial f(X; H)$ . Then, since  $\partial X = \partial \text{id}(X; H) = H$ , we can write Eq. (8b) simply as

$$\partial(X \odot_{\text{ax}} X) = 2X \odot_{\text{ax}} \partial X$$

so that the  $H$  has been “hidden” inside  $\partial X$ . More generally, we can derive rules like the following:

$$\partial(U + V) = \partial U + \partial V \quad (9a)$$

$$\partial(U \odot V) = U \odot \partial V + V \odot \partial U \quad (9b)$$

$$\partial\left(\frac{U}{V}\right) = \frac{V \odot \partial U - U \odot \partial V}{V^2} \quad (9c)$$

$$\partial \sum_{\text{ax}} U = \sum_{\text{ax}} \partial U \quad (9d)$$

$$\partial(U \odot_{\text{ax}} V) = U \odot_{\text{ax}} \partial V + V \odot_{\text{ax}} \partial U \quad (9e)$$

$$\partial U_s = [\partial U]_s \quad (9f)$$

$$\partial U_{\text{ax} \rightarrow \text{ax}'} = [\partial U]_{\text{ax} \rightarrow \text{ax}'}. \quad (9g)$$

The chain rule for differentials is

$$\partial f(U) = \left. \frac{\partial f(X)}{\partial X^*} \right|_{X=U} \odot_{S^*} \partial U_{S \rightarrow S^*} \quad f: \mathbb{R}^S \rightarrow \mathbb{R}^{\mathcal{T}}. \quad (9h)$$

Recall that  $f$  can be lifted to shapes larger than  $S$ . In that case, the rule above still applies, but note that the contraction will still be over  $S$ . A special case of this is when  $S = \mathcal{T} = \emptyset$ :

$$\partial f(U) = \left. \frac{df(x)}{dx} \right|_{x=U} \odot \partial U \quad f: \mathbb{R} \rightarrow \mathbb{R}. \quad (9i)$$

For example, letting  $f(x) = \exp x$  gives the rule

$$\partial(\exp U) = \exp U \odot \partial U. \quad (9j)$$

Using these rules we can compute the differential of a wide variety of expressions. For example, the softmax operator:

$$\begin{aligned} \partial(\text{softmax } X)_{\text{ax}} &\stackrel{(5a)}{=} \partial\left(\frac{\exp X}{\sum_{\text{ax}} \exp X}\right) \\ &\stackrel{(9c)}{=} \frac{(\sum_{\text{ax}} \exp X) \odot \partial(\exp X) - \exp X \odot \partial(\sum_{\text{ax}} \exp X)}{(\sum_{\text{ax}} \exp X)^2} \\ &\stackrel{(9d)}{=} \frac{(\sum_{\text{ax}} \exp X) \odot \partial(\exp X) - \exp X \odot \sum_{\text{ax}} \partial(\exp X)}{(\sum_{\text{ax}} \exp X)^2} \\ &\stackrel{(9j)}{=} \frac{(\sum_{\text{ax}} \exp X) \odot \exp X \odot \partial X - \exp X \odot \sum_{\text{ax}} (\exp X \odot \partial X)}{(\sum_{\text{ax}} \exp X)^2} \\ &\stackrel{(2)}{=} \frac{(\sum_{\text{ax}} \exp X) \odot \exp X \odot \partial X - \exp X \odot (\exp X \odot_{\text{ax}} \partial X)}{(\sum_{\text{ax}} \exp X)^2} \end{aligned}$$

$$\begin{aligned}
&= \frac{\exp X}{\sum_{\text{ax}} \exp X} \odot \partial X - \frac{\exp X}{\sum_{\text{ax}} \exp X} \odot \left( \frac{\exp X}{\sum_{\text{ax}} \exp X} \odot \partial X \right) \\
&\stackrel{(5a)}{=} \underset{\text{ax}}{\text{softmax } X} \odot \partial X - \underset{\text{ax}}{\text{softmax } X} \odot (\underset{\text{ax}}{\text{softmax } X} \odot \partial X) \\
&= \underset{\text{ax}}{\text{softmax } X} \odot (\partial X - \underset{\text{ax}}{\text{softmax } X} \odot \partial X). \tag{10}
\end{aligned}$$

We stop when the only differentials left are  $\partial X$ .

### 5.3 Derivatives via differentials

If we can get  $\partial f(X)$  into so-called *canonical form*,

$$\partial f(X) = D \underset{\mathcal{S}^*}{\odot} \partial X^* + \text{const.} \tag{11}$$

where ‘‘const.’’ stands for terms not depending on  $\partial X$ , then by Magnus & Neudecker’s first identification theorem (Theorem 1 in Appendix B), we can conclude that

$$\frac{\partial f(X)}{\partial X^*} = D.$$

When trying to get expressions into canonical form, one helpful fact is that renaming can be thought of as contraction with an identity matrix. First we define the identity matrix with shape  $\text{ax} \times \text{ax}'$ :

$$[I_{\text{ax}, \text{ax}'}]_{\text{ax}(i), \text{ax}'(j)} = \begin{cases} 1 & i = j \\ 0 & i \neq j. \end{cases}$$

Then for any tensor  $A$  with an axis  $\text{ax}$ ,

$$A_{\text{ax} \rightarrow \text{ax}'} = I_{\text{ax}, \text{ax}'} \underset{\text{ax}}{\odot} A. \tag{12}$$

More specifically, if  $\partial X \in \mathbb{R}^{\mathcal{S}}$ , then

$$\partial X = I_{\mathcal{S}, \mathcal{S}^*} \underset{\mathcal{S}^*}{\odot} \partial X^* \tag{13}$$

and then Eq. (4) can usually be used to move the  $\underset{\mathcal{S}^*}{\odot} \partial X_{\mathcal{S} \rightarrow \mathcal{S}^*}$  to the outermost level of the expression.

Above, we found the differential of the softmax function; now let us find its derivative.

$$\begin{aligned}
\partial(\underset{\text{ax}}{\text{softmax } X}) &\stackrel{(10)}{=} \underset{\text{ax}}{\text{softmax } X} \odot (\partial X - \underset{\text{ax}}{\text{softmax } X} \odot \partial X) \\
&\stackrel{(13)}{=} \underset{\text{ax}}{\text{softmax } X} \odot (I_{\text{ax}, \text{ax}^*} \underset{\text{ax}^*}{\odot} \partial X^* - \underset{\text{ax}}{\text{softmax } X} \odot (I_{\text{ax}, \text{ax}^*} \underset{\text{ax}^*}{\odot} \partial X^*)) \\
&\stackrel{(4)}{=} (\underset{\text{ax}}{\text{softmax } X} \odot (I_{\text{ax}, \text{ax}^*} - \underset{\text{ax}}{\text{softmax } X} \odot I_{\text{ax}, \text{ax}^*})) \underset{\text{ax}^*}{\odot} \partial X^* \\
&\stackrel{(12)}{=} \underset{\text{ax}}{\text{softmax } X} \odot (I_{\text{ax}, \text{ax}^*} - (\underset{\text{ax}}{\text{softmax } X})^*) \underset{\text{ax}^*}{\odot} \partial X^*.
\end{aligned}$$

This is in canonical form, so we have

$$\frac{\partial}{\partial X}(\underset{\text{ax}}{\text{softmax } X}) = \underset{\text{ax}}{\text{softmax } X} \odot (I_{\text{ax}, \text{ax}^*} - (\underset{\text{ax}}{\text{softmax } X})^*). \tag{14}$$

## 5.4 Lifting

Recall that  $f^{\mathcal{S}'}$  is the lift of  $f: \mathbb{R}^{\mathcal{S}} \rightarrow \mathbb{R}^{\mathcal{T}}$  with  $\mathcal{S}'$ , and in most contexts we can simply write  $f$  instead of  $f^{\mathcal{S}'}$ . However, derivatives are one place where some extra caution is in order. To lighten notation, let's write  $g$  for the derivative of  $f$ :

$$g(X) = \frac{\partial f(X)}{\partial X^*}.$$

Recall that the chain rule (9h) works under lifting, so

$$\partial f^{\mathcal{S}'}(X) = g^{\mathcal{S}'}(X) \odot_{\mathcal{S}^*} \partial X_{\mathcal{S} \rightarrow \mathcal{S}^*}.$$

But the contraction is only over  $\mathcal{S}^*$ , so it would be incorrect to conclude that  $\frac{\partial f^{\mathcal{S}'}(X)}{\partial X^*} = g^{\mathcal{S}'}(X)$ . The derivative of a lift is *not* the lift of a derivative. We must rename and contract  $\mathcal{S}'$  as well:

$$\begin{aligned} \partial f^{\mathcal{S}'}(X) &\stackrel{(9h)}{=} g^{\mathcal{S}'}(X) \odot_{\mathcal{S}^*} \partial X_{\mathcal{S} \rightarrow \mathcal{S}^*} \\ &\stackrel{(13)}{=} g^{\mathcal{S}'}(X) \odot_{\mathcal{S}^*} (I_{\mathcal{S}', \mathcal{S}'^*} \odot_{\mathcal{S}'^*} \partial X_{\mathcal{S} \cup \mathcal{S}' \rightarrow (\mathcal{S} \cup \mathcal{S}')^*}) \\ &\stackrel{(3)}{=} (g^{\mathcal{S}'}(X) \odot I_{\mathcal{S}', \mathcal{S}'^*}) \odot_{(\mathcal{S} \cup \mathcal{S}')^*} \partial X_{\mathcal{S} \cup \mathcal{S}' \rightarrow (\mathcal{S} \cup \mathcal{S}')^*} \\ \frac{\partial f^{\mathcal{S}'}(X)}{\partial X^*} &= g^{\mathcal{S}'}(X) \odot I_{\mathcal{S}', \mathcal{S}'^*}. \end{aligned}$$

In general, then, the derivative of a lift is the lift of the derivative, multiplied by the identity matrix for the new axes. Intuitively, this is because the derivative is a linear transformation—before lifting, a transformation from  $\mathcal{S}^*$  to  $\mathcal{T}$ . When lifting to  $\mathcal{S} \cup \mathcal{S}'$ , this transformation must also be lifted, which is what multiplication by  $I_{\mathcal{S}', \mathcal{S}'^*}$  accomplishes.

## 5.5 Extended example

As a more elaborate example, we find the derivative of self-attention. For brevity, we omit the factor  $\frac{1}{\sqrt{|\text{key}|}}$ , and we write  $\alpha = \text{softmax}_{\text{seq}}(Q \odot_{\text{key}} K)$ .

$$\begin{aligned} \partial \text{Att}(Q, K, V) &\stackrel{(7)}{=} \partial(\alpha \odot_{\text{seq}} V) \\ &\stackrel{(9e)}{=} \alpha \odot_{\text{seq}} \partial V + V \odot_{\text{seq}} \partial \alpha. \end{aligned} \tag{15}$$

Focus first on the first term, which is the only term depending on  $\partial V$ :

$$\begin{aligned} \alpha \odot_{\text{seq}} \partial V &\stackrel{(13)}{=} \alpha \odot_{\text{seq}} ((I_{\text{seq}, \text{seq}^*} \odot I_{\text{val}, \text{val}^*}) \odot_{\text{seq}^* \text{val}^*} \partial V^*) \\ &\stackrel{(4)}{=} ((\alpha \odot_{\text{seq}} I_{\text{seq}, \text{seq}^*}) \odot I_{\text{val}, \text{val}^*}) \odot_{\text{seq}^* \text{val}^*} \partial V^* \\ &\stackrel{(12)}{=} (\alpha_{\text{seq} \rightarrow \text{seq}^*} \odot I_{\text{val}, \text{val}^*}) \odot_{\text{seq}^* \text{val}^*} \partial V^* \\ \frac{\partial}{\partial V^*} \text{Att}(Q, K, V) &= \alpha_{\text{seq} \rightarrow \text{seq}^*} \odot I_{\text{val}, \text{val}^*}. \end{aligned}$$



Next, focus on the second term of Eq. (15):

$$\begin{aligned} V \odot_{\text{seq}} \partial \alpha &\stackrel{(10)}{=} V \odot_{\text{seq}} (\alpha \odot (\partial(Q \odot_{\text{key}} K) - \alpha \odot_{\text{seq}} \partial(Q \odot_{\text{key}} K))) \\ &\stackrel{(9e)}{=} V \odot_{\text{seq}} (\alpha \odot (Q \odot_{\text{key}} \partial K + K \odot_{\text{key}} \partial Q - \alpha \odot_{\text{seq}} (Q \odot_{\text{key}} \partial K + K \odot_{\text{key}} \partial Q))). \end{aligned} \quad (16)$$

Keeping only terms depending on  $\partial Q$ , we get

$$\begin{aligned} V \odot_{\text{seq}} (\alpha \odot (K \odot_{\text{key}} \partial Q - \alpha \odot_{\text{seq}} (K \odot_{\text{key}} \partial Q))) \\ &\stackrel{(13)}{=} V \odot_{\text{seq}} (\alpha \odot (K \odot_{\text{key}} (I_{\text{key}, \text{key}^*} \odot_{\text{key}^*} \partial Q^*) - \alpha \odot_{\text{seq}} (K \odot_{\text{key}} (I_{\text{key}, \text{key}^*} \odot_{\text{key}^*} \partial Q^*)))) \\ &\stackrel{(4)}{=} \left( V \odot_{\text{seq}} (\alpha \odot (K \odot_{\text{key}} I_{\text{key}, \text{key}^*} - \alpha \odot_{\text{seq}} (K \odot_{\text{key}} I_{\text{key}, \text{key}^*})) \right) \odot_{\text{key}^*} \partial Q^* \\ &\stackrel{(4)}{=} \left( V \odot_{\text{seq}} (\alpha \odot (K_{\text{key} \rightarrow \text{key}^*} - \alpha \odot_{\text{seq}} K_{\text{key} \rightarrow \text{key}^*})) \right) \odot_{\text{key}^*} \partial Q^* \\ \frac{\partial}{\partial Q^*} \text{Att}(Q, K, V) &= V \odot_{\text{seq}} (\alpha \odot (K_{\text{key} \rightarrow \text{key}^*} - \alpha \odot_{\text{seq}} K_{\text{key} \rightarrow \text{key}^*})). \end{aligned}$$

Similarly, keeping only terms from Eq. (16) depending on  $\partial K$ , we get

$$\begin{aligned} V \odot_{\text{seq}} (\alpha \odot (Q \odot_{\text{key}} \partial K - \alpha \odot_{\text{seq}} (Q \odot_{\text{key}} \partial K))) \\ = \left( V \odot_{\text{seq}} (\alpha \odot (Q^* \odot_{\text{seq}} I_{\text{seq}, \text{seq}^*}) - \alpha \odot_{\text{seq}} (Q^* \odot_{\text{seq}} I_{\text{seq}, \text{seq}^*})) \right) \odot_{\text{key}^*} \partial K^* \\ \frac{\partial}{\partial K^*} \text{Att}(Q, K, V) = V \odot_{\text{seq}} (\alpha \odot (Q^* \odot_{\text{seq}} I_{\text{seq}, \text{seq}^*}) - \alpha \odot_{\text{seq}} (Q^* \odot_{\text{seq}} I_{\text{seq}, \text{seq}^*})). \end{aligned}$$

## 6 Alternatives and Related Work

### 6.1 Index notations

Among alternatives to standard vector and matrix notation, the most common one is index notation as used in physics (Ricci & Levi-Civita, 1900). Related notations are used in other fields as well (Harshman, 2001). In this notation, axes are ordered, and every equation is written in terms of tensor components. If an index appears on both sides of an equation, then the equation must hold for each value of the index, and the Einstein summation convention (Einstein, 1916) is that if an index appears twice on one side and not on the other, there is an implicit summation over that index.

$$\text{Attention}: \mathbb{R}^{d_k} \times \mathbb{R}^{n \times d_k} \times \mathbb{R}^{n \times d_v} \rightarrow \mathbb{R}^{d_v}$$

$$[\text{Attention}(Q, K, V)]_k = \text{softmax}_i \left( \frac{Q_j K_{ij}}{\sqrt{d_k}} \right) V_{ik}.$$

Because  $k$  appears on both sides, the equation must hold over all values of this index. But because  $i$  and  $j$  occur twice on only the right-hand side, they are both summed over. We would have to define exactly what the  $i$  under the softmax means ( $i$  is bound inside the softmax and free outside it), and since softmax doesn't distribute over addition, we would need to modify the summation convention so that the summation over  $j$  occurs inside the softmax.

Aside from these correctable issues, this notation scales very well to more than two axes and is concise and unambiguous. But it does not solve the main problem we set out to solve, which is that ordered axes force the author and reader to remember the purpose of each axis. The indices do act as symbolic names for axes

(indeed, in *abstract* index notation (Penrose & Rindler, 1984), they really are symbols, not variables), but they are temporary names; they could be totally different in the next equation. It would be up to the author to choose to use consistent names, and to do so correctly.

A second issue is that because it depends on repetition of indices to work, index notation can be more verbose than our notation, particularly for reductions and contractions:

$$\begin{array}{ll} C = \max_i A_i & C = \max_{\text{ax}} A \\ C = A_i B_i & C = A \underset{\text{ax}}{\odot} B. \end{array}$$

Finally, index notation requires us to write out all indices explicitly. So if we wanted to lift attention to minibatches ( $b \in [B]$ ), multiple heads ( $h \in [H]$ ) and multiple query tokens ( $i' \in [n']$ ), we would write:

$$\begin{aligned} \text{Attention: } \mathbb{R}^{B \times H \times n' \times d_k} \times \mathbb{R}^{B \times H \times n \times d_k} \times \mathbb{R}^{B \times H \times n \times d_v} &\rightarrow \mathbb{R}^{B \times H \times n' \times d_v} \\ [\text{Attention}(Q, K, V)]_{bhi'k} &= \text{softmax}_i \left( \frac{Q_{bhi'j} K_{bhij}}{\sqrt{d_k}} \right) V_{bhik}. \end{aligned}$$

We could adopt a convention that lifts a function on tensors to tensors that have extra axes to the *left*, but such conventions tend to lead to messy reordering and squeezing/unsqueezing of axes. Named axes make such conventions unnecessary.

## 6.2 Graphical notations

In the graphical notation of Penrose (1971), a node in the graph stands for a tensor, and its incident edges stand for its indices. The edges are ordered from left to right. An edge connecting two nodes denotes contraction. The notation of Alsberg (1997) is similar, except that edges are named, not ordered.

Graphs are commonly used in machine learning for representing probability models (Koller & Friedman, 2009). A node in the graph stands for a random variable, and an edge or hyperedge stands for a dependency between variables. If random variables have finite domains, then a (hyper)edge with  $r$  endpoints can be thought of as an  $r$ -th order tensor. A graph can then be thought of as a product and contraction. Extensions that allow for a choice between two subgraphs (e.g., Minka & Winn, 2008) can be thought of as addition.

Our assessment of graphical notations like these is that, on the positive side, they have obvious value for visualization, and they at least have the potential to represent indices in a purely unordered way. On the negative side, these notations seem best suited for representing linear functions, and even for this purpose, some other practical considerations are that drawing pictures requires more effort from the author, and that pictures will have a less transparent relationship with their implementation in most programming languages.

## 6.3 Relational algebra

In relational algebra (Codd, 1970), the basic objects are sets of  $r$ -tuples, which could be thought of as tensors of order  $r$  with Boolean-valued entries. In the original formulation, the members of the tuples, which correspond to axes, were both named *and* ordered, although later definitions (e.g. Pirotte, 1982) made them unordered.

Probabilistic variants of relational algebra also exist (e.g. Dey & Sarkar, 1996; Fuhr & Rölleke, 1997), whose relations would correspond to tensors of probabilities.

While relational algebra and tensor notations are designed for totally different purposes, the notation of relational algebra generally has a similar flavor to ours (for example, our contraction operator is similar to the  $\bowtie$  operator, and our renaming operator is the same as the  $\rho$  operator).

## 6.4 Programming languages

One response to the notation presented here, as well as the alternative notations mentioned in this section, is that research papers in machine learning should simply present models as code rather than equations. But we argue that a model’s mathematical specification should abstract away from details of its implementation.

Conceptually, it is important to have a distinct specification to define what makes an implementation (both the original implementation and any reimplementations) correct or incorrect. If the implementation is its own specification, it cannot be correct or incorrect; it will be “not even wrong.”

Practically, abstracting away from implementation is important because we do not want the interpretation of research papers to be subject to differences across programming languages and libraries, or versions thereof. For example, PyTorch’s `Dropout2d` on order-3 tensors has one behavior in versions 1.11 and 1.13, but another behavior in 1.10, 1.12, and future versions. It would be problematic for correct understanding of a paper to depend on such differences.

## 7 Conclusions

Named tensor notation is a system of formal notation for representing operations between tensors in a non-ambiguous way while remaining intuitive for practitioners. The system is motivated by challenges that arise from taking notation designed for applied linear algebra and using it for representing neural networks, as demonstrated through examples of canonical deep-learning components such as attention and layer normalization. However, named tensors are not limited to specifying neural networks. We have also explained how to integrate our notation with Magnus & Neudecker (1985)’s method of differentials for matrix calculus. While there are other conventions that such as index notation that have some usage in the machine learning community, these conventions either lack the conciseness of named tensors or are not well-suited to non-linear operations. For these reasons, we encourage members of the machine learning community to try out named tensor notation for teaching, research, and software documentation.

## Acknowledgements

We would like to thank Ekin Akyürek, Justin Bayer, Tongfei Chen, Chu-Cheng Lin, Colin McDonald, Adam Poliak, Matt Post, Chung-chieh Shan, Nishant Sinha, and Yee Whye Teh for their input to this document (or the ideas in it). We also thank the anonymous TMLR reviewers for their feedback, which substantially improved the quality of the paper, especially Section 5.

This material is based upon work supported by the National Science Foundation under Grants No. CCF-2019291 and DMS-2134157, as well as a Sloan Fellowship, Simons Investigator Fellowship, DARPA grant W911NF2010021, and DOE grant DE-SC0022199. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding agencies.

## References

- Bjørn K. Alsberg. A diagram notation for n-mode array equations. *Journal of Chemometrics*, 11:251–266, 1997. doi: 10.1002/(SICI)1099-128X(199705)11:3<251::AID-CEM472>3.0.CO;2-Q.
- Matt Bauman. Axisarrays, 2018. URL <https://github.com/JuliaArrays/AxisArrays.jl>. Open-source software.
- Tongfei Chen. Typesafe abstractions for tensor operations. In *Proc. 8th ACM SIGPLAN International Symposium on Scala*, pp. 45–50, 2017. doi: 10.1145/3136000.3136001.
- E. F. Codd. A relational model of data for large shared data banks. *Communications of the ACM*, 13(6): 377–387, June 1970. doi: 10.1145/362384.362685.

- Zachary DeVito. Named tensors using first-class dimensions in PyTorch, 2023. URL <https://github.com/facebookresearch/torchdim>. Open-source software.
- Debabrata Dey and Sumit Sarkar. A probabilistic relational model and algebra. *ACM Transactions on Database Systems*, 21(3):339–369, September 1996. doi: 10.1145/232753.232796.
- A. Einstein. Die Grundlage der allgemeinen Relativitätstheorie. *Annalen der Physik*, 354(7):769–880, 1916. doi: 10.1002/andp.19163540702.
- Norbert Fuhr and Thomas Rölleke. A probabilistic relational algebra for the integration of information retrieval and database systems. *ACM Transactions on Information Systems*, 15(1):32–66, January 1997. doi: 10.1145/239041.239045.
- Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016.
- Richard A. Harshman. An index formalism that generalizes the capabilities of matrix notation and algebra to n-way arrays. *Journal of Chemometrics*, 15:689–714, 2001. doi: 10.1002/cem.665.
- Stephan Hoyer and Joe Hamman. xarray: N-D labeled arrays and datasets in Python. *Journal of Open Research Software*, 5(1):10, 2017. doi: 10.5334/jors.148.
- JAX authors. Named axes and easy-to-revise parallelism, 2021. URL [https://jax.readthedocs.io/en/latest/notebooks/xmap\\_tutorial.html](https://jax.readthedocs.io/en/latest/notebooks/xmap_tutorial.html).
- Daphne Koller and Nir Friedman. *Probabilistic Graphical Models: Principles and Techniques*. MIT Press, 2009.
- Soeren Laue, Matthias Mitterreiter, and Joachim Giesen. Computing higher order derivatives of matrix and tensor expressions. In *Proc. NeurIPS*, pp. 2750–2759, 2018. URL <https://proceedings.neurips.cc/paper/2018/file/0a1bf96b7165e962e90cb14648c9462d-Paper.pdf>.
- Jan R. Magnus and H. Neudecker. Matrix differential calculus with applications to simple, Hadamard, and Kronecker products. *Journal of Mathematical Psychology*, 29(4):474–492, 1985. doi: 10.1016/0022-2496(85)90006-9.
- Tom Minka and John Winn. Gates. In *Proc. NeurIPS*, pp. 1073–1080, 2008. URL <https://papers.nips.cc/paper/3379-gates>.
- Adam Paszke, Daniel D. Johnson, David Duvenaud, Dimitrios Vytiniotis, Alexey Radul, Matthew J. Johnson, Jonathan Ragan-Kelley, and Dougal Maclaurin. Getting to the point: Index sets and parallelism-preserving autodiff for pointful array programming. *Proc. ACM on Programming Languages*, 5(ICFP), August 2021. doi: 10.1145/3473593.
- R. Penrose and W. Rindler. *Spinors and space-time*, volume 1. Cambridge University Press, 1984.
- Roger Penrose. Applications of negative dimensional tensors. In D. J. A. Welsh (ed.), *Combinatorial Mathematics and its Applications*, pp. 221–244. Academic Press, 1971.
- Alain Pirotte. A precise definition of basic relational notations and of the relational algebra. *ACM SIGMOD Record*, 13(1):30–45, 1982. doi: 10.1145/984514.984516.
- PyTorch Contributors. PyTorch documentation, 2022. URL <https://pytorch.org/docs/1.12/>. version 1.12.
- M. M. G. Ricci and T. Levi-Civita. Méthodes de calcul différentiel absolu et leurs applications. *Mathematische Annalen*, 54:125–201, 1900. doi: 10.1007/BF01454201.
- Alex Rogozhnikov. Einops: Clear and reliable tensor manipulations with Einstein-like notation. In *Proc. ICLR*, 2022. URL <https://openreview.net/pdf?id=oapKSVM2bcj>.

Alexander Rush. Named tensors, 2019. URL <https://github.com/harvardnlp/NamedTensor>. Open-source software.

Nishant Sinha. Tensor shape (annotation) library, 2018. URL <https://github.com/ofnote/tsalib>. Open-source software.

Torch Contributors. Named tensors, 2019. URL [https://pytorch.org/docs/stable/named\\_tensor.html](https://pytorch.org/docs/stable/named_tensor.html). PyTorch documentation.

L. R. Tucker. The extension of factor analysis to three-dimensional matrices. In H. Gulliksen and N. Frederiksen (eds.), *Contributions to Mathematical Psychology*, pp. 110–127. Holt, Rinehart and Winston, New York, 1964.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Proc. NeurIPS*, pp. 5998–6008, 2017. URL <https://proceedings.neurips.cc/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf>.

Yuxin Wu and Kaiming He. Group normalization. In *Proc. ECCV*, 2018. URL [https://openaccess.thecvf.com/content\\_ECCV\\_2018/papers/Yuxin\\_Wu\\_Group\\_Normalization\\_ECCV\\_2018\\_paper.pdf](https://openaccess.thecvf.com/content_ECCV_2018/papers/Yuxin_Wu_Group_Normalization_ECCV_2018_paper.pdf).

## A Extended Examples

### A.1 Transformer

We define a Transformer used autoregressively as a language model. The input is a sequence of one-hot vectors, from which we compute word embeddings and positional encodings:

$$\begin{aligned}
 I &\in \{0, 1\}^{\text{seq} \times \text{vocab}} & \sum_{\text{vocab}} I &= 1 \\
 W &= (E \underset{\text{vocab}}{\odot} I) \sqrt{|\text{layer}|} & E &\in \mathbb{R}^{\text{vocab} \times \text{layer}} \\
 P &\in \mathbb{R}^{\text{seq} \times \text{layer}} \\
 P_{\text{seq}(p), \text{layer}(i)} &= \begin{cases} \sin((p-1)/10000^{(i-1)/|\text{layer}|}) & i \text{ odd} \\ \cos((p-1)/10000^{(i-2)/|\text{layer}|}) & i \text{ even.} \end{cases}
 \end{aligned}$$

Then we use  $L$  layers of self-attention and feed-forward neural networks:

$$\begin{aligned}
 X^0 &= W + P \\
 T^1 &= \text{LayerNorm}^1(\text{SelfAtt}^1(X^0) + X^0) \\
 X^1 &= \text{LayerNorm}^{1'}(\text{FFN}^1(T^1) + T^1) \\
 &\vdots \\
 T^L &= \text{LayerNorm}^L(\text{SelfAtt}^L(X^{L-1}) + X^{L-1}) \\
 X^L &= \text{LayerNorm}^{L'}(\text{FFN}^L(T^L) + T^L) \\
 O &= \underset{\text{vocab}}{\text{softmax}}(E \underset{\text{layer}}{\odot} X^L)
 \end{aligned}$$

where LayerNorm, SelfAtt and FFN are defined below.

Layer normalization ( $l = 1, 1', \dots, L, L'$ ):

$$\begin{aligned}
 \text{LayerNorm}^l &: \mathbb{R}^{\text{layer}} \rightarrow \mathbb{R}^{\text{layer}} \\
 \text{LayerNorm}^l(X) &= \underset{\text{layer}}{\text{standardize}}(X) \odot \gamma^l + \beta^l
 \end{aligned}$$

$$\beta^l, \gamma^l \in \mathbb{R}^{\text{layer}}$$

We defined attention in §4.3; the Transformer uses multi-head self-attention, in which queries, keys, and values are all computed from the same sequence.

$$\begin{aligned} \text{SelfAtt}^l: \mathbb{R}^{\text{seq} \times \text{layer}} &\rightarrow \mathbb{R}^{\text{seq} \times \text{layer}} \\ \text{SelfAtt}^l(X) &= Y \end{aligned}$$

where

$$\begin{aligned} |\text{seq}| &= |\text{seq}'| \\ |\text{key}| = |\text{val}| &= |\text{layer}|/|\text{heads}| \\ Q &= W^{l,Q} \underset{\text{layer}}{\odot} X_{\text{seq} \rightarrow \text{seq}'} & W^{l,Q} &\in \mathbb{R}^{\text{heads} \times \text{layer} \times \text{key}} \\ K &= W^{l,K} \underset{\text{layer}}{\odot} X & W^{l,K} &\in \mathbb{R}^{\text{heads} \times \text{layer} \times \text{key}} \\ V &= W^{l,V} \underset{\text{layer}}{\odot} X & W^{l,V} &\in \mathbb{R}^{\text{heads} \times \text{layer} \times \text{val}} \\ M &\in \mathbb{R}^{\text{seq} \times \text{seq}'} \\ M_{\text{seq}(i), \text{seq}'(j)} &= \begin{cases} 0 & i \leq j \\ -\infty & \text{otherwise} \end{cases} \\ Y &= W^{l,O} \underset{\substack{\text{heads} \\ \text{val}}}{\odot} \text{Attention}(Q, K, V, M)_{\text{seq}' \rightarrow \text{seq}} & W^{l,O} &\in \mathbb{R}^{\text{heads} \times \text{val} \times \text{layer}} \end{aligned}$$

Feedforward neural networks:

$$\begin{aligned} \text{FFN}^l: \mathbb{R}^{\text{layer}} &\rightarrow \mathbb{R}^{\text{layer}} \\ \text{FFN}^l(X) &= X^2 \end{aligned}$$

where

$$\begin{aligned} X^1 &= \text{relu}(W^{l,1} \underset{\text{layer}}{\odot} X + b^{l,1}) & W^{l,1} &\in \mathbb{R}^{\text{hidden} \times \text{layer}} & b^{l,1} &\in \mathbb{R}^{\text{hidden}} \\ X^2 &= W^{l,2} \underset{\text{hidden}}{\odot} X^1 + b^{l,2} & W^{l,2} &\in \mathbb{R}^{\text{layer} \times \text{hidden}} & b^{l,2} &\in \mathbb{R}^{\text{hidden}} \end{aligned}$$

## A.2 LeNet

$$\begin{aligned} X^0 &\in \mathbb{R}^{\text{batch} \times \text{chans}[c_0] \times \text{height} \times \text{width}} \\ T^1 &= \text{relu}(\text{Conv}^1(X^0)) \\ X^1 &= \text{MaxPool}^1(T^1) \\ T^2 &= \text{relu}(\text{Conv}^2(X^1)) \\ X^2 &= \text{MaxPool}^2(T^2)_{(\text{height}, \text{width}, \text{chans}) \rightarrow \text{layer}} \\ X^3 &= \text{relu}(W^3 \underset{\text{layer}}{\odot} X^2 + b^3) & W^3 &\in \mathbb{R}^{\text{hidden} \times \text{layer}} & b^3 &\in \mathbb{R}^{\text{hidden}} \\ O &= \text{softmax}(W^4 \underset{\text{classes}}{\odot} X^3 + b^4) & W^4 &\in \mathbb{R}^{\text{classes} \times \text{hidden}} & b^4 &\in \mathbb{R}^{\text{classes}} \end{aligned}$$

As an alternative to the flattening operation in the equation for  $X^2$ , we could have written

$$X^2 = \text{MaxPool}^2(T^2)$$

$$X^3 = \text{relu}(W^3 \underset{\substack{\text{height} \\ \text{width} \\ \text{chans}}}{\odot} X^2 + b^3) \quad W^3 \in \mathbb{R}^{\text{hidden} \times \text{height} \times \text{width} \times \text{chans}}$$

The convolution and pooling operations are defined as follows:

$$\text{Conv}^l(X) = \text{Conv2d}(X; W^l, b^l)_{\text{chans}' \rightarrow \text{chans}}$$

where

$$W^l \in \mathbb{R}^{\text{chans}'[c_l] \times \text{chans}[c_{l-1}] \times \text{kh}[kh_l] \times \text{kw}[kw_l]}$$

$$b^l \in \mathbb{R}^{\text{chans}'[c_l]}$$

and

$$\text{MaxPool}^l(X) = \text{MaxPool2d}_{ph^l, ph^l}(X).$$

## B Differentiation: Formal Definitions

The following definition and theorem come directly from the paper by Magnus & Neudecker (1985), but generalized to named tensors.

For any  $X \in \mathbb{R}^S$ , we write  $\|X\| = \text{norm}_S X$ .

**Definition 9.** Let  $f: S \rightarrow \mathbb{R}^T$  where  $S \subseteq \mathbb{R}^S$ . Let  $A$  be an interior point of  $\mathbb{R}^S$ , that is, for some  $r > 0$ ,  $B(A; r) = \{X \mid \|X - A\| < r\} \subseteq S$ . If there is a tensor  $D(A) \in \mathbb{R}^{S^* \cup T}$  and  $R(A, H) \in \mathbb{R}^T$  such that

$$f(A + H) = f(A) + D(A) \underset{S^*}{\odot} H_{S \rightarrow S^*} + R(A, H)$$

for all  $H \in \mathbb{R}^S$  with  $\|H\| < r$ , and

$$\lim_{H \rightarrow \mathbf{0}} \frac{R(A, H)}{\|H\|} = \mathbf{0},$$

then  $f$  is said to be *differentiable* at  $A$ ; the tensor

$$\partial f(A; H) = D(A) \underset{S^*}{\odot} H_{S \rightarrow S^*}$$

is then called the *(first) differential of  $f$  at  $A$  with increment  $H$* .

Magnus & Neudecker give their (first) identification theorem twice, once for vector-to-vector functions and once for matrix-to-matrix functions (but omitting vector-to-matrix and matrix-to-vector functions). Here, we only need one version, which works for functions from tensors to tensors of any shape.

**Theorem 1.** Let  $f: S \rightarrow \mathbb{R}^T$ , where  $S \subseteq \mathbb{R}^S$ , be differentiable at  $A \in S$ . Let  $D(X) \in \mathbb{R}^{S^* \cup T}$ . Then

$$\text{for all } H, \partial f(A; H) = D(X) \underset{S^*}{\odot} H_{S \rightarrow S^*} \quad \text{iff} \quad \left. \frac{\partial f(X)}{\partial X^*} \right|_{X=A} = D(X).$$

## C $\LaTeX$ Macros

Many of the  $\LaTeX$  macros used in this document are available in the style file `namedtensor.sty`, available on CTAN at <https://ctan.org/pkg/namedtensor>. To use it, put

```
\usepackage{namedtensor}
```

in the preamble of your L<sup>A</sup>T<sub>E</sub>X source file (after `\documentclass{article}` but before `\begin{document}`).

We write axis names in sans-serif font. To make this easier, `\name{ax}` prints an axis name (like this: `ax`), and `\ndef{ax}{ax}` defines a macro `ax` that does the same.

- Binary operators

- Use `A \ndot{ax} B` for contraction:  $A \underset{ax}{\odot} B$ . You can use `\` to stack up several names.

- In general, you can use `\nbin` to make a new binary operator with a name under it: `A \nbin{ax}{\star} B` gives you  $A \underset{ax}{\star} B$ .

- Functions

- Use `\nsum{ax} A` for summation:  $\sum_{ax} A$ .

- In general, you can use `\nfun` to make a function with a name under it: `\nfun{ax}{qux} A` gives you  $\underset{ax}{\text{qux}} A$ .