Vectors & Matrices : suffix notation

A common early vectors question is "prove that $\mathbf{a} \times (\mathbf{b} \times \mathbf{c}) = (\mathbf{a} \cdot \mathbf{c})\mathbf{b} - (\mathbf{a} \cdot \mathbf{b})\mathbf{c}$", and a typical first attempt will go something like:

"Since $\mathbf{x} \times \mathbf{y} = (x_2y_3 - x_3y_2, x_3y_1 - x_1y_3, x_1y_2 - x_2y_1)$, the first components of each side are

LHS : $a_2(\mathbf{b}\times\mathbf{c})_3 - a_3(\mathbf{b}\times\mathbf{c})_2 = a_2(b_1c_2 - b_2c_1) - a_3(b_3c_1 - b_1c_3) = a_2b_1c_2 - a_2b_2c_1 - a_3b_3c_1 + a_3b_1c_3$

RHS : $(a_1c_1 + a_2c_2 + a_3c_3)b_1 - (a_1b_1 + a_2b_2 + a_3b_3)c_1 = a_2b_1c_2 + a_3b_1c_3 - a_2b_2c_1 - a_3b_3c_1$

These agree. The other components are similar."

This is true, but wouldn't it be nicer if we could prove it more concisely, or for every component at the same time? Then we wouldn't have to rely on the other components being 'similar'. The idea is to try to prove a general result for the i^{th} component, where we can then substitute in $i = 1, 2, 3$ at the end, as we wish.

Note. I'm going to work in three dimensions throughout, to keep the explanations simpler.

Suffixes

Let's use suffixes to refer to the entries in a vector or a matrix. Thus, for example:

- Given a vector $\mathbf{v}$, let v_i be its i^{th} component.
- Given a matrix $\mathbf{M}$, let m_{ij} be its $(i,j)^{\text{th}}$ entry - that is, in the i^{th} row and j^{th} column.

Note. We don't have to (and won't always) use i or ij here. We could write v_k for the k^{th} component of $\mathbf{v}$, or $m_{\alpha\beta}$ for the entry in the α^{th} row and β^{th} column of $\mathbf{M}$. (Just don't try pronouncing α^{th}.)

Examples

So far, that's just notation, and it should be familiar. For example

1. Scalar ('dot') product: $\mathbf{a} \cdot \mathbf{b} = a_1b_1 + a_2b_2 + a_3b_3 = \sum_{i=1}^{3} a_ib_i$.
2. Multiplying a matrix by a vector. Given a matrix $\mathbf{M}$ and a vector $\mathbf{v}$, when we work out the first component of $\mathbf{Mv}$, we dot the first row of $\mathbf{M}$ with $\mathbf{v}$. Similarly, when we work out the i^{th} component of $\mathbf{Mv}$, we dot the i^{th} row of $\mathbf{M}$ with $\mathbf{v}$:

$$\mathbf{Mv} = \begin{pmatrix} \cdots & \cdots & \cdots \\ m_{i1} & m_{i2} & m_{i3} \\ \cdots & \cdots & \cdots \end{pmatrix} \begin{pmatrix} v_1 \\ v_2 \\ v_3 \end{pmatrix} = \begin{pmatrix} \cdots \\ m_{i1}v_1 + m_{i2}v_2 + m_{i3}v_3 \\ \cdots \end{pmatrix}$$

   So the i^{th} component of $\mathbf{Mv}$ is $(\mathbf{Mv})_i = m_{i1}v_1 + m_{i2}v_2 + m_{i3}v_3 = \sum_{j=1}^{3} m_{ij}v_j$.

   Note. This is an important example, and it appears often. For example, if we have the vector equation $\mathbf{v}' = \mathbf{Mv}$, then in suffixes it becomes $v'_i = \sum m_{ij}v_j$.
3. Multiplying a matrix by a matrix. Given matrices $\mathbf{M}$ and $\mathbf{N}$, when we work out the $(i,j)^{\text{th}}$ entry of $\mathbf{MN}$, we dot the i^{th} row of M with the j^{th} column of $\mathbf{N}$.

   Do this, and write it as $\sum$ of something concise involving suffixes.

   Common mistake. The answer is *not* $\sum M_{ij}N_{ij}$, and if you followed the "dotting a row with a column" reasoning, you shouldn't have got this. However, when dealing with suffix expressions later on, many people forget where $(MN)_{ij}$ comes from, and just write down this wrong expression. Don't do that!

Summation convention

In each of the examples above, the summation is performed over a suffix that appears twice: $\sum_{i=1}^{3} a_i b_i$ or $\sum_{j=1}^{3} m_{ij} v_j$. This happens so often that we turn it into a rule, and a very useful piece of shorthand: "if we appear to be multiplying two terms containing the same suffix, then actually we sum those terms over that suffix."

So, rather than writing $\sum_{i=1}^{3} a_i b_i$, we can just write $a_i b_i$, and think to ourselves "we are summing over i here". However, we see that $\sum_{i=1}^{3} a_i b_i = \sum_{j=1}^{3} a_j b_j = \sum_{\alpha=1}^{3} a_\alpha b_\alpha = \ldots = a_1 b_1 + a_2 b_2 + a_3 b_3$. The summation letter is just a name internal to the sum, and we get the same result no matter which we use. Therefore, in summation convention, $a_i b_i = a_j b_j = a_\alpha b_\alpha = \ldots$.

A suffix that appears twice like this is called a "dummy" suffix, as its name doesn't really matter – providing it doesn't cause confusion with other suffixed terms. (Confusion could arise if we had $a_i b_i c_j$ and wanted to rename i. We could pick any letter we liked – except j, that is.)

A suffix that appears on its own in a term is called a "free" suffix, because we are free to choose its value if we want to work out a specific component. For example: $\mathbf{a} = (\mathbf{b} \cdot \mathbf{c})\mathbf{d}$. Here, each component of $\mathbf{d}$ is scaled by the same factor to get $\mathbf{a}$, so $a_1 = (\mathbf{b} \cdot \mathbf{c}) d_1$, etc. In suffix notation, we write $a_i = b_j c_j d_i$. The suffix i is "free", and we are free to set $i = 1, 2, 3$ to work out specific components. But $b_j c_j$ is always the same value, and we have no choice to make.

Notes

Remember that the summation convention applies when "we appear to be *multiplying* two terms containing the same suffix". If we're *summing* two terms containing the same suffix, then there is no summation. For example, the i^{th} component of $\mathbf{a} + \mathbf{b}$ is $a_i + b_i$, just the sum of the individual i^{th} components. It certainly isn't $\sum(a_i + b_i)$.

So, if $a_i b_i$ means we sum, how do we refer to the single product of two i^{th} components, such as "a_i times b_i"? We could say "the i^{th} product $a_i b_i$" or "$a_i b_i$ (no summation)".

Common mistake. It is perfectly fine for a final answer to be, say, an i^{th} component. However, some people tend to panic and think that a final answer should be a definite number or vector, and end an answer with something like "$\ldots = a_i = a_1 + a_2 + a_3 = \mathbf{a}$". Both final equalities here are incorrect: a_i is simply the i^{th} component, not the sum of all three, and neither of those is the vector itself! If the answer is a_i, then so be it, and leave it as that.

Examples

1. The earlier examples: $\mathbf{a} \cdot \mathbf{b} = a_i b_i$ (or $a_j b_j$, etc), and $(\mathbf{M}\mathbf{v})_i = m_{ij} v_j$ (or $m_{i\alpha} v_\alpha$, etc).
2. $a_i b_j c_j = \sum_j a_i b_j c_j = a_i (\sum_j b_j c_j) = (\mathbf{b} \cdot \mathbf{c}) a_i$
3. In terms of vectors and matrices, what are $a_\alpha b_\beta c_\alpha d_\beta$ and $m_{jk} v_j v_k$?

"Beware the three i-ed monster"

What about three suffixes? What might $a_i b_i c_i$ mean? It *could* be defined to mean $\sum_{i=1}^{3} a_i b_i c_i$.

However, this turns out to cause more trouble than it's worth, as it very rarely happens that we multiply three things together simultaneously. For example, if we're multiplying three matrices $\mathbf{A}, \mathbf{B}, \mathbf{C}$, we actually just do $(\mathbf{AB})\mathbf{C}$ or $\mathbf{A}(\mathbf{BC})$.

It would also conflict with our current two-suffix summation convention, which would like to pair two of those i suffixes up into a sum. But then, does $a_i b_i c_i$ mean $(a_i b_i) c_i = (\mathbf{a} \cdot \mathbf{b}) c_i$ or $a_i (b_i c_i) = (\mathbf{b} \cdot \mathbf{c}) a_i$? It doesn't work.

So the simple rule here is: "if you are using this summation convention and have three (or more) subscripts the same, then you have gone wrong". It's no use claiming that *you* know which ones are being summed: you have gone wrong. If we really do need to sum over a three-suffixed thing, then we have to write $\sum_i a_i b_i c_i$. That's perfectly valid – we just can't use our handy summation convention.

What if we're given $x = a_i b_i$ and $y = c_i d_i$ and are asked to work out xy? Remember that these are *dummy* labels, and that $c_i d_i = c_j d_j = c_\alpha d_\alpha = \ldots$. So let's write y as $c_j d_j$, which is the same thing, and we get $xy = a_i b_i c_j d_j$.

If in doubt, we can revert to writing in the $\sum$ in the expressions. So $x = \sum_i a_i b_i$ and $y = \sum_i c_i d_i$, and we get $xy = (\sum_i a_i b_i)(\sum_i c_i d_i) = (\sum_i a_i b_i)(\sum_j c_j d_j)$. So because it's just being used to count through a sum, we can *always* relabel a dummy suffix. If we think that there's even a chance of a collision, rename the dummy i to p, or k to ξ. (But don't cause new collisions, of course!)

A useful check. If we're trying to work out, say, the $(i, j)^{\text{th}}$ entry of some matrix using suffix notation. Then at every stage in our calculations, the subscripts i and j must appear in each term being summed exactly once, while any other subscript that appears must do so exactly twice. If your answer is going wrong, stop and count the number of times each appears.

Exercise. Show that for matrices $\mathbf{A}, \mathbf{B}, \mathbf{C}$ (of suitable sizes), we do have $(\mathbf{AB})\mathbf{C} = \mathbf{A}(\mathbf{BC})$.

Kronecker Delta

Next, δ_{ij}. This has probably been defined as: $\delta_{ij} = 1$ if $i = j$, and 0 if $i \neq j$.

Consider the expression $\delta_{ij} v_j$. This is a sum over the j suffix, so is: $\delta_{i1} v_1 + \delta_{i2} v_2 + \delta_{i3} v_3$. We don't actually know what i is. If $i = 1$, then only the δ_{i1} term is non-zero, and the sum is v_1. Similarly if $i = 2$ or 3. In other words, whatever i is, the sum $\delta_{ij} v_j$ equals v_i.

Consider the sillier expression $\delta_{ij} b_{pqjr}$. This thing b is some nonsense I've made up – it's probably some weird four-dimensional thing. But we can do the sum: i is fixed at 'some value', and so for most of the time j won't equal it, giving us 0. But, for the one occasion when $j = i$, the δ part equals 1, and the b part is b_{pqir}, whatever that means. Thus $\delta_{ij} b_{pqjr} = b_{pqir}$.

We get the useful mnemonic: "δ_{ij} turns a j into an i".

What if we apply it to something with an i instead? What is $\delta_{ij} v_i$? For exactly the same reasons, we get v_j. So in fact, our useful mnemonic is "δ_{ij} turns a j into an i, or an i into a j, whichever it meets".

But hang on, what if it meets both of them? What is $\delta_{ij} a_i b_j$? Then we're summing over both i and j. We could do the i sum first, then the j sum. Or we could do the j sum first, then the i sum. Do both – and make sure you get the same answer!

You should find that the δ_{ij} is "used up" in turning an i into a j (or vice versa), and we can't use it again. (Indeed, $\delta_{ij} a_i b_j$ is *not* equal to $a_j b_i$.)

Delta as a matrix?

You might be thinking (hopefully) that the expression $\delta_{ij} v_j$ looks an awful lot like the matrix–vector expression from early on: $m_{ij} v_j$, which is the i^{th} component of $\mathbf{Mv}$.

So what is happening with $\delta_{ij} v_j$? Well, δ_{ij} has two suffixes, so looks like it might be the component of a matrix. And if it's 0 when $i \neq j$ and 1 when $i = j$, it sounds a lot like the identity matrix $\mathbf{I}$. So let's say that it is.

In which case, $\delta_{ij}v_j$ must surely be the i^{th} component of $\mathbf{Iv}$, which is just v_i. That agrees with our suffix calculation above, but have we really spent all this effort just to find a fiddly way of writing $\mathbf{Iv} = \mathbf{v}$?

It turns out to be more useful than that. For example, suppose we are trying to use suffixes to simplify some matrix–vector equation involving something like $\mathbf{v} + \mathbf{Mv}$. Then we could write it as $(\mathbf{I} + \mathbf{M})\mathbf{v}$, by replacing the $\mathbf{v}$ by $\mathbf{Iv}$, which is the effect of the δ. We will often be taking vectors and doing things to them, and this will be very helpful.

Epsilon

Now, this one looks a bit weird at first. It has three suffixes, so if we want to think of it as a sort of "matrix" then we'll have to write down a cube of numbers. Its definition is:

ϵ_{ijk} equals 1 if $ijk = 123, 231, 312$, equals -1 if $ijk = 321, 132, 213$, and equals 0 otherwise (i.e., if two suffixes are the same).

One very useful feature of ϵ_{ijk} is that it allows us to introduce anti-symmetry into our expressions. For example, when we work out a determinant or vector product, half of the terms are subtracted. And indeed, we can now write both of those in a concise summation form.

Example. We can write $(\mathbf{a} \times \mathbf{b})_i = \epsilon_{ijk}a_jb_k$.

The best way to convince yourself would be to try it. Let's set $i = 1$. The right hand side is a sum of nine terms, as we're summing over all $j, k = 1, 2, 3$. However, if two suffixes are ever equal, the ϵ term is 0. So, since $i = 1$, we lose the terms with j or k equal to 1. And we also lose the terms with $j = k = 2$ or $j = k = 3$. This leaves just $\epsilon_{123}a_2b_3 + \epsilon_{132}a_3b_2$, which (checking the definition of ϵ) is $a_2b_3 - a_3b_2$, which is correct.

Exercise. Let $\mathbf{M}$ be a matrix whose rows are the vectors $\mathbf{a}, \mathbf{b}, \mathbf{c}$. Then $\det \mathbf{M} = \epsilon_{ijk}a_ib_jc_k$. Verify this. (The sum contains 27 terms, but many of them are 0.)

A useful formula

δ and ϵ join up in the following very useful formula: $\epsilon_{ijk}\epsilon_{ilm} = \delta_{jl}\delta_{km} - \delta_{jm}\delta_{kl}$.

I've included this section entirely to try to help people remember the formula, since many often get the thing mixed up. It would be dangerous to remember it exactly as written there, since we might not get those exact subscripts: we might get $\epsilon_{\alpha\beta\gamma}\epsilon_{\alpha\mu\nu}$ or even $\epsilon_{imk}\epsilon_{ijl}$.

Below is the mnemonic I use, which you might find helpful.

By the definition of ϵ, we can cycle the labels around: so $\epsilon_{abc} = \epsilon_{bca} = \epsilon_{cab}$. So, if we have a double-ϵ expression with a shared suffix then we can cycle the subscripts around until the shared suffix is the first of each. That is, if we see $\epsilon_{kij}\epsilon_{lmi}$, write it as $\epsilon_{ijk}\epsilon_{ilm}$.

Then think to yourself "same minus different". That is, once the first subscripts of each ϵ are the same, the first two δ terms are formed from the two second-place suffixes paired up, then the two third-place suffixes. The second two δ terms (the ones we subtract) are from the "differents", i.e. the two pairings of a second-place suffix with a third-place suffix.

I hope that makes sense!

Exercise. We began with $\mathbf{a} \times (\mathbf{b} \times \mathbf{c}) = (\mathbf{a} \cdot \mathbf{c})\mathbf{b} - (\mathbf{a} \cdot \mathbf{b})\mathbf{c}$. Verify this with suffixes.

Please send any corrections or comments to me at glt1000@cam.ac.uk