Never Be Confused by a Change of Basis Again

(Maybe)

Richard Chapling v1 26 November 2019

In this handout, we're going to introduce a little-considered concept, that makes it much easier to understand change of basis. It is the author's hope that the combination of these with some basic diagram algebra will be enough to reduce the usual confusion this subject induces to nearly null.

1 Some basic diagram algebra

In this section I am going to discuss some simple ideas about *diagrams.*[1]

Definition 1. A *diagram* is a finite, labelled directed graph. That is, it is a set of vertices, each given a name (which will be that of a specific space or object), joined by arrows, each of which is also given a name (which we will be taking as the names of maps between these spaces).

We will assume our diagrams are connected: there is always a sequence of arrows linking one vertex to another (even if some of these arrows go in "the wrong direction"). We call a path *valid* if the arrows point "the right way" along it.

All of this so far has sounded very complicated, but that is because it is so general. We only need a few general principles and some basic examples here, but it is worth introducing the *ideas* as early as possible, to give you a chance to get used to them.

1.1 The most basic diagrams

One vertex with no arrows forms a diagram, but this is not very interesting. The most basic diagram with an arrow in it is

$$A \xrightarrow{f} B.$$

You probably have seen this without even thinking of it as a diagram. We interpret this as *the function f sends every element of A to an element of B.* The vertex at the tail of an arrow is the *domain* of the function, that at the head is the *codomain.*[2]

Compositions are defined by following one arrow, then another:

$$A \xrightarrow{f} B \xrightarrow{g} C$$

this means that $g \circ f$ makes sense. On the other hand, if the right-hand arrow points to B instead of C, we can't compose the maps: the domain and codamain don't match as they must.

[1] I promise that this is both related, and useful.

[2] Not "range"! The range is the *actual* image set $f(A)$ of the function, the codomain is the set that f *may* send elements of A to.

This is the main use we have for these diagrams: they tell us what goes from where to where and what compositions make sense. They are a "map" of the region our objects and and functions live in, telling us how to get from one place to another.[3]

However, these sorts of diagrams still don't tell us very much: their most important rôle is in demonstrating clearly what maps can be composed. The real power of diagrams come when we allow them to contain loops.

1.2 Diagrams with loops

Definition 2. If there are two paths from A to B in a diagram, the loop formed by them is said to *commute* if the composition along each one of them results the an identical function $A \to B$. If all such loops commute, the diagram is called a *commutative diagram.*

These are by far the most common diagrams: it is often very surprising to meet a diagram that does not commute. In particular, *every diagram in this handout will commute.*

The simplest commutative diagrams tend to be triangle-shaped: the definition of composition is equivalent to this diagram commuting:

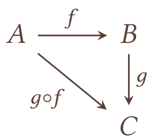

Another example we can give is the inverse function: suppose that $f : A \to B$ and $g : B \to A$. Then g is an inverse to f if and only if the following two diagrams commute:

We can combine these into a single diagram:

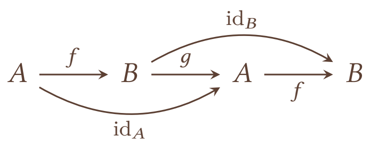

Note that only the way the diagram is connected matters: it may be clearer if it is drawn in a different way: the previous diagram is the same as

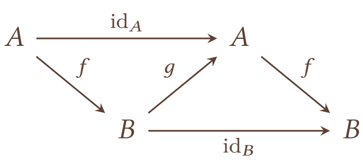

and

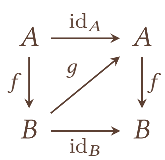

We are thus dealing with a "Tube map" map, not an "OS map" map or a "road map" map.

[3] To unnecessarily quote Roger Penrose, "where by 'map' here I mean the thing made of stiff paper that you take with you when you go hiking, not the mathematical notion of 'map' " (*The Road to Reality*, p. 189).

Moreover, when a function is the identity, it is common to replace the two copies of the set joined by the identity function by one copy of the set, since the identity "does nothing":

where the ellipses and arrows represent the rest of the diagram. The inverse diagram is thus simply written as

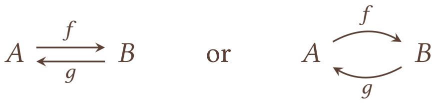

depending on preference.

The useful thing about commutative diagrams is that they give us *more equality*: we can stick some information about things being equal into a commutative diagram, and discover that a whole lot more is true. The rather simple example above gives a property of the identitiy maps, for example: both triangles commute, so the square commutes and hence $f \circ \mathrm{id}_A = \mathrm{id}_B \circ f$.

One last thing we can do with diagrams is *glue them together*: while attaching one vertex to another to compose maps is useful, if two diagrams share not only a vertex, but two vertices and the arrow joining them: for example,

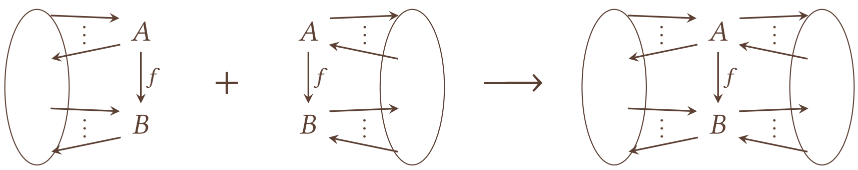

This idea is entirely general: *any* diagrams containing a shared subdiagram[4] can be put together in this way. Notice, in particular, that to keep the new diagram commutative, the functions must match as well as the vertices.

2 Choosing a basis

After that deviation into the abstract, let us return to the (reasonably) concrete.

Before we deal with two bases, we had better understand what one basis does. Recall that a *basis* B for a vector space V is a subset of V which is

Spanning Every element of V can be written as a linear combination of elements of B.

Linearly independent The only linear combination of elements of B that gives the zero vector is the one in which every coefficient is 0.

[4] A subdiagram is a collection of vertices in a diagram together with all the arrows between them.

This means that each vector in V can be written *in a unique way* as a linear combination of elements of B: there is exactly one way of writing[5]

$$u = \sum_{i=1}^{n} b_i u_i, \qquad b_i \in B, u_i \in F.$$

The *dimension* of V is the cardinality of a basis. As you will learn in IB Linear Algebra, all bases have the same cardinality, so this definition makes sense.

Unusually among vector spaces, F^n (the vector space of n-tuples of elements of F with elementwise addition and scalar multiplication) has a distinguished, canonical basis, namely the n vectors

$$\mathbf{e}_i = (0, \dots, 0, \underbrace{1}_{i\text{th position}}, 0, \dots, 0) \qquad i \in \{1, \dots, n\}.$$

In three dimensions, for example, these vectors are

$$\mathbf{e}_1 = (1, 0, 0), \quad \mathbf{e}_2 = (0, 1, 0), \quad \mathbf{e}_3 = (0, 0, 1),$$

and a general vector in F^n can be written as

$$\mathbf{u} = (u_1, u_2, u_3) = \mathbf{e}_1 u_1 + \mathbf{e}_2 u_2 + \mathbf{e}_3 u_3.$$

These tuples are frequently written as *column vectors*: an $n \times 1$ table of numbers. For example, in three dimensions,

$$\mathbf{u} = (u_1, u_2, u_3) = \begin{pmatrix} u_1 \\ u_2 \\ u_3 \end{pmatrix}.$$

As you should know, a general n-dimensional vector space V over a field F is isomorphic to F^n.[6] But there are lots of different ways of choosing this isomorphism.

Indeed, it turns out that picking an isomorphism is "the same" as picking an ordered basis. We make this precise as follows:

Proposition 3. *Let V be an n-dimensional vector space over F.*

1. *Each isomorphism $F^n \to V$ produces a unique ordered basis of V.*
2. *Each ordered basis of V gives rise to a unique isomorphism $V \to F^n$.*

This result is essentially why we can use matrices; we will explain this in the next section.

We ask that the associated basis is ordered because the standard basis of F^n is ordered, an extra piece of structure that is worth preserving, since it makes the isomorphism unique rather than just unique up to permutations.

[5] I am deliberately writing this with the basis vectors before the coefficients; why will become clear later, when matrices are introduced.

[6] Remember that an homomorphism is a structure-preserving map. The structure on an F-vector space is linearity, so a vector space homomorphism is an F-linear map. An isomorphism is a bijective homomorphism, and in the context of vector spaces, this is an invertible linear map.

Proof. 1. Let ϕ be the isomorphism, and define $b_i = \phi(\mathbf{e}_i)$ for each $i \in \{1, \dots, n\}$. The uniqueness of the $\mathbf{e}_i$ ensures that this definition is unique, and now all we have to do is show that b_i are a basis for V. But ϕ is linear, so

$$\phi(\mathbf{u}) = \phi\Big(\sum_{i=1}^{n} \mathbf{e}_i u_i\Big) = \sum_{i=1}^{n} \phi(\mathbf{e}_i) u_i = \sum_{i=1}^{n} b_i u_i.$$

Since ϕ is surjective, all vectors in V are $\phi(\mathbf{u})$ for some $\mathbf{u} \in F^n$, while since ϕ is injective, such a $\mathbf{u}$ is unique (and hence so are the coefficients u_i). Thus $\{b_i\}_{i=1}^{n}$ is a basis.

2. Let $(b_i)_{i=1}^{n}$ be an ordered basis. Define a function $\psi \colon V \to F^n$ by $\psi(b_i) = \mathbf{e}_i$ and extending linearly, i.e.

$$\psi(u) = \psi\Big(\sum_{i=1}^{n} b_i u_i\Big) = \sum_{i=1}^{n} \psi(b_i) u_i = \sum_{i=1}^{n} \mathbf{e}_i u_i.$$

Since the ordering of the basis is fixed, the function ψ is unique. We now need to show that ψ is an isomorphism. We can cheat here: its inverse function is actually ϕ from the first part of the proof, and we know that that is an isomorphism, so so too must ψ be. □

We can therefore write ϕ_B for the isomorphism $F^n \to V$ given by $\phi(\mathbf{e}_i) = b_i$. The numbers u_i that have appeared throughout in $u = \sum_{i=1}^{n} b_i u_i$ are called the *components* of u, although strictly they are the *components of u with respect to the basis B*, or the B-components of u.

We write $\mathbf{u} = \phi_B(u)$ for the vector in F^n corresponding to u. We can also write $\phi_B(u) = [u]_B$ if we need to be especially clear about which basis we are using.

The diagram for the basis isomorphism can be written as

$$\begin{array}{c} V \\ {\scriptstyle \phi_B}\big\uparrow \\ F^n \end{array}$$

where we will keep this orientation consistent for readability.

3 Matrices

An $m \times n$ *matrix* is a linear map $F^n \to F^m$. It is conventional to represent matrices by two-dimensional arrays of numbers, just as elements of F^n are written as column vectors.

Since we have isomorphisms to general vector spaces, this suggests that a matrix can be used to describe a linear map in a given basis. Indeed this is the case: suppose we have

- An n-dimensional F-vector space V and an m-dimensional space W,
- A linear map $\alpha \colon V \to W$,
- A basis B of V and a basis C of W,
- An $m \times n$ matrix $\mathbf{A}$ that we want to be the matrix of α.

We can put all this together the following diagram:

$$\begin{array}{ccc} V & \xrightarrow{\alpha} & W \\ {\scriptstyle \phi_B}\uparrow & & \uparrow{\scriptstyle \phi_C} \\ F^n & \xrightarrow[\mathbf{A}]{} & F^m \end{array}$$

(and, notice this is the *only possible way* to make such a diagram!) If this diagram commutes, we find that

$$\alpha\phi_B = \phi_C\mathbf{A}.$$

In particular, this has to hold for every element of F^n. For a such a vector, $\mathbf{u} \in F^n$, this means that

$$\alpha(\phi_B(\mathbf{u})) = \alpha(u) = \sum_{j=1}^{n} \alpha(b_j)u_j$$

is equal to

$$\phi_C(\mathbf{Au}) = \phi_C\left(\sum_{j=1}^{n} u_j\mathbf{Ae}_j\right) = \sum_{j=1}^{n} \phi_C(\mathbf{Ae}_j)u_j$$

To go further, we need to know what $\mathbf{A}$ does to the $\mathbf{e}_j$. Since $\mathbf{Ae}_j \in F^m$, and $\{\mathbf{e}_i\}_{i=1}^m$ is a basis of F^m, there *must* be a collection of elements $A_{ij} \in F$ so that

$$\mathbf{Ae}_j = \sum_{i=1}^{m} \mathbf{e}_i A_{ij}.$$

These A_{ij} then slot into the rows and columns of an array, the rows being indexed by i, the columns by j:

$$\mathbf{A} = \begin{pmatrix} A_{11} & \cdots & A_{1j} & \cdots & A_{1n} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ A_{i1} & \cdots & A_{ij} & \cdots & A_{in} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ A_{m1} & \cdots & A_{mj} & \cdots & A_{mn} \end{pmatrix}$$

Then $\mathbf{Au} = \mathbf{A}\sum_{j=1}^m \mathbf{e}_j u_j = \sum_{i=1}^m \sum_{j=1}^n \mathbf{e}_i A_{ij} u_j$, or in terms of arrays,

$$\mathbf{A} = \begin{pmatrix} A_{11} & \cdots & A_{1j} & \cdots & A_{1n} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ A_{i1} & \cdots & A_{ij} & \cdots & A_{in} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ A_{m1} & \cdots & A_{mj} & \cdots & A_{mn} \end{pmatrix} \begin{pmatrix} u_1 \\ \vdots \\ u_j \\ \vdots \\ u_n \end{pmatrix} = \begin{pmatrix} A_{11}u_1 + \cdots + A_{1j}u_j + \cdots + A_{1n}u_n \\ \vdots \\ A_{i1}u_1 + \cdots + A_{ij}u_j + \cdots + A_{in}u_n \\ \vdots \\ A_{m1}u_1 + \cdots + A_{mj}u_j + \cdots + A_{mn}u_n \end{pmatrix}$$

Going back to the linear map, we find

$$\sum_{j=1}^{n} \alpha(b_j)u_j = \sum_{j=1}^{n} \phi_C(\mathbf{Ae}_j)u_j = \sum_{j=1}^{n} \phi_C(\mathbf{e}_i A_{ij})u_j = \sum_{j=1}^{n} \phi_C(\mathbf{e}_i)A_{ij}u_j = \sum_{j=1}^{n} c_i A_{ij} u_j.$$

These formulae hold for any u_j, so we must have

$$\alpha(b_j) = c_i A_{ij},$$

which we then make the *definition* of the components of α with respect to the bases B and C. We call $\mathbf{A}$ the matrix of α with respect to these bases; we sometimes write $\mathbf{A} = [\alpha]_C^B$.

The last equation has another interpretation: the product is the linear combination $\mathbf{A}_{1\cdot}u_1 + \cdots + \mathbf{A}_{n\cdot}u_n$, where $\mathbf{A}_{i\cdot}$ is the ith column of $\mathbf{A}$. Hence the image in F^m is the span of the columns. This is not totally unexpected: they are the images of the basis vectors, after all.

3.1 Operations on matrices

Addition and scalar multiplication Addition and scalar multiplication of linear maps is defined as one might expect: given linear maps $\sigma, \tau \colon V \to W$, we can define their sum and scalar multiplication by

$$(\sigma + \tau)(u) = \sigma(u) + \tau(u) \qquad (\lambda.\sigma)(u) = \lambda.(\sigma(u)).$$

Hence the set of linear maps $V \to W$, which we write $L(V, W)$, is also a vector space.

The corresponding results for the matrices are that they are a vector space, with termwise operations:

$$(S + T)_{ij} = S_{ij} + T_{ij} \qquad (\lambda S)_{ij} = \lambda S_{ij}.$$

But we can only do any of this when the spaces on which the matrices act are the same. Therefore, *we can add matrices when their dimensions are the same.*

Composition How do we multiply matrices? This is the wrong question at present: what we need to ask first is: what should multiplying matrices mean?

"Multiplication" of linear maps is really composition: we just aren't writing the brackets around the argument. Suppose we have vector spaces U, V, W with bases $(b_i)_{k=1}^p, (c_j)_{j=1}^m, (d_i)_{i=1}^n$, and linear maps $\sigma \colon U \to V$ and $\tau \colon V \to W$. Then, writing $[\sigma]_{ij}$ for the components of σ and so on, we have

$$(\sigma\tau)(u) = \sum_{i=1}^{m}\sum_{k=1}^{p} d_i[\sigma\tau]_{ik}u_k,$$

by definition, but also

$$\sigma(\tau(u)) = \sum_{k=1}^{p}\sigma(\tau(b_k))u_k = \sum_{k=1}^{p}\sigma\Big(\sum_{j=1}^{n}c_j[\tau]_{jk}\Big)u_k = \sum_{j=1}^{n}\sum_{k=1}^{p}\sigma(c_j)[\tau]_{jk}u_k = \sum_{i=1}^{m}\sum_{j=1}^{n}\sum_{k=1}^{p}d_i[\sigma]_{ij}[\tau]_{jk}u_k,$$

and since this is true for each u_k and the d_i are a basis, we find

$$[\sigma\tau]_{ik} = \sum_{j=1}^{m}[\sigma]_{ij}[\tau]_{jk},$$

and this was the only possible way the composition could work given our previous results. But of course $[\sigma]_{ij}$ are the elements of the associated matrix, so we also obtain the method of matrix multiplication,

$$(ST)_{ik} = \sum_{j=1}^{m} S_{ij}T_{jk} :$$

the ikth entry of ST is the found by taking the products of elements of the ith row of S with the corresponding ones in the kth column of T, then summed.

Unwinding all this, we have found that given an $m_1 \times n_1$ matrix $\mathbf{S}$ and a $m_2 \times n_2$ matrix $\mathbf{T}$, they can be multiplied if $n_1 = m_2$. If this multiplication can be carried out, the result is an $m_1 \times n_2$ matrix.

4 Change of basis

Usually the topic that people find most confusing is changing basis. Hopefully, with all the machinery that we have built up in this handout, we can avoid most of the confusing parts now. We will be using both diagrams and component calculations.

Suppose first that we have a vector space V and two bases $B = (b_i)$ and $B' = (b'_{i'})$: we will denote indices in the second basis with primes to further distinguish associated objects. Firstly, we can write down a diagram with both basis isomorphisms, and the matrix $\mathbf{P}$ that maps between them:

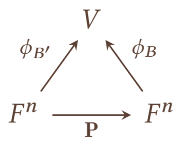

so $\phi_B = \mathbf{P}\phi_{B'}$. What linear map is $\mathbf{P}$ the matrix of? We can obtain it by using separating V into two:

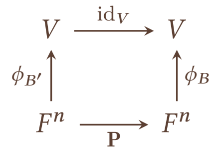

So we see immediately that ψ is the matrix of the identity with respect to the bases B' and B:

$$\mathbf{P} = [\mathrm{id}_V]_B^{B'},$$

In terms of components, we have

$$u = u_j b_j = u'_{i'} b'_{i'} = b_j P_{ji'} u'_{i'},$$

and from this we can read off the relationship between components:

$$u_j = P_{ji'} u'_{i'} \iff u'_{i'} = (P^{-1})_{i'j} u_j$$

which can be written in terms of the column vectors as

$$\mathbf{u}' = \mathbf{P}^{-1}\mathbf{u} \iff \mathbf{u} = \mathbf{P}\mathbf{u}'. \tag{1}$$

In the component matrix notation,

$$[u]_{B'} = [\mathrm{id}_V]_{B'}^{B} [u]_B = ([\mathrm{id}_V]_B^{B'})^{-1} [u]_B.$$

By taking u itself to be a basis vector, we also find the relationship between the basis vectors:

$$b'_{i'} = b_j P_{ji'}.$$

the basis vectors transform in the *opposite* way to the components of vectors. This has to be the case, since the vector u itself does not depend on the basis chosen:

$$u' = b'_{i'} u'_{i'} = (b_j P_{ji'})((P^{-1})_{i'k} u_k) = b_j (P_{ji'} (P^{-1})_{i'k}) u_k = b_j \delta_{jk} u_k = b_j u_j,$$

as expected.

Now suppose we have a linear map $\alpha\colon V \to W$, and bases B, B' for V, C, C' for W, related by

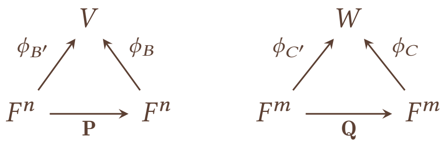

Now, at last, everything comes together. Suppose that $[\alpha]_C^B = \mathbf{A}$. Then we have

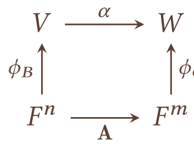

as before. To find $[\alpha]_{C'}^{B'}$, we want to be able to relate this diagram to the corresponding one for the new basis:

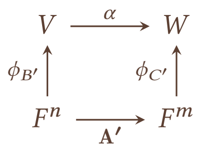

But we can do this precisely by pasting them both together with the two change-of-basis diagrams. *And there is only one way to do this!*

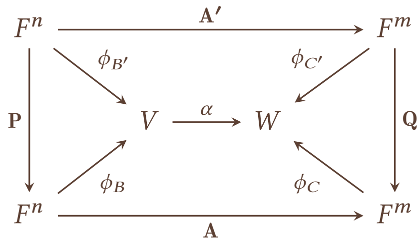

There are several other ways to write this diagram, however, one of which might be more memorable:

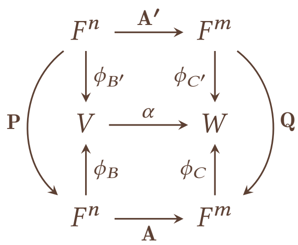

or

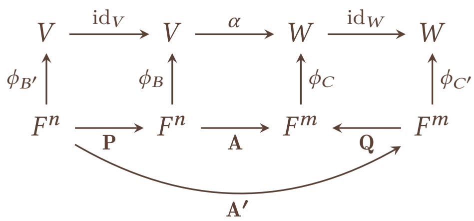

Whichever way we draw this diagram, that it is commutative implies that $\mathbf{QA'} = \mathbf{AP}$, or

$$\mathbf{A'} = \mathbf{Q}^{-1}\mathbf{AP}. \tag{2}$$

In terms of the component matrix notation,

$$[\alpha]_{C'}^{B'} = [\mathrm{id}_W]_{C'}^{C}[\alpha]_C^B[\mathrm{id}_V]_B^{B'} = ([\mathrm{id}_W]_C^{C'})^{-1}[\alpha]_C^B[\mathrm{id}_V]_B^{B'}.$$

Notice that, as it must be, (2) is compatible with the formulae (1) for how the components change: since $\mathbf{u}' = \mathbf{P}^{-1}\mathbf{u}$, if $\mathbf{w} = \mathbf{A}\mathbf{u}$, then

$$\mathbf{w}' = \mathbf{Q}^{-1}\mathbf{w} = \mathbf{Q}^{-1}\mathbf{A}\mathbf{u} = \mathbf{Q}^{-1}\mathbf{A}\mathbf{P}\mathbf{P}^{-1}\mathbf{u} = \mathbf{A}'\mathbf{u}'$$

as expected.

Summary Let B, B' be bases of V, Define change-of-basis matrices by $b_i = b'_{j'}P_{j'i}$ and $c_k = c'_{l'}Q_{l'k}$. Then we have the following change-of-basis formulae for components of vectors in V and matrices of linear maps $V \to W$:

$$u'_{i'} = (P^{-1})_{i'j}u_j \qquad\qquad A'_{i'j'} = (Q^{-1})_{i'j}A_{kl}P_{lj'}$$
$$\mathbf{u}' = \mathbf{P}^{-1}\mathbf{u} \qquad\qquad \mathbf{A}' = \mathbf{Q}^{-1}\mathbf{A}\mathbf{P}.$$

In terms of the component matrix notation for vectors and linear maps,

$$[u]_{B'} = [\mathrm{id}_V]^B_{B'}[u]_B \qquad\qquad [\alpha]^{B'}_{C'} = [\mathrm{id}_W]^C_{C'}[\alpha]^B_C[\mathrm{id}_V]^{B'}_B.$$

N.B. We have chosen to use the convention that the *basis vectors* are the thing transformed by $\mathbf{P}$, $b'_{i'} = b_jP_{ji'}$; then the components must be transformed by $\mathbf{P}^{-1}$ to preserve the value of u. There is an alternative, equally legitimate viewpoint, where $\mathbf{P}$ transforms the *components*, and the basis vectors are then transformed by $\mathbf{P}^{-1}$. We would then replace our formulae by

$$b'_{i'} = b_j(P^{-1})_{ji'}, \qquad u'_{i'} = P_{i'j}u_j, \qquad A'_{i'j'} = Q_{i'j}A_{kl}(P^{-1})_{lj'}$$
$$\mathbf{u}' = \mathbf{P}\mathbf{u} \qquad \mathbf{A}' = \mathbf{Q}\mathbf{A}\mathbf{P}^{-1}.$$

Always check how the transformation is defined.[7]

[7]**N.B. to the N.B.** Beware: this is not the same as the distinction between *active* and *passive transformations* in applied mathematics and physics. What we are discussing here are passive transformations: the vector u *itself* does not change, but we transform the basis vectors in B and so the components $[u]_B$ also transform. By contrast, active transformations *actually change the vector*, $u \mapsto f(u)$, while the basis stays the same.