

WMEx

WARWICK MATHEMATICS EXCHANGE

MA251

ALGEBRA I

2023, April 20th

Desync, aka The Big Ree

Contents

Table of Contents	i
1 Review	1
1.1 Mathematical Interpretation	1
1.2 Basis Vectors, Span & Linear Independence	2
1.3 Linear Transformations	3
1.3.1 Transformations as Matrix-Vector Multiplication	3
1.4 Composition as Matrix-Matrix Multiplication	5
1.5 Change of Basis	7
1.6 Transformations in Different Bases	10
1.7 Eigenvectors	13
2 Jordan Canonical Form	19
2.1 Generalised Eigenspaces	19
2.2 Cayley-Hamilton Theorem	20
2.3 Calculating Minimal Polynomials	20
2.4 Jordan Chains	21
2.5 Computing the Jordan Canonical Form	23
2.6 Review	32
3 Matrix Functions	32
3.1 Matrix Powers	32
3.2 Lagrange Interpolation	33
3.3 Matrix Exponentials	34
3.3.1 Recurrence Relations	34
3.3.2 Differential Equations	35
4 Bilinear Maps	37
4.1 Bilinear Forms	39
4.2 Quadratic Forms	40
4.3 Bases for Quadratic Forms	41
4.4 The Gram-Schmidt Process	43
4.5 Orthogonal Transformations	46
4.6 Orthonormal Bases for Bilinear Forms	48
4.7 Reduction of Second Degree Polynomial Equations	50
4.8 Singular Value Decomposition	52
5 Sesquilinear Forms	55
6 Operators on Hilbert Spaces	55
7 Finitely Generated Abelian Groups	55
7.1 Review	55
7.2 Free Abelian Groups	58
7.3 Unimodular Smith Normal Form	60
7.4 Subgroups of Free Abelian Groups	63
7.5 General Finitely Generated Abelian Groups	64
7.6 Finite Abelian Groups	66

Introduction

In *Algebra I: Advanced Linear Algebra*, we continue our work on vector spaces from MA106 *Linear Algebra*. We begin by introducing the concept of matrix polynomials before dipping into the spectral theory of matrices with generalised eigenspaces and the Jordan canonical form. We then cover bilinear maps and quadratic forms, before decomposing and classifying finitely generated abelian groups.

This document is intended to broadly cover all the topics within the Algebra I module. All knowledge and algorithms contained within the module guide for MA106 will be assumed as prior knowledge in this document. For a recap of Linear Algebra, you can see the module guides for MA106, but I also recommend viewing [my reference book](#), where you can find connections between different areas conveniently hyperlinked together in one file.

Disclaimer: I make *absolutely no guarantee* that this document is complete nor without error. In particular, any content covered exclusively in lectures (if any) will not be recorded here. This document was written during the 2022 academic year, so any changes in the course since then may not be accurately reflected.

Notes on formatting

New terminology will be introduced in *italics* when used for the first time. Named theorems will also be introduced in *italics*. Important points will be **bold**. Common mistakes will be underlined. The latter two classifications are under my interpretation. YMMV.

Content not taught in the course will be outlined in the margins like this. Anything outlined like this is not examinable, but has been included as it may be helpful to know alternative methods to solve problems.

The table of contents above, and any inline references are all hyperlinked for your convenience.

Scalars are written in lowercase italics, *c*, or using greek letters.

Vectors are written in lowercase bold, \mathbf{v} , or rarely overlined, \overleftarrow{v} , where more contrast or clarity is required.

Matrices are written in uppercase bold, \mathbf{A} .

Note: transformations represented by matrices may be written in just italics, as functions often are, i.e., $s(\mathbf{v}) = \mathbf{A}\mathbf{v}$.

History

First Edition: 2023-04-06*

Current Edition: 2023-04-20

Authors

This document was written by R.J. Kit L., a maths student. I am not otherwise affiliated with the university, and cannot help you with related matters.

Please send me a PM on Discord @Desync#6290, a message in the WMX server, or an email to Warwick.Mathematics.Exchange@gmail.com for any corrections. (If this document somehow manages to persist for more than a few years, these contact details might be out of date, depending on the maintainers. Please check the most recently updated version you can find.)

*Storing dates in big-endian format is clearly the superior option, as sorting dates lexicographically will also sort dates chronologically, which is a property that little and middle-endian date formats do not share. See ISO-8601 for more details. This footnote was made by the computer science gang.

If you found this guide helpful and want to support me, you can [buy me a coffee!](#)

(Direct link for if hyperlinks are not supported on your device/reader: ko-fi.com/desync.)

1 Review

The sections on change of basis and eigenvectors from MA106 is repeated here as they are particularly important, but for a full review of past material, see the [MA106 guide](#) or the [reference book](#).

1.1 Mathematical Interpretation

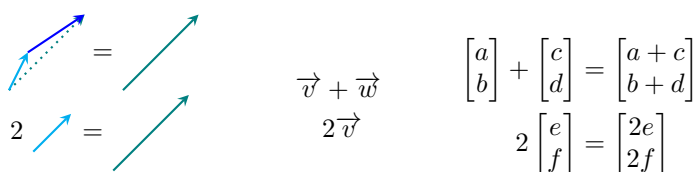
In physics, vectors are often treated as *arrows* pointing in space - some kind of quantity which has a *magnitude* and a *direction*. As long as the length and direction of a vector are the same, it's the same vector, no matter where it is. For example, you might model the velocity of an object as a vector, and consider the velocity as staying the same if the length and direction remain constant.

Sets of vectors that all lie within a plane are two-dimensional, and those in the space we live in are three-dimensional. Picturing an arbitrary n -dimensional vector in this context can be rather tricky, due to the limitations of our reality.

On the other hand, in computer science, vectors are ordered lists of numbers, or *tuples*. For example, you might model the population of two species of animals, say, foxes and rabbits, in a given area with a pair of numbers, the first representing the number of foxes, and the second representing the number of rabbits. Note that order matters; two vectors are not equal if the numbers are swapped around.

In this context, we'd be modelling the populations as a two-dimensional vector. What makes the vector two-dimensional is that the list has two elements within it.

In maths, we are much more general. A vector is anything where we have some kind of notion of adding two objects, our *vectors*, and multiplying those vectors by a number, called a *scalar*. A *vector space* is just a set whose elements are vectors.



$$\vec{v} + \vec{w} = \begin{bmatrix} a \\ b \end{bmatrix} + \begin{bmatrix} c \\ d \end{bmatrix} = \begin{bmatrix} a+c \\ b+d \end{bmatrix}$$

$$2\vec{v} = 2 \begin{bmatrix} e \\ f \end{bmatrix} = \begin{bmatrix} 2e \\ 2f \end{bmatrix}$$

But, in a sense, what makes a vector space, a vector space, are these fundamental operations, independent of how we represent the vectors themselves - it doesn't matter whether you think about vectors fundamentally being arrows, which happen to have a nice numerical representation, or a list of numbers, which happen to have a nice visual representation as arrows. The usefulness of linear algebra is less to do with specific representations of vectors, and more to do with the ability to translate between and equate these different views.

This very general view encompasses both the arrows and ordered lists, and more, but in exchange, is very abstract, and can be possibly more difficult to pick up.

For now, we will first focus on a geometric interpretation of vectors, before moving on to more abstract vector spaces.

When we say a vector, for now, picture an arrow within a coordinate system, with the tail rooted at the origin. Note that this is somewhat distinct from the physics viewpoint discussed above, as vectors in that sense aren't tied to a specific coordinate system, and are free to move about.

This specific view is very helpful as we can then use matrix algebra in our calculations, and changing to a computer science tuple view is just as easy as reading off the coordinates of the head of the vector.

1.2 Basis Vectors, Span & Linear Independence

When we write a vector as a pair of coordinates, say,

$$\mathbf{v} = \begin{bmatrix} 2 \\ 3 \end{bmatrix}$$


you can think of these coordinates as scalars scaling two vectors.

In the Cartesian coordinate system, there are two very special vectors we often use; the vector pointing to the right with length 1, denoted $\hat{\mathbf{i}}$, and the vector pointing up with length 1, denoted $\hat{\mathbf{j}}$.

Thinking of the coordinates as scalars, we scale $\hat{\mathbf{i}}$ by 2, and $\hat{\mathbf{j}}$ by 3, before adding them together to give \mathbf{v} , so \mathbf{v} is the sum of two scaled vectors. Though this is an extremely simple example, this concept of adding two scaled vectors is worth keeping in mind, as it will soon come up, a lot. Any time we scale up vectors and add them together, it's called a *linear combination* of the vectors.

Together, $\hat{\mathbf{i}}$ and $\hat{\mathbf{j}}$ have a special name. They are the *basis vectors* of the Cartesian coordinate system. Specifically, we call them the *canonical* or *standard* basis vectors.


Informally, what it means to form a basis is that, when you use coordinates as scalars, the basis vectors are what the scalars act on. But this brings up a pretty interesting question. What if we picked other basis vectors?



Think about all the arrows you can get by picking two scalars, using them to scale these two arrows, then adding the results together. For these particular two arrows, the answer is that you can reach every possible two-dimensional arrow. You can see this by the fact that the transformed grid covers all of the 2D plane.

A new pair of basis vectors like this also gives us a valid way to translate between pairs of numbers and arrows in the plane. But notice that this translation is different from the canonical basis. $[1,1]$ in this new basis certainly points to a different place than in the canonical basis. We will go into more detail later on, on how coordinates in different bases are related, but for now, just appreciate that any time we describe vectors numerically, it depends on some implicit arbitrary choice of basis vectors.

Now, if we allow the scalars to vary through all possible pairs of values, considering the linear combination given by each pair, we have three possible situations. For most pairs of vectors, we can reach every point in the plane, like in the example above. But if your two vectors line up and are parallel, then the resulting vector is also forced onto the line passing through the origin, parallel to the vectors.



Compared to the previous case, here, the transformed grid is compressed onto a single line.

Additionally, if both vectors are the zero vector, you're just stuck on the origin. The set of all possible vectors you can reach with a linear combination of a set of vectors is the *span* of the vectors. So, we can say that the span of the first pair of vectors above is the entire Cartesian plane (or equivalently, we say that the Cartesian plane is *spanned by* those two vectors, or that those two vectors form a *spanning set* of the Cartesian plane), while the span of the second pair of vectors is just a line, and the span of two zero vectors is just the single point on the origin.

In this second case, we note that one of the vectors is somewhat redundant. We can still access the full line, just using one of the vectors. In this case, we say that the vectors are *linearly dependent* - one of the vectors in the set can be expressed as a linear combination of the others, since it already lies within the span of the others.

Conversely, if each new vector adds a new dimension to the span, we say that the vectors are *linearly independent*.

Symbolically, we say that a set of vectors, $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3 \dots, \mathbf{v}_n$ are linearly independent if the equation,

$$a_1\mathbf{v}_1 + a_2\mathbf{v}_2 + \dots + a_n\mathbf{v}_n = \mathbf{0}$$

holds only if $a_1 = a_2 = \dots = a_n = 0$. In other words, if you can find a way to add up scaled versions of your vectors to get back to the origin, they are not linearly independent.

Now, we can more formally define a basis of a vector space as a set of linearly independent vectors that span the space, and the *dimension* of a vector space is the number of vectors in its basis.

If a linearly independent set of vectors span a space, then every vector in that space can be written as a unique linear combination of those vectors. If this spanning set is not linearly independent, then this linear combination representation of vectors will not be unique (which is why such a set is not usable as a basis - coordinates are not unique). Any two bases of the same vector space contain the same number of vectors. (Take a moment to think about why these properties are true, given the definitions we have just seen.)

1.3 Linear Transformations

1.3.1 Transformations as Matrix-Vector Multiplication

As the name suggests, in linear algebra, we only consider transformations that are linear.

Let V and W be vector spaces over a field, K (we will discuss what a field is later). A function $f : V \rightarrow W$, is linear if for any two vectors, $\mathbf{u}, \mathbf{v} \in V$, and any scalar, $c \in K$,

- $f(\mathbf{u} + \mathbf{v}) = f(\mathbf{u}) + f(\mathbf{v})$ (*additivity* or *operation of vector addition*)
- $f(c\mathbf{u}) = cf(\mathbf{u})$ (*degree 1 homogeneity* or *operation of scalar multiplication*)

In other words, it does not matter whether the linear map is applied before or after the operations of vector addition and scalar multiplication. In particular, linear maps preserves linear combinations. Geometrically, a transformation is linear if the origin is fixed in place, and all lines remain lines under the transformation.

Although this is a rather restrictive condition, there are still a vast range of linear transformations. So, how do we represent these transformations numerically? Given a pair of numbers - a point, a coordinate - how do we find the image of that pair under any given transformation?

Looking back at the definition of a linear transformation, it doesn't matter whether we apply the map before or after the operations of vector addition and scalar multiplication, and, as discussed earlier, every vector can be seen as scaling and adding up the basis vectors - so, if we keep track of where the basis vectors are mapped under the transformation, everything else immediately follows on.

For example, if we know that,

$$\hat{\mathbf{i}} \mapsto \begin{bmatrix} 1 \\ -3 \end{bmatrix} \quad \hat{\mathbf{j}} \mapsto \begin{bmatrix} -2 \\ 4 \end{bmatrix}$$

then we can easily tell where any arbitrary vector,

$$\mathbf{v} = \begin{bmatrix} x \\ y \end{bmatrix}$$

is mapped, by using the linear properties of these maps and breaking it down into its constituent parts, $\mathbf{v} = x\hat{\mathbf{i}} + y\hat{\mathbf{j}}$, so,

$$\begin{bmatrix} x \\ y \end{bmatrix} \mapsto x \begin{bmatrix} 1 \\ -3 \end{bmatrix} + y \begin{bmatrix} -2 \\ 4 \end{bmatrix} = \begin{bmatrix} 1x - 2y \\ -3x + 4y \end{bmatrix}$$

In doing so, we see that every two-dimensional linear map is completely determined by just 4 numbers - the coordinates of the image of $\hat{\mathbf{i}}$ and $\hat{\mathbf{j}}$. Or more generally, the coordinates of the image of the basis vectors of the relevant space. But sticking with $\hat{\mathbf{i}}$ and $\hat{\mathbf{j}}$ for now, we often like to package these coordinates into an array of numbers - a *matrix*.

We do this in such a way that the first column contains the coordinates of where $\hat{\mathbf{i}}$ lands, and the second, the coordinates for $\hat{\mathbf{j}}$.

$$\underbrace{\begin{bmatrix} 1 \\ -3 \end{bmatrix}}_{\hat{\mathbf{i}}} \quad \underbrace{\begin{bmatrix} -2 \\ 4 \end{bmatrix}}_{\hat{\mathbf{j}}}$$

If you have the matrix for some linear transformation, and you want to know the image of any given vector, you take the coordinates of that vector, multiply them by the respective columns of the matrix, and sum the results. In other words, we are adding up the scaled versions of the new basis vectors.

For an arbitrary matrix and vector,

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix}, \begin{bmatrix} x \\ y \end{bmatrix}$$

the image of the vector is given by,

$$\begin{bmatrix} x \\ y \end{bmatrix} \mapsto x \begin{bmatrix} a \\ c \end{bmatrix} + y \begin{bmatrix} b \\ d \end{bmatrix} = \begin{bmatrix} ax + by \\ cx + dy \end{bmatrix}$$

Since the matrix really represents a linear map - a kind of function - let's write it to the left of the vector like we normally do with functions, and give the vector as the function variable.

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} \left(\begin{bmatrix} x \\ y \end{bmatrix} \right) = \begin{bmatrix} ax + by \\ cx + dy \end{bmatrix}$$

But the brackets are somewhat clumsy, so we often drop them from this expression, and read the function as multiplication,


$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} ax + by \\ cx + dy \end{bmatrix}$$

and, we've just discovered matrix-vector multiplication. If you've ever wondered why matrix-vector multiplication is what it is, this is why: it stems from linear transformations being applied to vectors.


1.4 Composition as Matrix-Matrix Multiplication

Now, we often don't study transformations in isolation: what happens if we apply two transformations to a vector, one after another?

For example, consider the transformations given by shearing parallel to the horizontal axis, and rotating by 90° clockwise.



Because of linearity, the overall effect of applying the shear then rotation is another linear transformation, distinct from both the shear and rotation alone. The new transformation is the *composition* of the two original transformations.



(We read right to left for composition. This notation stems from function notation; $(r \circ s)(\hat{\mathbf{i}}) = r(s(\hat{\mathbf{i}}))$, so we apply s first.)

Now, being a linear transformation, this composition also has a matrix representation. Above, we see that,

$$\hat{\mathbf{i}} \mapsto \begin{bmatrix} 0 \\ -1 \end{bmatrix} \quad \hat{\mathbf{j}} \mapsto \begin{bmatrix} 1 \\ -1 \end{bmatrix} \quad (1)$$

so the composition matrix is given by,

$$\begin{bmatrix} 0 & 1 \\ -1 & -1 \end{bmatrix}$$

This matrix gives the effect of shearing, then rotating, in a single transformation - one action, instead of two successive ones.

We can otherwise write the composition out in terms of the original transformations by multiplying a vector on the left by the shear, to give the image under the shear, then multiplying again by the rotation, to apply it after.

$$\underbrace{\begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}}_{\text{Rotation}} \left(\underbrace{\begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}}_{\text{Shear}} \left(\begin{bmatrix} x \\ y \end{bmatrix} \right) \right)$$

Given our definition of matrix-vector multiplication, this is exactly what it means to apply linear transformations as matrices to a vector.

But, given that this pair of transformations has the same overall effect as the composition matrix on any vector, it seems sensible to write

$$\begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix} \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ -1 & -1 \end{bmatrix}$$

More generally, two arbitrary transformation matrices will give another transformation matrix. You probably know how matrix multiplication is defined, but put that knowledge aside for a second, and we will rederive that definition.

$$\underbrace{\begin{bmatrix} a & b \\ c & d \end{bmatrix}}_{M_2} \underbrace{\begin{bmatrix} A & B \\ C & D \end{bmatrix}}_{M_1} = \begin{bmatrix} ? & ? \\ ? & ? \end{bmatrix}$$

To figure out the overall matrix, we need to follow where $\hat{\mathbf{i}}$ goes. By how we construct matrices in the first place, the image of $\hat{\mathbf{i}}$ is just the first column of M_1 . To see where that column is mapped, we then multiply that column by M_2 :

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} A \\ C \end{bmatrix}$$

Using our definition of matrix-vector multiplication we defined earlier, this gives

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} A \\ C \end{bmatrix} = A \begin{bmatrix} a \\ c \end{bmatrix} + C \begin{bmatrix} b \\ d \end{bmatrix} = \begin{bmatrix} Aa + Cb \\ Ac + Cd \end{bmatrix}$$

which is the first column of the composition matrix. Similarly, for $\hat{\mathbf{j}}$, we have,

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} B \\ D \end{bmatrix} = B \begin{bmatrix} a \\ c \end{bmatrix} + D \begin{bmatrix} b \\ d \end{bmatrix} = \begin{bmatrix} Ba + Db \\ Bc + Dd \end{bmatrix}$$

so the composition matrix is,

$$\begin{bmatrix} Aa + Cb & Ba + Db \\ Ac + Cd & Bc + Dd \end{bmatrix}$$

This is where the arbitrary-feeling “rows into columns” definition matrix multiplication you learned at A-level actually comes from. It’s just how the numbers work out when we compose linear transformations together.

Furthermore, seeing matrix multiplication as composition of transformations makes the various properties of matrix multiplication much easier to understand.


For example, rotating then shearing, and, shearing then rotating, clearly give different results, so matrix multiplication is not commutative. This is a trivial property you can verify in your head, without having to compute anything at all.

Similarly, matrix multiplication is clearly associative: applying transformation A , then $(B$ then $C)$ is clearly the same thing as applying transformation $(A$ then $B)$, then C . There's nothing to prove here; it's the same three transformations being applied in the same order both ways.

Trying to prove these properties symbolically is a nightmare, but, as transformation compositions, they're trivial. Not only are these valid proofs, they're good intuitive explanations as to why these properties should be true.

1.5 Change of Basis


Assigning numbers to vectors (that are interpreted as arrows rooted at the origin) depends on some choice of basis vectors to provide a meaningful translation between geometry and algebra.



In our standard system, we would say that this green vector has coordinates,

$$\begin{bmatrix} 3 \\ 2 \end{bmatrix}$$


because going from its tail to its tip requires moving 3 units to the right, and 2 units up. We think of these coordinates as scalars - something that scales up a vector. In this case, we implicitly take the first coordinate to scale \hat{i} , and the second to scale \hat{j} , before adding up the result, with all the information about distance and direction tied up in our choice of basis vectors.



We call these ways to translate between these arrows and sets of numbers a *coordinate system*. The choice of \hat{i} being the target of the first scalar, and \hat{j} being the target of the second scalar gives us the standard Cartesian coordinate system.

But of course, other basis vectors are available.


Say we have a friend, Alice, who uses a different set of basis vectors: \mathbf{v}_1 , which points to the bottom-right and is acted upon by the first scalar of a coordinate, and \mathbf{v}_2 , that points to the top right and is acted upon by the second scalar.



We can draw the same green vector again - the one we would describe as $[3, 2]$ - on to her grid. Alice would then describe this green vector as,

$$\begin{bmatrix} 1 \\ \frac{3}{2} \end{bmatrix}$$

What this means is that, to get to the tip of that vector using her basis vectors, is to scale up \mathbf{v}_1 by 1, \mathbf{v}_2 by $\frac{3}{2}$, then add up the results



Whenever Alice uses coordinates to describe a vector, she thinks of the first coordinate scaling \mathbf{v}_1 , and the second, scaling \mathbf{v}_2 , just like how we scale $\hat{\mathbf{i}}$ and $\hat{\mathbf{j}}$, respectively.

We note that, although the two coordinates look different, they actually represent the same vector, just in two different coordinate systems. We're both describing the same things, but in a different language.

$$\begin{bmatrix} 3 \\ 2 \end{bmatrix} = \begin{bmatrix} 1 \\ \frac{3}{2} \end{bmatrix}$$


I've been showing the choices of bases using colour (and will continue doing so), but it is helpful to have notation for this as well. To do this, we give a label to each choice of basis, and subscript our vectors with that label. Often, this is done with set brackets (for example, $\{\mathbf{e}_i\}$) to indicate that the basis is a set of vectors, but here, for clarity, I will label the Cartesian coordinate system as E , and Alice's coordinate system as A .

$$\begin{bmatrix} 3 \\ 2 \end{bmatrix}_E = \begin{bmatrix} 1 \\ \frac{3}{2} \end{bmatrix}_A$$

But how did we find that second set of coordinates? More generally, how do we find the coordinates of some vector in some given different coordinate system? Well, we should first look at the basis vectors of the coordinate systems in question.

We can describe the basis vectors of the target coordinate system in terms of our standard one. In E , Alice's basis vectors are,

$$\mathbf{v}_1 = \begin{bmatrix} \frac{3}{2} \\ -1 \end{bmatrix}_E, \text{ and } \mathbf{v}_2 = \begin{bmatrix} 1 \\ 2 \end{bmatrix}_E$$



But it is important to note that, in her system, these vectors are

$$\mathbf{v}_1 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}_A, \text{ and } \mathbf{v}_2 = \begin{bmatrix} 0 \\ 1 \end{bmatrix}_A$$

since this is exactly what it means to even be a basis vector. They are what *define* the meaning of $[1,0]$ and $[0,1]$ in her system.

Both systems look at the same vector in space, but assign it different coordinates. The point is, the grids both of us use are just artificial constructs. Space does not intrinsically have a grid.

So, if we want to translate from our standard basis to Alice's, we want to find some kind of function that maps $[\frac{3}{2}, -1]$ in our system to $[1,0]$ in Alice's system, and similarly, $[1,2]$ to $[0,1]$. We also note that $[0,0]$ is exactly the same in both coordinate systems: we both agree on where the origin is, since scaling any vector by 0 should always give the same result, regardless of coordinate system.

Now, on the surface, this seems rather difficult. However, it might be easier to find the translation from the Alice's basis to our standard basis, where we're looking to map $[1,0]$ to $[\frac{3}{2}, -1]$, and $[0,1]$ to $[1,2]$, and you might already see where we're going with this.

If we were given some vector, say $[1, -2]$, given in Alice's coordinates, A , how would we go about translating this into our standard coordinates, E ? Well, the first coordinate scales Alice's first basis vector, and similarly to the second, and we know how to express those basis vectors in our coordinate system, so we have,

$$1 \begin{bmatrix} \frac{3}{2} \\ -1 \end{bmatrix} - 2 \begin{bmatrix} 1 \\ 2 \end{bmatrix} = \begin{bmatrix} -2 \\ 4 \end{bmatrix}$$

So $[1, -2]_A$ is expressed as $[-2,4]_E$ in our standard coordinate system. But...

Doesn't this look familiar?

Recalling from a long way back (§ 1.3.1), we've already done this exact same thing before! It's matrix-vector multiplication, with the matrix containing Alice's basis vectors expressed in our coordinate system.

$$1 \begin{bmatrix} \frac{3}{2} \\ -1 \end{bmatrix} - 2 \begin{bmatrix} 1 \\ 2 \end{bmatrix} = \begin{bmatrix} \frac{3}{2} & 1 \\ -1 & 2 \end{bmatrix} \begin{bmatrix} 1 \\ -2 \end{bmatrix}$$

This mapping between bases is actually a linear transformation in and of itself, and we often label its associated matrix as \mathbf{P} .

$$\mathbf{P} = \begin{bmatrix} \frac{3}{2} & 1 \\ -1 & 2 \end{bmatrix}$$

In general, this matrix is given by the basis vectors of the coordinate system being converted *from*, expressed in the coordinate system being converted *to* (\mathbf{P} here converts from Alice's coordinates to ours, so we use Alice's basis vectors written in our language.)

This matrix, \mathbf{P} , is called a *change of basis* matrix. In this case, from A to E . To convert any vector given in F to its representation in E , we left multiply by this matrix.

Since we wanted to find the coordinates for the green vector in A , given that we know its coordinates in E , we simply take the inverse, \mathbf{P}^{-1} , and left multiply by that instead.

$$\mathbf{P}^{-1} = \begin{bmatrix} \frac{1}{2} & -\frac{1}{4} \\ \frac{1}{4} & \frac{3}{8} \end{bmatrix}$$

And you can verify yourself that this matrix, multiplied with $[3,2]_E$, gives $[\frac{3}{2},1]_A$.

From our geometric interpretation of matrix-vector multiplication from before, this is actually a very reasonable thing to do. The matrix containing the coordinates of the new basis vectors moves our basis vectors, $\hat{\mathbf{i}}$ and $\hat{\mathbf{j}}$ - the things we think of as $[1,0]$ and $[0,1]$, - over to Alice's basis vectors - the things she thinks of as $[1,0]$ and $[0,1]$.

For example, if Alice was talking about a vector, say, $[1,2]$, then multiplying $[1,2]$ by \mathbf{P} transforms our basis vectors over to Alice's, where the process of scaling then adding basis vectors by the coordinates $[1,2]$ works in our favour, as we're effectively now working with Alice's basis vectors.

Geometrically, this matrix transforms our grid into Alice's, but numerically, it translates a vector in Alice's system into our system.

$$\mathbf{P} = \underbrace{\begin{bmatrix} a & b \\ c & d \end{bmatrix}}_{\substack{\text{Alice's basis vectors} \\ \text{in our coordinates}}} \quad \mathbf{P} \underbrace{\begin{bmatrix} x_0 \\ y_0 \end{bmatrix}}_{\substack{\text{Vector in} \\ \text{Alice's coordinates}}} = \underbrace{\begin{bmatrix} x_1 \\ y_1 \end{bmatrix}}_{\substack{\text{The same vector,} \\ \text{in our coordinates}}}$$

1.6 Transformations in Different Bases

Now we can translate vectors between bases, how about transformations? If we have the 90° clockwise rotation matrix,

$$\mathbf{U} = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$$

how would Alice represent this same transformation in her own coordinate system? To be clear, we're trying to write down a matrix that takes Alice's grid, and rotates it 90° clockwise.

The columns of the matrix encode information about where our basis vectors, $\hat{\mathbf{i}}$ and $\hat{\mathbf{j}}$ go, so just translating the columns into Alice's coordinates is not enough. That would just give a matrix that tells her where our basis vectors would land, written in her coordinate system.

She wants a matrix that gives where *her* basis vectors land, and it needs to describe those landing spots in her coordinate system as well.

Let's first consider what the rotation matrix does to a single specific vector, given in her coordinate system, say,

$$\begin{bmatrix} 3 \\ 4 \end{bmatrix}$$

Since we don't know the rotation matrix in her system, let's first convert this vector into our coordinate system. We do this by using the change of basis matrix - the matrix containing her basis vectors in our coordinate system as columns.

$$\underbrace{\begin{bmatrix} 1 & \frac{3}{2} \\ 2 & -1 \end{bmatrix}}_{\substack{\text{The same vector,} \\ \text{but in our language}}} \begin{bmatrix} 3 \\ 4 \end{bmatrix}$$

Change of basis matrix, \mathbf{P}

And now we have the vector in a form we can work with. I've left the multiplication unexpanded, but keep in mind that the whole right hand side represents a vector - the exact same vector as before, just described in our language.

Since we know the rotation matrix in our system, we can just multiply this whole thing by it:

$$\underbrace{\begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}}_{\text{Transformation matrix in our language}} \underbrace{\begin{bmatrix} 1 & \frac{3}{2} \\ 2 & -1 \end{bmatrix} \begin{bmatrix} 3 \\ 4 \end{bmatrix}}_{\text{Transformed vector in our language}}$$

This tells us where the vector should go, but it's still in our language, so we convert it back into Alice's basis with the inverse change of basis matrix:

$$\underbrace{\begin{bmatrix} 1 & \frac{3}{2} \\ 2 & -1 \end{bmatrix}^{-1}}_{\text{Inverse change of basis matrix, } \mathbf{P}^{-1}} \underbrace{\begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix} \begin{bmatrix} 1 & \frac{3}{2} \\ 2 & -1 \end{bmatrix} \begin{bmatrix} 3 \\ 4 \end{bmatrix}}_{\text{Transformed vector in Alice's language}}$$

And we've just figured out where some specific vector, given in Alice's language, should go, under a 90° clockwise rotation. But, since the choice of vector was arbitrary, we've found the transformation we wanted!

We apply the change of basis matrix to get the vector into a workable form, then the transformation (which we know in our language), then the inverse change of basis matrix to translate back.

$$\underbrace{\begin{bmatrix} 1 & \frac{3}{2} \\ 2 & -1 \end{bmatrix}^{-1} \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix} \begin{bmatrix} 1 & \frac{3}{2} \\ 2 & -1 \end{bmatrix}}_{\text{Transformation matrix in Alice's language}} \mathbf{v}$$


This composition of three matrices, together, gives the rotation matrix in Alice's coordinate system. It takes in a vector, in her language, and returns the transformed version of that vector, in her language.

We can represent all of this with the (mis)use of commutative diagrams:

$$\begin{array}{ccc} V_E & \xrightarrow{\mathbf{R}} & T(V_E) \\ \uparrow \mathbf{P} & \xlongequal{T} & \uparrow \mathbf{P} \\ V_A & \xrightarrow{\mathbf{S}} & T(V_A) \end{array}$$

where the transformation, $T : V \rightarrow V$ is acting on V , with the choice of basis indicated using subscripts. Note that the change of basis matrix works the same before and after the transformation - it still translates between Alice's system and ours. The transformation of space, T (notice that this is not written in bold, as it is a transformation not tied to a specific matrix), can be represented as the matrices, \mathbf{R} and \mathbf{S} , specific to the bases E and A .

Note, since we're only dealing with one transformation that is pretty obviously an endomorphism, and only one vector space is being considered, in an exam, I would just draw the diagram above as:




to save time. Some of the diagrams in this section will contain additional detail to aid my explanations, but by no means do you have to include every little extraneous detail when using these diagrams yourself.

Some people define the change of basis matrix to be the opposite way as is defined here, with \mathbf{P} being the change of basis from E to A (so what we would call \mathbf{P}^{-1}), but as long as you are consistent with your arrows, the diagram makes everything clear.

On the top diagram, we want S , which directly transforms V_A to $T(V_A)$. An alternative route there, is to take P , then R , then to go along the second P arrow, but against the direction it is pointing, which indicates we should take an inverse of the matrix representing P . So, in terms of matrices, $\mathbf{S} = \mathbf{P}^{-1}\mathbf{R}\mathbf{P}$ (reading right to left, as per function notation), matching the result from before.


Now, let's say we have another friend, Bob, who uses yet another coordinate system, distinct from both Alice's and ours. How might Alice and Bob communicate? How do we find a change of basis from Bob to Alice, and vice versa?

We can use our standard basis as an intermediary:



Where R and S are the change of basis transformations from Alice's and Bob's systems to ours, as found earlier. We want P here, so we travel along R , then backwards along S , so $P = S^{-1}R$.

How would Bob give a transformation to Alice?



Following the arrows, we have,

$$U = S^{-1}RTR^{-1}S$$


or,

$$U = P^{-1}TP$$

In general, when we give a transformation between vector spaces, $T : V \rightarrow W$, we have to be careful when turning this transformation into a matrix. V and W , being abstract spaces, don't intrinsically have grids mapping them out - we have to assign basis vectors to each.

For example, let $T : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be a 90° clockwise rotation. The matrix for $\mathbb{R}_E^2 \rightarrow \mathbb{R}_E^2$ is just the rotation matrix we're familiar with, and we found the matrix for $\mathbb{R}_A^2 \rightarrow \mathbb{R}_A^2$ earlier. But a matrix giving the rotation from $\mathbb{R}_B^2 \rightarrow \mathbb{R}_A^2$ is an equally valid representation of that same transformation. In the diagram above, this matrix could be given as $TR^{-1}S$, or $R^{-1}SU$.


Now, for some more terminology. Going back to the first diagram,



Because \mathbf{R} and \mathbf{S} represent the same transformation, T , within the same space, V , just with respect to different bases, we call them *similar* matrices. In general, two $n \times n$ matrices, \mathbf{A} and \mathbf{B} are similar if you can write $\mathbf{B} = \mathbf{P}^{-1}\mathbf{A}\mathbf{P}$ for some (usually change of basis) matrix \mathbf{P} . Two similar matrices must be square.

More generally, for possibly rectangular $m \times n$ matrices, we have an analogous concept of *equivalence* (this equivalence is a type of equivalence relation, if you have done binary relations). Two rectangular matrices, \mathbf{A} and \mathbf{B} are equivalent if you can write $\mathbf{B} = \mathbf{Q}\mathbf{A}\mathbf{P}$ for two invertible matrices \mathbf{P} and \mathbf{Q} (the change of basis matrices for each of the pairs of coordinate systems for each space).

On a diagram, this would be,



Here, V and W are vector spaces of different dimension, with subscripts indicating choice of basis. There are 4 bases in play here, as Alice and Bob each choose their own bases for both V and W . Note that, unlike the previous diagram, $\mathbf{P} \neq \mathbf{Q}$, since they are change of basis matrices in different spaces. Here, R and S both represent the same transformation of space, then they are equivalent, as you could write $\mathbf{R} = \mathbf{Q}^{-1}\mathbf{S}\mathbf{P}$. As a side note, to find \mathbf{P} , you would do the same procedure as a few diagrams ago, and use the standard basis as an intermediary, but the diagram was already getting cluttered enough, so this is left as an exercise for the reader.


In general, similar matrices are equivalent, but equivalent matrices are not necessarily similar.

1.7 Eigenvectors

Consider the linear transformation given by the matrix,

$$\mathbf{A} = \begin{bmatrix} 3 & 1 \\ 0 & 2 \end{bmatrix}$$

and how it acts on some arbitrary vector. In particular, think about the span of that vector.




For most vectors in the plane, they get knocked off of their span during the transformation. But some special vectors do remain on their own span, meaning the transformation has no rotational effect on that vector, only scaling it by some amount.

As you might have guessed, such a vector is called an *eigenvector* of the transformation, and the amount by which it is scaled is its associated *eigenvalue*. (The prefix, *-eigen*, from German, means *own*, so an eigen-vector is an “own-vector”, which makes sense, given that its image is just itself, up to scaling.)

For the transformation above, $\hat{\mathbf{i}}$ is one such vector. The span of $\hat{\mathbf{i}}$ is just the horizontal axis, and the image of $\hat{\mathbf{i}}$ clearly remains on that axis after the transformation. From the matrix, we can see that $\hat{\mathbf{i}}$ lands on $[3,0]$, so it is scaled by a factor of 3. We say that $[1,0]$ is an eigenvector of \mathbf{A} , with an eigenvalue of 3.

Furthermore, due to linearity, *any* vector on the horizontal axis is also similarly scaled by a factor of 3, also remaining on their own spans.

But there are more, slightly less obvious, eigenvectors to this particular transformation:




This vector, $[1, -1]$ lands on $[2, -2]$, being scaled by a factor of 2, so $[1, -1]$ is also an eigenvector of \mathbf{A} , with eigenvalue 2. And again, due to linearity, any vector on that line will also be an eigenvector with eigenvalue 2.

For this transformation, those are all the eigenvectors there are. Every other vector in the plane will get moved off of their spans under this transformation.

But you can also have more or fewer eigenvectors - any rotation matrix,

$$\mathbf{R} = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$$

will move every vector off of its span.




So this transformation has zero eigenvectors.

This shear,

$$\mathbf{S} = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}$$


on the other hand, has every vector on the horizontal axis as an eigenvector, with eigenvalue 1, as they are unchanged by the transformation.



But every other vector is moved off of its span, so this transformation just has a single line of eigenvectors, the horizontal axis, with eigenvalue 1.

Eigenvalues don't have to be unique either - you can have multiple lines of eigenvectors with the same eigenvalue, as given by,

$$\mathbf{T} = \begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix}$$



This scaling matrix just stretches every vector in the plane by a factor of 2, so *every* vector is an eigenvector of this transformation, all with the same eigenvalue of 2.

By definition, the effect of a transformation, \mathbf{A} , on an eigenvector, \mathbf{v} , is just to scale it by some amount, λ , the eigenvalue. We can write this definition symbolically, as,

$$\mathbf{A}\mathbf{v} = \lambda\mathbf{v}$$

The left side is matrix-vector multiplication, while the right is scalar multiplication, so we tend to do some rearranging of this expression, by writing the right side as a matrix-vector product.

We want to scale \mathbf{v} by a scalar, λ . The columns of the desired matrix need to scale each basis vector by λ , so this matrix will have λ across the diagonal, and zeros everywhere else.

$$\begin{bmatrix} \lambda & 0 & 0 & \cdots & 0 \\ 0 & \lambda & 0 & \cdots & 0 \\ 0 & 0 & \lambda & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & \lambda \end{bmatrix}$$

Factoring out the λ , this is just the identity matrix.

$$\mathbf{A}\mathbf{v} = (\lambda\mathbf{I})\mathbf{v}$$

And now both sides are a matrix-vector product. We can then subtract the right side, and factor out the \mathbf{v} ,

$$\begin{aligned} \mathbf{A}\mathbf{v} - (\lambda\mathbf{I})\mathbf{v} &= \mathbf{0} \\ (\mathbf{A} - \lambda\mathbf{I})\mathbf{v} &= \mathbf{0} \end{aligned}$$

The expression inside the bracket is just a matrix - the original transformation matrix, \mathbf{A} , but with a λ being subtracted from the diagonal, and would look something like,

$$\begin{bmatrix} 3 - \lambda & 1 & 4 \\ 1 & 5 - \lambda & 9 \\ 2 & 6 & 5 - \lambda \end{bmatrix}$$

We're looking for a vector, \mathbf{v} , such that this new matrix maps \mathbf{v} to the zero vector.

If $\mathbf{v} = \mathbf{0}$, then this is trivially true and isn't particularly helpful, so we're looking for non-zero solutions for \mathbf{v} - a non-zero eigenvector.

That is, we're looking for a non-zero vector that lies in the null space of $(\mathbf{A} - \lambda\mathbf{I})$. If the null space of $(\mathbf{A} - \lambda\mathbf{I})$ is non-empty, i.e., there exists a non-zero eigenvector, it follows that $(\mathbf{A} - \lambda\mathbf{I})$ cannot be full rank, and therefore has a zero determinant (if this doesn't make immediate sense, reread the definitions of null space and rank, and take a few moments to consider what it means for a transformation to have a non-empty null space).

In other words, the only way for a non-zero vector to be mapped to the origin, is if the transformation collapses space down into a lower dimension, corresponding to a zero determinant.

So, we're trying to solve,

$$\det(\mathbf{A} - \lambda\mathbf{I}) = 0$$

for λ . This is called the *characteristic equation* of the matrix. The left side by itself is called the *characteristic polynomial*.

For example, earlier, we had,

$$\begin{aligned} \mathbf{A} &= \begin{bmatrix} 3 & 1 \\ 0 & 2 \end{bmatrix} \\ \mathbf{A} - \lambda\mathbf{I} &= \begin{bmatrix} 3 - \lambda & 1 \\ 0 & 2 - \lambda \end{bmatrix} \\ \det(\mathbf{A} - \lambda\mathbf{I}) &= (3 - \lambda)(2 - \lambda) - (1 \cdot 0) \\ \det(\mathbf{A} - \lambda\mathbf{I}) &= 0 \\ 0 &= (3 - \lambda)(2 - \lambda) \\ \lambda &= 3, 2 \end{aligned}$$

matching the results from before. Then, to find the actual eigenvectors, multiply the modified matrix by an arbitrary vector, and set it equal to 0.

For $\lambda = 2$, we have

$$\begin{aligned} \begin{bmatrix} 3 - 2 & 1 \\ 0 & 2 - 2 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} &= \begin{bmatrix} 0 \\ 0 \end{bmatrix} \\ \begin{bmatrix} 1 & 1 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} &= \begin{bmatrix} 0 \\ 0 \end{bmatrix} \\ \begin{bmatrix} x + y \\ 0 \end{bmatrix} &= \begin{bmatrix} 0 \\ 0 \end{bmatrix} \\ x + y &= 0 \\ y &= -x \end{aligned}$$

So any vector of the form,

$$\begin{bmatrix} t \\ -t \end{bmatrix}$$

is an eigenvector with eigenvalue 2. We generally just pick one single eigenvector as a representative for this entire line, so letting $t = 1$, we have $[1, -1]$, as before.

For a transformation which has multiple eigenvectors with the same eigenvalue, you'll find that the simultaneous equations in the final step will have multiple solutions, corresponding to the multiple eigenvectors.

Doing the same process for the rotation matrix,

$$\begin{aligned}\mathbf{R} - \lambda\mathbf{I} &= \begin{bmatrix} 0 - \lambda & 1 \\ -1 & 0 - \lambda \end{bmatrix} \\ \det(\mathbf{R} - \lambda\mathbf{I}) &= (-\lambda)(-\lambda) - (1 \cdot (-1)) \\ \det(\mathbf{R} - \lambda\mathbf{I}) &= 0 \\ \lambda^2 + 1 &= 0 \\ \lambda &= \pm i\end{aligned}$$

we find that there are no real eigenvalues for this transformation. The eigenvalues of $\pm i$ correspond to the fact that multiplying by i represents a 90° rotation in the complex plane, and the magnitude of the eigenvalues being 1 corresponds to the fact that vectors aren't scaled under this transformation. In general, imaginary components of eigenvalues correspond to some kind of rotation. We can still solve for eigenvectors, but they will have complex components.


As one very basic application of eigenvectors, if you can find an eigenvector of a 3D rotation, you've found the axis of rotation - and it's much easier to think of rotations in 3D as an angle around an axis, rather than the entire 3×3 rotation matrix.

This is a common theme throughout linear algebra - with any linear transformation given as a matrix, we can interpret what it is doing by looking at its columns and seeing where the basis vectors are mapped. But this puts a lot of emphasis on coordinate systems - another way, less dependent on coordinate systems, is to look at the eigenvectors and eigenvalues.

Two similar matrices - two matrices representing the same linear transformations, but in different coordinate systems - will have the same characteristic equation, and the same eigenvalues. Changing the coordinate system doesn't change the eigenvalues of a transformation - regardless of how you label space, eigenvectors are scaled the same way.

For another, much more general application, consider what happens if our basis vectors both happen to be eigenvectors. Let's start in the canonical coordinate system, and say we have a linear transformation such that,

$$\hat{\mathbf{i}} \mapsto \begin{bmatrix} 2 \\ 0 \end{bmatrix} = 2\hat{\mathbf{i}} \quad \hat{\mathbf{j}} \mapsto \begin{bmatrix} 0 \\ -1 \end{bmatrix} = -1\hat{\mathbf{j}}$$



So the matrix associated with that transformation would be,

$$\begin{bmatrix} 2 & 0 \\ 0 & -1 \end{bmatrix}$$

Notice how the eigenvalues of the basis vectors lie along the diagonal of the matrix, and every other entry is zero. Any matrix with this property is called a *diagonal matrix*, and we've met one before - the identity matrix is a diagonal matrix.

The way to interpret a diagonal matrix, is that all the basis vectors are eigenvectors, with the eigenvalues written along the diagonals.

There are many reasons why diagonal matrices are much nicer to work with. One application is in taking powers of matrices, or equivalently, applying a transformation to a vector many times. Since a diagonal matrix only scales each basis vector by some eigenvalue, applying that matrix to a vector n times, just means you scale each basis vector by the eigenvalue to the power of n .

$$\begin{aligned} \begin{bmatrix} 2 & 0 \\ 0 & 3 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} &= \begin{bmatrix} 2x \\ 3y \end{bmatrix} \\ \begin{bmatrix} 2 & 0 \\ 0 & 3 \end{bmatrix} \begin{bmatrix} 2 & 0 \\ 0 & 3 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} &= \begin{bmatrix} 2^2 x \\ 3^2 y \end{bmatrix} \\ \underbrace{\begin{bmatrix} 2 & 0 \\ 0 & 3 \end{bmatrix} \cdots \begin{bmatrix} 2 & 0 \\ 0 & 3 \end{bmatrix} \begin{bmatrix} 2 & 0 \\ 0 & 3 \end{bmatrix}}_{100} \begin{bmatrix} x \\ y \end{bmatrix} &= \begin{bmatrix} 2^{100} x \\ 3^{100} y \end{bmatrix} \\ &= \begin{bmatrix} 2^{100} & 0 \\ 0 & 3^{100} \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} \end{aligned}$$

Just looking at the transformation overall, we can write an exceedingly simple formula for the n th power of a diagonal matrix:

$$\begin{bmatrix} a & 0 \\ 0 & b \end{bmatrix}^n = \begin{bmatrix} a^n & 0 \\ 0 & b^n \end{bmatrix}$$

In contrast, try calculating the 100th power of a non-diagonal matrix. There is no simple pattern to find.

For a few more nice properties of diagonal matrices, the determinant of a diagonal matrix is just the product of the diagonal. This is because each entry on the diagonal tells us how much the basis vector is scaled in that direction, so the product of all of these entries gives us how much measure is scaled overall.

Of course, all of this is only useful when the matrix we're working with is diagonal - when our basis vectors just happen to both be eigenvectors.

However, if your transformation has a lot of eigenvectors, enough so we can choose a set that spans the space the transformation is acting on, then we could use a change of basis matrix to change those eigenvectors to be our basis.

For the matrix earlier,

$$\mathbf{A} = \begin{bmatrix} 3 & 1 \\ 0 & 2 \end{bmatrix}$$

we found two eigenvectors,

$$\mathbf{v}_1 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \text{ and } \mathbf{v}_2 = \begin{bmatrix} 1 \\ -1 \end{bmatrix}$$

with eigenvalues $\lambda_1 = 3$, and $\lambda_2 = 2$, respectively.

We use the eigenvectors as the columns of a change of basis matrix, and change the transformation matrix into our new basis, as before.

$$\begin{bmatrix} 1 & 1 \\ 0 & -1 \end{bmatrix}^{-1} \begin{bmatrix} 3 & 1 \\ 0 & 2 \end{bmatrix} \begin{bmatrix} 1 & 1 \\ 0 & -1 \end{bmatrix}$$

Because we've chosen the basis vectors to be eigenvectors, we know that resulting matrix will be a diagonal matrix, with the corresponding eigenvalues along the diagonal, without even doing any calculations.

$$\begin{bmatrix} 3 & 0 \\ 0 & 2 \end{bmatrix}$$

This is because we're now working in a basis, where the effect of this specific transformation on these basis vectors is just to scale them by these eigenvalues.

A basis where every basis vector is an eigenvector is called an *eigenbasis*, and this process of changing a matrix to an eigenbasis is called *diagonalisation*.

If we wanted to calculate the 100th power of \mathbf{A} , we could change to an eigenbasis, calculate the power there, using our simple formula, then change back.

But not every transformation admits an eigenbasis. The shear we saw earlier, for example, only has a single line of eigenvectors, which isn't enough to span all of 2D space.

2 Jordan Canonical Form

In this section, we will take V to be an n -dimensional vector space over a field K . We will take $T : V \rightarrow V$ to be a linear map from V to V (an *endomorphism*) and \mathbf{A} will be the matrix representing T with respect to a fixed ordered basis $E = (\mathbf{e}_i)_{i=1}^n$.

Our goal is to find a new basis $e = (\mathbf{e}_i)_{i=1}^n$ such that the matrix of T with respect to this new basis is as simple as possible (or equivalently, a change of basis matrix \mathbf{P} such that $\mathbf{P}^{-1}\mathbf{A}\mathbf{P}$ is as simple as possible).

One particularly simple form of a matrix is a diagonal matrix, but as mentioned above, not every transformation admits a diagonal matrix representation.

However, if K is \mathbb{C} (or is a field extension of \mathbb{C}), then every matrix \mathbf{A} is similar to a form that is almost as good as diagonal: the *Jordan canonical form* or *Jordan normal form*.

2.1 Generalised Eigenspaces

Theorem 2.1. *Let $T : V \rightarrow V$ be a linear map. Then, the matrix of T is diagonalisable if and only if V has an eigenbasis.*

Theorem 2.2. *Let $(\lambda_i)_{i=1}^r$ be distinct eigenvalues of $T : V \rightarrow V$, and let $(\mathbf{v}_i)_{i=1}^r$ be corresponding eigenvectors – that is, $T(\mathbf{v}_i) = \lambda_i\mathbf{v}_i$ for all $1 \leq i \leq r$. Then, $(\mathbf{v}_i)_{i=1}^r$ are linearly independent.*

Theorem 2.3. *Let $\mathbf{A} \in K^{n \times n}$ be a $n \times n$ matrix over K . Then, there is some non-zero polynomial $p \in K[x]$ of degree at most n^2 such that $p(\mathbf{A}) = \mathbf{0}_n$.*

A polynomial is *monic* if the coefficient of the highest degree term is 1.

Theorem 2.4. *Let $\mathbf{A} \in K^{n \times n}$ represent the linear transformation $T : V \rightarrow V$. Then,*

- *There is a unique monic non-zero polynomial p with minimal degree and coefficients in K such that $p(\mathbf{A}) = \mathbf{0}_n$.*
- *If q is any polynomial with $q(\mathbf{A}) = \mathbf{0}_n$, then p divides q .*

This unique polynomial is called the *minimal polynomial* of \mathbf{A} , and is denoted $\mu_{\mathbf{A}}$.

Theorem 2.5. *Similar matrices have the same minimal polynomial.*

Theorem 2.6. Let \mathbf{D} be a diagonal matrix with distinct diagonal entries $(\delta_i)_{i=1}^r$. Then,

$$\begin{aligned}\mu_{\mathbf{D}}(x) &= \prod_{i=1}^r (x - \delta_i) \\ &= (x - \delta_1)(x - \delta_2) \cdots (x - \delta_r)\end{aligned}$$

Corollary 2.6.1. If \mathbf{A} is diagonalisable, then $\mu_{\mathbf{A}}$ is a product of linear factors.

2.2 Cayley-Hamilton Theorem

Recall the *characteristic equation* of a matrix \mathbf{A} is defined as follows:

$$\det(\mathbf{A} - \lambda \mathbf{I}) = 0$$

The left side by itself is called the *characteristic polynomial* of \mathbf{A} , denoted $c_{\mathbf{A}}$.

Theorem 2.7 (Cayley-Hamilton). Let $A \in K^{n \times n}$, and let $c_{\mathbf{A}}$ be the characteristic polynomial of A . Then, $c_{\mathbf{A}}(\mathbf{A}) = \mathbf{0}_n$.

Corollary 2.7.1. For any $\mathbf{A} \in K^{n \times n}$, $\mu_{\mathbf{A}}$ divides $c_{\mathbf{A}}$, and in particular, $\deg(\mu_{\mathbf{A}}) \leq n$.

2.3 Calculating Minimal Polynomials

Lemma 2.8. Let λ be any eigenvalue of \mathbf{A} . Then, $\mu_{\mathbf{A}}(\lambda) = 0$.

Using this lemma with the Cayley-Hamilton theorem lets us reduce the possibilities for the minimal polynomial.

Algorithm 1 Top Down Algorithm

- 1: Calculate the characteristic polynomial, $c_{\mathbf{A}}$.
 - 2: Factorise the characteristic polynomial by inspection, or using the factor and remainder theorem with polynomial division.
 - 3: Evaluate each possible combination of factors that include all eigenvalues as roots in order of increasing degree. The first to return $\mathbf{0}_n$ is the minimal polynomial.
-

As an example for the last step, if $c_{\mathbf{A}}(x) = (x - 1)(x - 2)^2(x - 3)^3$, then,

$$\mu_{\mathbf{A}}(x) \in \begin{cases} (x - 1)(x - 2)(x - 3) \\ (x - 1)(x - 2)(x - 3)^2 \\ (x - 1)(x - 2)^2(x - 3) \\ (x - 1)(x - 2)^2(x - 3)^2 \\ (x - 1)(x - 2)(x - 3)^3 \\ (x - 1)(x - 2)^2(x - 3)^3 \end{cases}$$

Then, evaluate the polynomials in this list at \mathbf{A} from top to bottom (the list is sorted in degree order), and the first one to return the zero matrix is the minimal polynomial.

Algorithm 2 Bottom Up Algorithm

- 1: Pick some simple non-zero vector, \mathbf{v} (the standard basis vector \mathbf{e}_1 is often a good choice).
- 2: Apply \mathbf{A} to \mathbf{v} repeatedly to form a chain of vectors

$$\mathbf{v}_0 \xrightarrow{\mathbf{A}} \mathbf{v}_1 \xrightarrow{\mathbf{A}} \mathbf{v}_2 \xrightarrow{\mathbf{A}} \cdots$$

- 3: At some point, these image vectors will become linearly dependent, say, after d applications of \mathbf{A} , so there exists coefficients (α_i) such that

$$\sum_{i=0}^d \alpha_i \mathbf{v}_i = \mathbf{0}$$

with $\alpha_d = 1$.

- 4: Then, the monic polynomial

$$\sum_{i=0}^d \alpha_i x^i$$

divides the minimal polynomial.

- 5: Repeat this process with different starting vectors that do not lie in the image of previous generated chains, until all the generated chains span V . The minimal polynomial is then the least common multiple of these polynomials.

2.4 Jordan Chains

Recall that a non-zero vector \mathbf{v} that satisfies $(\mathbf{A} - \lambda \mathbf{I}_n)\mathbf{v} = \mathbf{0}$ is an eigenvector of \mathbf{A} with eigenvalue λ . We weaken this notion to classify a more general type of vector.

A non-zero vector \mathbf{v} that satisfies $(\mathbf{A} - \lambda \mathbf{I}_n)^i \mathbf{v} = \mathbf{0}$ for some $i > 0$ is a *generalised eigenvector* of \mathbf{A} with eigenvalue λ , and, for a fixed $i > 0$ and fixed λ , the collection of these generalised eigenvectors,

$$N_i(\mathbf{A}, \lambda) := \{\mathbf{v} \in V : (\mathbf{A} - \lambda \mathbf{I}_n)^i \mathbf{v} = \mathbf{0}\}$$

is the nullspace of $(\mathbf{A} - \lambda \mathbf{I}_n)^i$, and is called the *generalised eigenspace of index i* with respect to λ .

The *full* generalised eigenspace of \mathbf{A} with respect to λ is defined as

$$\mathbf{0} \cup \bigcup_{i \in \mathbb{N}} N_i(\mathbf{A}, \lambda)$$

That is, it is the union of all generalised eigenspaces with respect to λ , along with the zero vector.

A *Jordan chain* of length k is a sequence of non-zero vectors $(\mathbf{v}_i)_{i=1}^k \subset K^{n,1}$ such that, for some eigenvalue λ of \mathbf{A} ,

$$\mathbf{A}\mathbf{v}_1 = \lambda\mathbf{v}_1, \quad \mathbf{A}\mathbf{v}_i = \lambda\mathbf{v}_i + \mathbf{v}_{i-1}, \quad 2 \leq i \leq k$$

or equivalently,

$$(\mathbf{A} - \lambda \mathbf{I}_n)\mathbf{v}_1 = \mathbf{0}, \quad (\mathbf{A} - \lambda \mathbf{I}_n)\mathbf{v}_i = \mathbf{v}_{i-1}, \quad 1 \leq i \leq k$$

thus all vectors in a Jordan chain are generalised eigenvectors with $\mathbf{v}_i \in N_i(\mathbf{A}, \lambda)$.

Lemma 2.9. *The vectors in a Jordan chain are linearly independent.*

Theorem 2.10. *The dimensions of corresponding generalised eigenspaces of similar matrices are the same.*

We define the *Jordan block* of degree k with eigenvalue λ to be the $k \times k$ matrix $\mathbf{J}_{\lambda,k}$ given by

$$\mathbf{J}_{\lambda,k} = (J_{i,j}) = \begin{cases} \lambda & \text{if } j = i \\ 1 & \text{if } j = i + 1 \\ 0 & \text{otherwise} \end{cases}$$

That is, the main diagonal has values λ , and the superdiagonal has values 1.

For example,

$$\mathbf{J}_{2,3} = \begin{bmatrix} 2 & 1 & 0 \\ 0 & 2 & 1 \\ 0 & 0 & 2 \end{bmatrix}, \quad \mathbf{J}_{i,2} = \begin{bmatrix} i & 1 \\ 0 & i \end{bmatrix}, \quad \mathbf{J}_{-2,4} = \begin{bmatrix} -2 & 1 & 0 & 0 \\ 0 & -2 & 1 & 0 \\ 0 & 0 & -2 & 1 \\ 0 & 0 & 0 & -2 \end{bmatrix}$$

A matrix \mathbf{A} of a transformation T with respect to the basis $(\mathbf{v}_i)_{i=1}^n \subset K^n$ is a Jordan block of degree n if and only if $(\mathbf{v}_i)_{i=1}^n$ is a Jordan chain for \mathbf{A} .

The minimal and characteristic polynomials of $\mathbf{J}_{\lambda,k}$ are given by,

$$\begin{aligned} \mu_{\mathbf{J}_{\lambda,k}}(x) &= (x - \lambda)^k \\ c_{\mathbf{J}_{\lambda,k}}(x) &= (\lambda - x)^k \end{aligned}$$

We denote the $m \times n$ zero matrix by $\mathbf{0}_{m,n}$. If \mathbf{A} is an $m \times m$ matrix, and \mathbf{B} is an $n \times n$ matrix, we define their *direct sum* $\mathbf{A} \oplus \mathbf{B}$ to be the $(m + n) \times (m + n)$ matrix with block form

$$\left[\begin{array}{c|c} \mathbf{A} & \mathbf{0}_{m,n} \\ \hline \mathbf{0}_{n,m} & \mathbf{B} \end{array} \right]$$

A *Jordan basis* for \mathbf{A} is a basis of K^n consisting of one or more Jordan chains. The matrix of a transformation with respect to a Jordan basis is the direct sum of the corresponding Jordan blocks.

Lemma 2.11. *Suppose that $\mathbf{M} = \mathbf{A} \oplus \mathbf{B}$. Then, $c_{\mathbf{M}} = c_{\mathbf{A}} \times c_{\mathbf{B}}$ and $\mu_{\mathbf{M}} = \text{lcm}(\mu_{\mathbf{A}}, \mu_{\mathbf{B}})$.*

Theorem 2.12. *Let \mathbf{A} be an $n \times n$ matrix over \mathbb{C} . Then, there exists a Jordan basis for \mathbf{A} , and hence \mathbf{A} is similar to a matrix $\mathbf{J} = \bigoplus \mathbf{J}_{\lambda,k}$, where the Jordan blocks $\mathbf{J}_{\lambda,k}$ are uniquely determined by \mathbf{A} .*

The matrix \mathbf{J} in the above is called the *Jordan canonical form* or *Jordan normal form* of \mathbf{A} , and is determined uniquely up to the order of the blocks. The field has to be at least \mathbb{C} (or an extension of \mathbb{C}) so that \mathbf{A} has at least one eigenvalue, since \mathbb{C} is algebraically closed.

Theorem 2.13. *Let $A \in \mathbb{C}^{n \times n}$, and suppose $\{\lambda_i\}_{i=1}^r$ are the eigenvalues of A . Then,*

- $$c_{\mathbf{A}}(x) = (-1)^n \prod_{i=1}^r (x - \lambda_i)^{a_i}$$

where a_i is the sum of the degrees of the Jordan blocks of \mathbf{A} of eigenvalue λ_i ;
- $$\mu_{\mathbf{A}}(x) = \prod_{i=1}^r (x - \lambda_i)^{b_i}$$

where b_i is the largest among the degrees of the Jordan blocks of \mathbf{A} of eigenvalue λ_i ;
- \mathbf{A} is diagonalisable if and only if $\mu_{\mathbf{A}}(x)$ has no repeated factors.

2.5 Computing the Jordan Canonical Form

Some terminology may be new or differ slightly from lecture notes, but makes the following algorithm easier to follow.

Suppose a matrix $\mathbf{A} \in \mathbb{C}^{10,10}$ has a characteristic polynomial,

$$c_{\mathbf{A}}(x) = (x - 1)^3(x - 2)^4(x - 3)^2(x - 4)$$

and thus has eigenvalues $\lambda_1 = 1$, $\lambda_2 = 2$, $\lambda_3 = 3$ and $\lambda_4 = 4$.

We say that the eigenvalue 1 has *algebraic multiplicity* $\alpha(\lambda_1) = 3$, because it is repeated as a root of the characteristic polynomial 3 times. The other algebraic multiplicities are then $\alpha(\lambda_2) = 4$, $\alpha(\lambda_3) = 2$, and $\alpha(\lambda_4) = 1$. The sum of the algebraic multiplicities over all eigenvalues is equal to the dimension of the matrix:

$$\sum_{i=1}^r \alpha(\lambda_i) = n$$

The *geometric multiplicity* $\gamma(\lambda_i)$ of an eigenvalue λ_i is the dimension of the kernel of $\mathbf{A} - \lambda_i \mathbf{I}_n$, or, $\text{null}(\mathbf{A} - \lambda_i \mathbf{I}_n)$.

The *generalised geometric multiplicity* $\gamma_k(\lambda_i)$ of an eigenvalue λ_i is the dimension of the kernel of $(\mathbf{A} - \lambda_i \mathbf{I}_n)^k$, or, $\text{null}((\mathbf{A} - \lambda_i \mathbf{I}_n)^k)$.

Now, the JCF of \mathbf{A} will have the eigenvalues along the diagonal:

$$\mathbf{J} = \begin{bmatrix} \lambda_1 & & & & & & & & & & \\ & \lambda_1 & & & & & & & & & \\ & & \lambda_1 & & & & & & & & \\ & & & \lambda_2 & & & & & & & \\ & & & & \lambda_2 & & & & & & \\ & & & & & \lambda_2 & & & & & \\ & & & & & & \lambda_2 & & & & \\ & & & & & & & \lambda_3 & & & \\ & & & & & & & & \lambda_3 & & \\ & & & & & & & & & \lambda_3 & \\ & & & & & & & & & & \lambda_4 \end{bmatrix}$$

with each eigenvalue λ_i appearing $\alpha(\lambda_i)$ times. Note that there are many different possibilities for the orderings of these eigenvalues, but our convention will be to group the same eigenvalues together, and (where possible) to order these groups in increasing order. If the eigenvalues are complex, just pick any sensible ordering.

We will call these groups *Jordan boxes* (not to be confused with Jordan blocks), highlighted below:

$$\mathbf{J} = \begin{bmatrix} \boxed{\lambda_1} & & & & & & & & & & \\ & \boxed{\lambda_1} & & & & & & & & & \\ & & \boxed{\lambda_1} & & & & & & & & \\ & & & \boxed{\lambda_2} & & & & & & & \\ & & & & \boxed{\lambda_2} & & & & & & \\ & & & & & \boxed{\lambda_2} & & & & & \\ & & & & & & \boxed{\lambda_2} & & & & \\ & & & & & & & \boxed{\lambda_3} & & & \\ & & & & & & & & \boxed{\lambda_3} & & \\ & & & & & & & & & \boxed{\lambda_3} & \\ & & & & & & & & & & \boxed{\lambda_4} \end{bmatrix}$$

A Jordan box is like a Jordan block, but we don't necessarily know where the 1s on the superdiagonal are yet. That is, we need to fill a Jordan box with Jordan blocks, and once we have done so for all boxes, we will have determined a Jordan canonical form for \mathbf{A} .

To begin with, the geometric multiplicity tells us how many blocks are in each box. For instance,

$$\begin{bmatrix} \lambda_3 & & \\ & \lambda_3 & \\ & & \lambda_3 \end{bmatrix} \rightarrow \underbrace{\begin{bmatrix} [\lambda_3] & & \\ & [\lambda_3] & \\ & & [\lambda_3] \end{bmatrix}}_{\gamma(\lambda_3)=2} \quad \text{or} \quad \underbrace{\begin{bmatrix} [\lambda_3 & 1] \\ & [\lambda_3] \end{bmatrix}}_{\gamma(\lambda_3)=1}$$

Here, we can have either two 1×1 blocks, if the geometric multiplicity of λ_3 is 2, or a single 2×2 block, if $\gamma(\lambda_3) = 1$.

For 3×3 ,

$$\begin{bmatrix} \lambda_1 & & \\ & \lambda_1 & \\ & & \lambda_1 \end{bmatrix} \rightarrow \underbrace{\begin{bmatrix} [\lambda_1] & & \\ & [\lambda_1] & \\ & & [\lambda_1] \end{bmatrix}}_{\gamma(\lambda_1)=3} \quad \text{or} \quad \underbrace{\begin{bmatrix} [\lambda_1 & 1] \\ & [\lambda_1] & \\ & & \lambda_1 \end{bmatrix}}_{\gamma(\lambda_1)=2} \quad \text{or} \quad \underbrace{\begin{bmatrix} [\lambda_1 & 1 & \\ & \lambda_1 & 1 \\ & & \lambda_1] \end{bmatrix}}_{\gamma(\lambda_1)=1}$$


the Jordan box can contain three 1×1 Jordan blocks, one 2×2 Jordan block and one 1×1 Jordan block, or a single 3×3 Jordan block, if the geometric multiplicity is 3, 2, or 1, respectively.

However, for 4×4 boxes or larger, the geometric multiplicity alone is not sufficient to determine the blocks within the box. For instance,

$$\begin{bmatrix} \lambda & & & \\ & \lambda & & \\ & & \lambda & \\ & & & \lambda \end{bmatrix} \rightarrow \underbrace{\begin{bmatrix} [\lambda & 1 & \\ & \lambda & 1 \\ & & \lambda] \\ & & & [\lambda] \end{bmatrix}}_{\gamma(\lambda_1)=2} \quad \text{or} \quad \underbrace{\begin{bmatrix} [\lambda & 1] \\ & [\lambda] \\ & & [\lambda & 1] \\ & & & [\lambda] \end{bmatrix}}_{\gamma(\lambda_1)=2}$$

are both consistent with a geometric multiplicity of 2. So, we have to calculate generalised geometric multiplicities to gain more information.

The generalised geometric multiplicities of index k tell us how many chains exist in each generalised eigenspace of index k , allowing us to determine the lengths of the Jordan chains. For instance, suppose $\alpha(\lambda) = 7$, so we have a 7×7 Jordan box. If $\gamma_1(\lambda) = 4$, $\gamma_2(\lambda) = 6$, $\gamma_3(\lambda) = 7 = \alpha(\lambda)$, then the chains would be:



and the lengths of the chains indicate the dimensions of the Jordan blocks within the Jordan box for λ . We have one chain of length 3, one chain of length 2, and two chains of length 1, so we have,

$$\begin{bmatrix} \begin{bmatrix} \lambda & & \\ & \lambda & \\ & & \lambda \end{bmatrix} & & & \\ & \begin{bmatrix} \lambda & \\ & \lambda \end{bmatrix} & & \\ & & [\lambda] & \\ & & & [\lambda] \end{bmatrix}$$

$$\begin{aligned}
&= \dim \ker \begin{bmatrix} 3 & -2 & 1 & -7 & 1 & 5 \\ 0 & 1 & 4 & -4 & -1 & 1 \\ 1 & -1 & -1 & -1 & 1 & 2 \\ 1 & -1 & -2 & 0 & 1 & 2 \\ 0 & 0 & 1 & -1 & 0 & 0 \\ 0 & 0 & -1 & 1 & 0 & 1 \end{bmatrix} \\
&\xrightarrow{\text{row reduce}} \dim \ker \begin{bmatrix} 1 & 0 & 0 & -2 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}
\end{aligned}$$

There are 5 pivot columns, so $\dim \ker(\mathbf{A} - \lambda_1 \mathbf{I}_6) = 6 - 5 = 1 = \gamma(\lambda_1)$. This is sufficient information to determine a 2×2 Jordan box, so we may stop here, but we will continue to compute the generalised geometric multiplicity $\gamma_2(\lambda_1)$, as it will be helpful later for computing the transformation matrix.

$$\begin{aligned}
\gamma_2(\lambda_1) &= \dim \ker((\mathbf{A} - \lambda_1 \mathbf{I}_6)^2) \\
&= \dim \ker \begin{bmatrix} 3 & -2 & 1 & -7 & 1 & 5 \\ 0 & 1 & 4 & -4 & -1 & 1 \\ 1 & -1 & -1 & -1 & 1 & 2 \\ 1 & -1 & -2 & 0 & 1 & 2 \\ 0 & 0 & 1 & -1 & 0 & 0 \\ 0 & 0 & -1 & 1 & 0 & 1 \end{bmatrix}^2 \\
&= \dim \ker \begin{bmatrix} 3 & -2 & 4 & 10 & -1 & 6 \\ 0 & 1 & 6 & -6 & -1 & 2 \\ 1 & -1 & -1 & -1 & 0 & 2 \\ 1 & -1 & -2 & 0 & 0 & 2 \\ 0 & 0 & 1 & -1 & 0 & 0 \\ 0 & 0 & -2 & 2 & 0 & 1 \end{bmatrix} \\
&\xrightarrow{\text{row reduce}} \dim \ker \begin{bmatrix} 1 & 0 & 0 & -2 & -1 & 0 \\ 0 & 1 & 0 & 0 & -1 & 0 \\ 0 & 0 & 1 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}
\end{aligned}$$

Important: when calculating $(\mathbf{A} - \lambda_1 \mathbf{I}_6)^2$, do not square the row reduced form we found earlier. You must use the original non-reduced matrix.

There are 4 pivot columns, so $\gamma_2(\lambda_1) = 6 - 4 = 2$, and we have reached the algebraic multiplicity $\alpha(\lambda_1)$, so we may stop here.

Next, we similarly calculate the geometric multiplicity $\gamma(\lambda_2)$:

$$\begin{aligned}
\gamma(\lambda_2) &= \dim \ker(\mathbf{A} - \lambda_2 \mathbf{I}_6) \\
&= \dim \ker \begin{bmatrix} 2 & -2 & 1 & -7 & 1 & 5 \\ 0 & 0 & 4 & -4 & -1 & 1 \\ 1 & -1 & -2 & -1 & 1 & 2 \\ 1 & -1 & -2 & -1 & 1 & 2 \\ 0 & 0 & 1 & -1 & -1 & 0 \\ 0 & 0 & -1 & 1 & 0 & 0 \end{bmatrix}
\end{aligned}$$

$$\xrightarrow{\text{row reduce}} \dim \ker \begin{bmatrix} 1 & -1 & 0 & -3 & 0 & 0 \\ 0 & 0 & 1 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

There are 4 pivot columns, so $\dim \ker(\mathbf{A} - \lambda_2 \mathbf{I}_6) = 6 - 4 = 2 = \gamma(\lambda_2)$. This is insufficient information to determine the Jordan box of λ_2 , so we calculate higher index generalised geometric multiplicities:

$$\begin{aligned} \gamma_2(\lambda_2) &= \dim \ker((\mathbf{A} - \lambda_2 \mathbf{I}_6)^2) \\ &= \dim \ker \begin{bmatrix} 2 & -2 & 1 & -7 & 1 & 5 \\ 0 & 0 & 4 & -4 & -1 & 1 \\ 1 & -1 & -2 & -1 & 1 & 2 \\ 1 & -1 & -2 & -1 & 1 & 2 \\ 0 & 0 & 1 & -1 & -1 & 0 \\ 0 & 0 & -1 & 1 & 0 & 0 \end{bmatrix}^2 \\ &= \dim \ker \begin{bmatrix} -2 & 2 & 2 & 4 & -3 & -4 \\ 0 & 0 & -2 & 2 & 1 & 0 \\ -1 & 1 & 2 & 1 & -2 & -2 \\ -1 & 1 & 2 & 1 & -2 & -2 \\ 0 & 0 & -1 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \\ &\xrightarrow{\text{row reduce}} \dim \ker \begin{bmatrix} 1 & -1 & 0 & -3 & 0 & 2 \\ 0 & 0 & 1 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \end{aligned}$$

Again, ensure you use the original matrix, and not the row reduced matrix we found before.

There are 3 pivot columns, so $\gamma_2(\lambda_2) = 6 - 3 = 3$. We still have not reached the algebraic multiplicity of λ_2 , so we continue to the next generalised eigenspace:

$$\begin{aligned} \gamma_3(\lambda_2) &= \dim \ker((\mathbf{A} - \lambda_2 \mathbf{I}_6)^3) \\ &= \dim \ker \begin{bmatrix} 2 & -2 & 1 & -7 & 1 & 5 \\ 0 & 0 & 4 & -4 & -1 & 1 \\ 1 & -1 & -2 & -1 & 1 & 2 \\ 1 & -1 & -2 & -1 & 1 & 2 \\ 0 & 0 & 1 & -1 & -1 & 0 \\ 0 & 0 & -1 & 1 & 0 & 0 \end{bmatrix}^3 \\ &= \dim \ker \begin{bmatrix} 2 & -2 & -5 & -1 & 5 & 4 \\ 0 & 0 & 1 & -1 & -1 & 0 \\ 1 & -1 & -3 & 0 & 3 & 2 \\ 1 & -1 & -3 & 0 & 3 & 2 \\ 0 & 0 & 1 & -1 & -1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \end{aligned}$$

$$\xrightarrow{\text{row reduce}} \dim \ker \begin{bmatrix} 1 & -1 & 0 & -3 & 0 & 2 \\ 0 & 0 & 1 & -1 & -1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

There are 2 pivot columns, so $\gamma_2(\lambda_2) = 6 - 2 = 4$ and we have reached the algebraic multiplicity of λ_2 , so we stop here.

At this point, we will draw our Jordan chains for λ_2 to keep track of our results:

$$\begin{array}{l|l} i = 3 & \mathbf{w}_3 \\ & \downarrow \\ i = 2 & \mathbf{w}_2 \\ & \downarrow \\ i = 1 & \mathbf{w}_1 \quad \mathbf{u}_1 \end{array}$$

This indicates that we have one 3×3 Jordan block and one 1×1 Jordan block within the 4×4 Jordan box for λ_2 .

For λ_1 , the Jordan chains would be:

$$\begin{array}{l|l} i = 2 & \mathbf{v}_2 \\ & \downarrow \\ i = 1 & \mathbf{v}_1 \end{array}$$

as $\gamma_1(\lambda_1) = 1$ and $\gamma_2(\lambda_1) = 2$. (Doing this is unnecessary as $\gamma_1(\lambda_1)$ alone is sufficient to determine a 2×2 Jordan box, but this generalised procedure will work for boxes of any size.)

Thus, the blocks are:

$$\mathbf{J} = \begin{bmatrix} \begin{bmatrix} \lambda_1 & 1 \\ & \lambda_1 \end{bmatrix} & & & & & \\ & \begin{bmatrix} \lambda_2 & & & & & \\ & \lambda_2 & 1 & & & \\ & & \lambda_2 & 1 & & \\ & & & \lambda_2 & & \end{bmatrix} & & & & & \\ & & & & & & & & & & \end{bmatrix} = \begin{bmatrix} 2 & 1 & 0 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 3 & 0 & 0 & 0 \\ 0 & 0 & 0 & 3 & 1 & 0 \\ 0 & 0 & 0 & 0 & 3 & 1 \\ 0 & 0 & 0 & 0 & 0 & 3 \end{bmatrix}$$

Next, we will compute the transformation matrix. To calculate the Jordan chains, we will begin by finding bases for each of the generalised eigenspaces we have found so far (again, using MA106 techniques).

For λ_1 , we have,

$$(\mathbf{A} - \lambda_1 \mathbf{I}_6) \mathbf{v} = \mathbf{0} \xrightarrow{\text{row reduce}} \begin{bmatrix} 1 & 0 & 0 & -2 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} a \\ b \\ c \\ d \\ e \\ f \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}$$

$$a = 2d$$

$$c = d$$

$$\begin{bmatrix} a \\ b \\ c \\ d \\ e \\ f \end{bmatrix} = d \begin{bmatrix} 2 \\ 0 \\ 1 \\ 1 \\ 0 \\ 0 \end{bmatrix} \longrightarrow \ker(\mathbf{A} - \lambda_1 \mathbf{I}_6) = \text{span} \left(\begin{bmatrix} 2 \\ 0 \\ 1 \\ 1 \\ 0 \\ 0 \end{bmatrix} \right)$$

$$((\mathbf{A} - \lambda_1 \mathbf{I}_6)^2) \mathbf{v} = \mathbf{0} \xrightarrow{\text{row reduce}} \begin{bmatrix} 1 & 0 & 0 & -2 & -1 & 0 \\ 0 & 1 & 0 & 0 & -1 & 0 \\ 0 & 0 & 1 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} a \\ b \\ c \\ d \\ e \\ f \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}$$

$$a = 2d + e$$

$$b = e$$

$$c = d$$

$$\begin{bmatrix} a \\ b \\ c \\ d \\ e \\ f \end{bmatrix} = d \begin{bmatrix} 2 \\ 0 \\ 1 \\ 1 \\ 0 \\ 0 \end{bmatrix} + e \begin{bmatrix} 1 \\ 1 \\ 0 \\ 0 \\ 1 \\ 0 \end{bmatrix} \rightarrow \ker(\mathbf{A} - \lambda_1 \mathbf{I}_6) = \text{span} \left(\begin{bmatrix} 2 \\ 0 \\ 1 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \\ 0 \\ 0 \\ 1 \\ 0 \end{bmatrix} \right)$$

For λ_2 , we have,

$$(\mathbf{A} - \lambda_2 \mathbf{I}_6) \mathbf{v} = \mathbf{0} \xrightarrow{\text{row reduce}} \begin{bmatrix} 1 & -1 & 0 & -3 & 0 & 0 \\ 0 & 0 & 1 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} a \\ b \\ c \\ d \\ e \\ f \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}$$

$$a = b + 3d$$

$$c = d$$

$$\begin{bmatrix} a \\ b \\ c \\ d \\ e \\ f \end{bmatrix} = b \begin{bmatrix} 1 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} + d \begin{bmatrix} 3 \\ 0 \\ 1 \\ 1 \\ 0 \\ 0 \end{bmatrix} \rightarrow \ker(\mathbf{A} - \lambda_2 \mathbf{I}_6) = \text{span} \left(\begin{bmatrix} 1 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 3 \\ 0 \\ 1 \\ 1 \\ 0 \\ 0 \end{bmatrix} \right)$$

$$((\mathbf{A} - \lambda_2 \mathbf{I}_6)^2) \mathbf{v} = \mathbf{0} \xrightarrow{\text{row reduce}} \begin{bmatrix} 1 & -1 & 0 & -3 & 0 & 2 \\ 0 & 0 & 1 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} a \\ b \\ c \\ d \\ e \\ f \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}$$

$$a = b + 3d - 2f$$

$$c = d$$

$$\begin{bmatrix} a \\ b \\ c \\ d \\ e \\ f \end{bmatrix} = b \begin{bmatrix} 1 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} + d \begin{bmatrix} 3 \\ 0 \\ 1 \\ 1 \\ 0 \\ 0 \end{bmatrix} + f \begin{bmatrix} -2 \\ 0 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} \rightarrow \ker((\mathbf{A} - \lambda_2 \mathbf{I}_6)^2) = \text{span} \left(\begin{bmatrix} 1 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 3 \\ 0 \\ 1 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} -2 \\ 0 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} \right)$$

$$((\mathbf{A} - \lambda_2 \mathbf{I}_6)^3) \mathbf{v} = \mathbf{0} \xrightarrow{\text{row reduce}} \begin{bmatrix} 1 & -1 & 0 & -3 & 0 & 2 \\ 0 & 0 & 1 & -1 & -1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} a \\ b \\ c \\ d \\ e \\ f \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}$$

$$a = b + 3d - 2f$$

$$c = d + e$$

$$\begin{bmatrix} a \\ b \\ c \\ d \\ e \\ f \end{bmatrix} = b \begin{bmatrix} 1 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} + d \begin{bmatrix} 3 \\ 0 \\ 1 \\ 1 \\ 0 \\ 0 \end{bmatrix} + f \begin{bmatrix} -2 \\ 0 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} + e \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \\ 1 \\ 0 \end{bmatrix} \rightarrow \ker((\mathbf{A} - \lambda_2 \mathbf{I}_6)^3) = \text{span} \left(\begin{bmatrix} 1 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 3 \\ 0 \\ 1 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} -2 \\ 0 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \\ 1 \\ 0 \end{bmatrix} \right)$$

I have coloured the basis vectors according to which index they first appeared in, as this will be helpful for our next step.

Recall the Jordan chains we have found:

	λ_1	λ_2
$i = 3$		\mathbf{w}_3
		\downarrow
$i = 2$	\mathbf{v}_2	\mathbf{w}_2
	\downarrow	\downarrow
$i = 1$	\mathbf{v}_1	$\mathbf{w}_1 \quad \mathbf{u}_1$

We begin by choose a vector to be \mathbf{v}_2 . From the diagram above, it lies in the generalised eigenspace of index 2 for λ_1 , but not of 1. We have one obvious option,

$$\mathbf{v}_2 = \begin{bmatrix} 1 \\ 1 \\ 0 \\ 0 \\ 1 \\ 0 \end{bmatrix}$$

We then compute \mathbf{v}_1 :

$$\mathbf{v}_1 = (\mathbf{A} - \lambda_1 \mathbf{I}_n) \mathbf{v}_2$$

$$= \begin{bmatrix} 2 \\ 0 \\ 1 \\ 1 \\ 0 \\ 0 \end{bmatrix}$$

completing the chain, and the eigenvalue.

Moving on to λ_2 , we select the largest chain, so we have to choose $\mathbf{w}_3 \in \ker((\mathbf{A} - \lambda_2 \mathbf{I}_n)^3) \setminus \ker((\mathbf{A} -$

$$= \begin{bmatrix} 2 & 1 & 3 & -1 & 2 & 0 \\ 0 & 1 & 0 & -1 & 3 & 0 \\ 1 & 0 & 1 & 0 & -1 & 1 \\ 1 & 0 & 1 & 0 & -1 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & -1 & 0 \end{bmatrix}$$

2.6 Review

Theorem 2.14. Let $\mathbf{A} \in \mathbb{C}^{n \times n}$ be a square matrix with complex entries, and let $p \in \mathbb{C}[x]$ be any polynomial. Then if λ is an eigenvalue of \mathbf{A} , then $p(\lambda)$ is an eigenvalue of $p(\mathbf{A})$, and any eigenvalue of $p(\mathbf{A})$ is of this form.

Theorem 2.15. For $A \in \mathbb{C}^{n \times n}$, define

$$N_i(\mathbf{A}, \lambda) := \ker((\mathbf{A} - \lambda \mathbf{I})^i)$$

Suppose $\mathbf{A}, \mathbf{B} \in \mathbb{C}^{n \times n}$ are similar. That is, there exists an invertible matrix $\mathbf{S} \in \mathbb{C}^{n \times n}$ such that $\mathbf{B} = \mathbf{S}^{-1}\mathbf{A}\mathbf{S}$. Then, \mathbf{A} and \mathbf{B} share the same set of eigenvalues,

$$\lambda_1 = \mu_1, \dots, \lambda_k = \mu_k$$

and moreover,

$$\dim N_i(\mathbf{A}, \lambda_j) = \dim N_i(\mathbf{B}, \mu_j) \quad \text{for all } i, j$$

Or, using our earlier notation,

$$\gamma_i(\lambda_j) = \gamma_i(\mu_j) \quad \text{for all } i, j$$

The converse of this result also holds. That is, if $\mathbf{A}, \mathbf{B} \in \mathbb{C}^{n \times n}$ share the same eigenvalues and satisfy the equations above, then \mathbf{A} and \mathbf{B} are similar.

3 Matrix Functions

3.1 Matrix Powers

Suppose we wish to compute \mathbf{A}^n for a general matrix \mathbf{A} and large exponent $n \gg 1$.

Ideally, \mathbf{A} is diagonalisable, and we may compute,

$$\begin{aligned} \mathbf{A}^n &= (\mathbf{P}^{-1}\mathbf{D}\mathbf{P})^n \\ &= (\mathbf{P}^{-1}\mathbf{D}\mathbf{P})(\mathbf{P}^{-1}\mathbf{D}\mathbf{P}) \dots (\mathbf{P}^{-1}\mathbf{D}\mathbf{P}) \\ &= \mathbf{P}^{-1}\mathbf{D}(\mathbf{P}\mathbf{P}^{-1})\mathbf{D}(\mathbf{P} \dots \mathbf{P}^{-1})\mathbf{D}\mathbf{P} \\ &= \mathbf{P}^{-1}\mathbf{D}^n\mathbf{P} \end{aligned}$$

and powers of diagonal matrices are trivial to compute:

$$\mathbf{D}^n = \begin{bmatrix} \lambda_1 & & & \\ & \lambda_2 & & \\ & & \ddots & \\ & & & \lambda_k \end{bmatrix}^n = \begin{bmatrix} \lambda_1^n & & & \\ & \lambda_2^n & & \\ & & \ddots & \\ & & & \lambda_k^n \end{bmatrix}$$

But, as we have already seen, not every matrix admits an eigenbasis. On the other hand, every matrix does have a Jordan canonical form, and by a similar telescoping sum, we have,

$$\mathbf{A}^n = \mathbf{P}^{-1}\mathbf{J}^n\mathbf{P}$$

The problem is now to efficiently compute powers of a matrix in Jordan canonical form.

Theorem 3.1. For any square matrices \mathbf{A}, \mathbf{B} ,

$$(\mathbf{A} \oplus \mathbf{B})^n = \mathbf{A}^n \oplus \mathbf{B}^n$$

and more generally, for any collection of square matrices $(\mathbf{A}_i)_{i=1}^k$

$$\left(\bigoplus_{i=1}^k \mathbf{A}_i \right)^n = \bigoplus_{i=1}^k \mathbf{A}_i^n$$

Recall that \mathbf{J} is the direct sum of Jordan blocks, so if these blocks are sufficiently small or simple, this can simplify the calculation greatly. We also have a general formula for larger Jordan blocks:

Theorem 3.2. For any Jordan block $\mathbf{J}_{\lambda,k}$,

$$\mathbf{J}_{\lambda,k}^n = \begin{bmatrix} \binom{n}{0}\lambda^n & \binom{n}{1}\lambda^{n-1} & \cdots & \binom{n}{k-2}\lambda^{n-k+2} & \binom{n}{k-1}\lambda^{n-k+1} \\ 0 & \binom{n}{0}\lambda^n & \cdots & \binom{n}{k-3}\lambda^{n-k+3} & \binom{n}{k-2}\lambda^{n-k+2} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & \binom{n}{0}\lambda^n & \binom{n}{1}\lambda^{n-1} \\ 0 & 0 & \cdots & 0 & \binom{n}{0}\lambda^n \end{bmatrix}$$

noting that $\binom{n}{k} = 0$ whenever $k > n$.

3.2 Lagrange Interpolation

Another way to compute arbitrary powers of matrices is to use *Lagrange interpolation*.

Theorem 3.3 (Lagrange Interpolation). Suppose $\psi(\mathbf{A}) = \mathbf{0}_n$ for a polynomial $\psi \in \mathbb{C}[x]$, and furthermore suppose ψ has roots $(\alpha_i)_{i=1}^k$ with corresponding (algebraic) multiplicities $(m_i)_{i=1}^k$. (In practice, we would choose $\psi = c_{\mathbf{A}}$ or $\psi = \mu_{\mathbf{A}}$.)

Then, for any sufficiently well-behaved* function $f : \mathbb{C} \rightarrow \mathbb{C}$, there exists a function q such that

$$f = q\psi + r$$

where r is a polynomial of degree strictly lower than ψ and

$$f^{(t)}(\alpha_i) = r^{(t)}(\alpha_i)$$

for all $1 \leq j \leq k$, $0 \leq t < m_j$, and furthermore, $f(\mathbf{A}) = r(\mathbf{A})$.

The idea is that we compute a polynomial, r , that acts effectively like a Taylor polynomial near the roots of ψ .

Example. Find a general formula for $f(\mathbf{A}) = \mathbf{A}^n$, where

$$\mathbf{A} = \begin{bmatrix} 3 & 1 & 0 & 1 \\ -1 & 5 & 4 & 1 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 4 \end{bmatrix}$$

You may use that $\mu_{\mathbf{A}}(x) = (2 - x)(4 - x)^2$ without proof.

* f must be analytic in a neighbourhood around every $(\alpha_i)_{i=1}^k$.

Here, $f(z) = z^n$, which is certainly a well-behaved function with easily computable derivatives so we may attempt Lagrange interpolation.

$\mu_{\mathbf{A}}$ is a cubic, so we may choose $r(z) = \alpha z^2 + \beta z + \gamma$.

We know that f and $\mu_{\mathbf{A}}$ agree at the roots of $\mu_{\mathbf{A}}$, so we compute,

$$\begin{aligned} f(2) = 2^n &= r(2) = 4\alpha + 2\beta + \gamma \\ f(4) = 4^n &= r(4) = 16\alpha + 4\beta + \gamma \\ f'(4) = n4^{n-1} &= r'(4) = 8\alpha + \beta \end{aligned}$$

Then,

$$\begin{aligned} r(\mathbf{A}) &= \alpha \mathbf{A}^2 + \beta \mathbf{A} + \gamma \mathbf{I}_4 \\ &= \alpha \begin{bmatrix} 8 & 8 & 4 & 8 \\ -8 & 24 & 28 & 8 \\ 0 & 0 & 4 & 0 \\ 0 & 0 & 0 & 16 \end{bmatrix} + \beta \begin{bmatrix} 3 & 1 & 0 & 1 \\ -1 & 5 & 4 & 1 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 4 \end{bmatrix} + \gamma \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \\ &= \begin{bmatrix} 8\alpha + 3\beta + \gamma & 8\alpha + \beta & 4\alpha & 8\alpha + \beta \\ -8\alpha - \beta & 24\alpha + 5\beta + \gamma & 28\alpha + 4\beta & 8\alpha + \beta \\ 0 & 0 & 4\alpha + 2\beta + \gamma & 0 \\ 0 & 0 & 0 & 16\alpha + 4\beta + \gamma \end{bmatrix} \\ &= \begin{bmatrix} 4^n - n4^{n-1} & n4^{n-1} & 2^n - 4^n + 2n4^{n-1} & n4^{n-1} \\ -n4^{n-1} & 4^n + n4^{n-1} & 4^n - 2^n + 2n4^{n-1} & n4^{n-1} \\ 0 & 0 & 2^n & 0 \\ 0 & 0 & 0 & 4^n \end{bmatrix} \end{aligned}$$

Note that we did not have to calculate the individual values of α , β and γ , (which are

$$\begin{aligned} \alpha &= \frac{1}{4} \cdot 2^n - \frac{1}{4} \cdot 4^n + \frac{1}{2} n4^{n-1} \\ \beta &= -2 \cdot 2^n - 2 \cdot 4^n + 5n4^{n-1} \\ \gamma &= 4 \cdot 2^n + 5 \cdot 4^n - 12n4^{n-1} \end{aligned}$$

for those interested) as the entries in $r(\mathbf{A})$ are simple linear combinations of the values of $r(2)$, $r(4)$, and $r'(4)$. For instance, in the first entry, we have,

$$\begin{aligned} 8\alpha + 3\beta + \gamma &= (16\alpha + 4\beta + \gamma) - (8\alpha + \beta) \\ &= r(4) - r'(4) \\ &= 4^n - n4^{n-1} \end{aligned}$$

with every other entry being computed similarly.

We've only been using Lagrange interpolation to calculate powers of matrices here, but the technique works identically for any sufficiently well-behaved function $f : \mathbb{C} \rightarrow \mathbb{C}$.

3.3 Matrix Exponentials

3.3.1 Recurrence Relations

Consider a vector-valued first-order recurrence relation,

$$\mathbf{x}_n = \mathbf{A}\mathbf{x}_{n-1}$$

where $(\mathbf{x}_i)_{i=1}^{\infty} \subset K^m$ is a sequence of vectors. We will only be considering autonomous recurrence relations – that is, the matrix \mathbf{A} is not a function of n . If a value for \mathbf{x}_0 is given, then these equations are also called (*discrete*) *initial value problems*.

These recurrence relations can be solved analogously to the scalar-valued case with back substitution:

$$\begin{aligned}\mathbf{x}_n &= \mathbf{A}\mathbf{x}_{n-1} \\ &= \mathbf{A}^2\mathbf{x}_{n-2} \\ &= \mathbf{A}^3\mathbf{x}_{n-3} \\ &\vdots \\ &= \mathbf{A}^n\mathbf{x}_0\end{aligned}$$

However, in this case, we now have to calculate an arbitrary power of a matrix, but we can use techniques from the last section to do so.

One application of this is in solving higher-order scalar-valued autonomous recurrence relations. For example, consider the Fibonacci sequence, given by the second-order recurrence relation,

$$F_n = F_{n-1} + F_{n-2}$$

We can easily rewrite this as,

$$\begin{aligned}\begin{bmatrix} F_n \\ F_{n+1} \end{bmatrix} &= \begin{bmatrix} 0 & 1 \\ 1 & 1 \end{bmatrix} \begin{bmatrix} F_{n-1} \\ F_n \end{bmatrix} \\ \mathbf{v}_n &= \begin{bmatrix} 0 & 1 \\ 1 & 1 \end{bmatrix} \mathbf{v}_{n-1}\end{aligned}$$

and now we have a single first-order vector-valued recurrence relation, and this works more generally – we can transform an n th-order recurrence relation into a first-order vector-valued recurrence relation in n dimensions.

3.3.2 Differential Equations

Now, suppose we have a system of first-order linear autonomous simultaneous differential equations, say,

$$\begin{aligned}x' &= 7x - 2y + 9z \\ y' &= 7x + 3y - 5z \\ z' &= 2x + 5y + 6z\end{aligned}$$

We can alternatively interpret these separate variables as a single vector:

$$\begin{bmatrix} x' \\ y' \\ z' \end{bmatrix} = \begin{bmatrix} 7x - 2y + 9z \\ 7x + 3y - 5z \\ 2x + 5y + 6z \end{bmatrix}$$

and, factoring out the matrix, we can represent this system as a single first-order vector-valued differential equation:

$$\mathbf{v}' = \mathbf{A}\mathbf{v}$$

where,

$$\mathbf{v} = \begin{bmatrix} x \\ y \\ z \end{bmatrix}, \quad \text{and} \quad \mathbf{A} = \begin{bmatrix} 7 & -2 & 9 \\ 7 & 3 & -5 \\ 2 & 5 & 6 \end{bmatrix}$$

Compare this to the case of an ordinary first-order differential equation,

$$x' = ax$$

where a is some constant. The solution to this differential equation is given by,

$$x(t) = e^{at}x(0)$$

Similarly, the vector-valued differential equation,

$$\mathbf{v}' = \mathbf{A}\mathbf{v}$$

has a solution given by

$$\mathbf{v}(t) = e^{\mathbf{A}t}\mathbf{v}(0)$$

But, what does $e^{\mathbf{A}t}$ mean? Clearly,

$$e^{\begin{bmatrix} 3 & 1 & 4 \\ 1 & 5 & 9 \\ 2 & 6 & 5 \end{bmatrix}} = \underbrace{e \times e \times \cdots \times e}_{\begin{bmatrix} 3 & 1 & 4 \\ 1 & 5 & 9 \\ 2 & 6 & 5 \end{bmatrix} \text{ times?}}$$

is meaningless.

Instead, recall the Taylor series of $f(x) = e^x$ for real inputs $x \in \mathbb{R}$:

$$e^x = \sum_{n=0}^{\infty} \frac{x^n}{n!}$$

Unlike the expression on the left, it does make sense for us to input things other than real numbers into the series on the right, even if those objects do not immediately make sense as exponents. For instance, we could input complex numbers, or even matrices to this expression.

While the equation above is a *theorem* for real numbers, it's a *definition* for more exotic inputs, like complex numbers or matrices,

$$e^x := \sum_{n=0}^{\infty} \frac{x^n}{n!}, \quad x \in \mathbb{C}, K^{n \times n}, \dots$$

and we sometimes prefer using the notation $\exp(x)$ instead of e^x to emphasise this point more. There are some issues of convergence – after all, why should we expect this series to converge for matrix inputs just because it converges for real inputs – but that is a relatively easy exercise in analysis.

We can use similar techniques from calculating matrix powers before. In particular, \exp is entire, so Lagrange interpolation also applies, and is, in general (for non-diagonalisable matrices), simpler than using a JCF decomposition.

Example. Solve the system of differential equations,

$$\begin{aligned} x' &= x - 3z \\ y' &= x - y - 6z \\ z' &= -x + 2y + 5z \end{aligned}$$

with initial conditions,

$$\begin{aligned} x(0) &= 1 \\ y(0) &= 1 \\ z(0) &= 0 \end{aligned}$$

First, we write the system as a vector-valued differential equation:

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix}' = \underbrace{\begin{bmatrix} 1 & 0 & -3 \\ 1 & -1 & -6 \\ -1 & 2 & 5 \end{bmatrix}}_{\mathbf{A}} \begin{bmatrix} x \\ y \\ z \end{bmatrix}$$

Then, we find the characteristic polynomial:

$$\begin{aligned} c_{\mathbf{A}}(z) &= \det(\mathbf{A} - z\mathbf{I}_3) \\ &= -z^3 + 5z^2 - 8z + 4 \\ &= (1 - z)(2 - z)^2 \end{aligned}$$

so we have roots $\lambda_1 = 1$, and $\lambda_2 = 2$, with multiplicities $\alpha(\lambda_1) = 1$ and $\alpha(\lambda_2) = 2$.

We interpolate $f(z) = e^{zt}$ with $r(z) = \alpha z^2 + \beta z + \gamma$:

$$\begin{aligned} f(1) = e^t &= r(1) = \alpha + \beta + \gamma \\ f(2) = e^{2t} &= r(2) = 4\alpha + 2\beta + \gamma \\ f'(2) = te^{2t} &= r'(2) = 4\alpha + \beta \end{aligned} \quad \begin{cases} \alpha = (t-1)e^{2t} + e^t \\ \beta = (4-3t)e^{2t} - 4e^t \\ \gamma = (2t-3)e^{2t} + 4e^t \end{cases}$$

$$\begin{aligned} r(\mathbf{A}) &= \alpha \mathbf{A}^2 + \beta \mathbf{A} + \gamma \mathbf{I}_3 \\ &= \alpha \begin{bmatrix} 4 & -6 & -18 \\ 6 & -11 & -27 \\ -4 & 8 & 16 \end{bmatrix} + \beta \begin{bmatrix} 1 & 0 & -3 \\ 1 & -1 & -6 \\ -1 & 2 & 5 \end{bmatrix} + \gamma \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \\ &= \begin{bmatrix} 4\alpha + \beta + \gamma & -6\alpha & -18\alpha - 3\beta \\ 6\alpha + \beta & -11\alpha - \beta - \gamma & -27\alpha - 6\beta \\ -4\alpha - \beta & 8\alpha + 2\beta & 16\alpha + 5\beta + \gamma \end{bmatrix} \\ &= \begin{bmatrix} (3t-3)e^{2t} + 4e^t & (6-6t)e^{2t} - 6e^t & (6-9t)e^{2t} - 6e^t \\ (3t-2)e^{2t} + 2e^t & (4-6t)e^{2t} - 3e^t & (3-9t)e^{2t} - 3e^t \\ -te^{2t} & 2te^{2t} & (3t+1)e^{2t} \end{bmatrix} \end{aligned}$$

and so,

$$\begin{aligned} \begin{bmatrix} x(t) \\ y(t) \\ z(t) \end{bmatrix} &= e^{\mathbf{A}t} \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} \\ &= \begin{bmatrix} (3-3t)e^{2t} - 2e^t \\ (2-3t)e^{2t} - e^t \\ te^{2t} \end{bmatrix} \end{aligned}$$

4 Bilinear Maps

Let V and W be vector spaces over a field K . A *bilinear map* on V and W is a map $\tau : V \times W \rightarrow K$ such that for all $\mathbf{v}, \mathbf{v}_1, \mathbf{v}_2 \in V$, $\mathbf{w}, \mathbf{w}_1, \mathbf{w}_2 \in W$, and $\alpha, \beta \in K$,

1. $\tau(\alpha \mathbf{v}_1 + \beta \mathbf{v}_2, \mathbf{w}) = \alpha \tau(\mathbf{v}_1, \mathbf{w}) + \beta \tau(\mathbf{v}_2, \mathbf{w})$;
2. $\tau(\mathbf{v}, \alpha \mathbf{w}_1 + \beta \mathbf{w}_2) = \alpha \tau(\mathbf{v}, \mathbf{w}_1) + \beta \tau(\mathbf{v}, \mathbf{w}_2)$.

That is, for a fixed \mathbf{v} , $\tau(\mathbf{v}, \mathbf{w})$ is linear in \mathbf{w} , and for a fixed \mathbf{w} , $\tau(\mathbf{v}, \mathbf{w})$ is linear in \mathbf{v} .

So, if we fix bases of V and W , a bilinear map is completely determined by its actions on the basis vectors. Let $(\mathbf{e}_i)_{i=1}^n$ and $(\mathbf{f}_i)_{i=1}^m$ be bases of V and W , respectively. Then, the $n \times m$ matrix $\mathbf{A} = (\alpha_{i,j})$ defined by $\alpha_{i,j} = \tau(\mathbf{e}_i, \mathbf{f}_j)$ is said to be the matrix of τ with respect to the bases $(\mathbf{e}_i)_{i=1}^n$ and $(\mathbf{f}_i)_{i=1}^m$.

Then, for any vectors,

$$\mathbf{v} = \sum_{i=1}^n a_i \mathbf{e}_i, \quad \mathbf{w} = \sum_{i=1}^m b_i \mathbf{f}_i$$

we have by bilinearity,

$$\begin{aligned} \tau(\mathbf{v}, \mathbf{w}) &= \tau\left(\sum_{i=1}^n a_i \mathbf{e}_i, \sum_{j=1}^m b_j \mathbf{f}_j\right) \\ &= \sum_{i=1}^n a_i \tau\left(\mathbf{e}_i, \sum_{j=1}^m b_j \mathbf{f}_j\right) \\ &= \sum_{i=1}^n \sum_{j=1}^m a_i \tau(\mathbf{e}_i, \mathbf{f}_j) b_j \\ &= \sum_{i=1}^n \sum_{j=1}^m a_i \alpha_{i,j} b_j \\ &= \mathbf{v}^\top \mathbf{A} \mathbf{w} \end{aligned}$$

So, for any fixed bases of V and W , every bilinear map on V and W corresponds to a unique $n \times m$ matrix, and conversely, every $n \times m$ matrix determines a bilinear map.

Example. Write down the matrix corresponding to the bilinear map τ defined by

$$\tau(\mathbf{v}, \mathbf{w}) := v_1 w_1 - v_1 w_2 + 2v_2 w_1$$

First, expand out the formula above with a general matrix:

$$\begin{aligned} \tau(\mathbf{v}, \mathbf{w}) &= \begin{bmatrix} v_1 \\ v_2 \end{bmatrix}^\top \begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} w_1 \\ w_2 \end{bmatrix} \\ &= \begin{bmatrix} v_1 & v_2 \end{bmatrix} \begin{bmatrix} aw_1 + bw_2 \\ cw_1 + dw_2 \end{bmatrix} \\ &= av_1 w_1 + bv_1 w_2 + cv_2 w_1 + dv_2 w_2 \end{aligned}$$

Equating coefficients, we have,

$$\begin{aligned} a &= 1 \\ b &= -1 \\ c &= 2 \\ d &= 0 \end{aligned}$$

$$\mathbf{A} = \begin{bmatrix} 1 & -1 \\ 2 & 0 \end{bmatrix}$$

4.1 Bilinear Forms

Theorem 4.1. Let \mathbf{A} be the matrix of the bilinear map $\tau : V \times W \rightarrow K$ with respect to the bases $(\mathbf{e}_i)_{i=1}^n$ and $(\mathbf{f}_i)_{i=1}^m$ of V and W , and let \mathbf{B} be its matrix with respect to the bases $(\mathbf{e}'_i)_{i=1}^n$ and $(\mathbf{f}'_i)_{i=1}^m$ of V and W . Let P and Q be the change of basis matrices. Then,

$$\mathbf{B} = \mathbf{P}^\top \mathbf{A} \mathbf{Q}$$

Now, we consider the case where $W = V$. Then, a bilinear map $\tau : V \times V \rightarrow K$ is called a *bilinear form* on V .

The previous theorem then becomes,

Theorem 4.2. Let \mathbf{A} be the matrix of the bilinear form τ on V with respect to the basis $(\mathbf{e}_i)_{i=1}^n$ of V , and let \mathbf{B} be its matrix with respect to the basis $(\mathbf{e}'_i)_{i=1}^n$, and let P be the change of basis matrix. Then,

$$\mathbf{B} = \mathbf{P}^\top \mathbf{A} \mathbf{P}$$

If \mathbf{A} and \mathbf{B} satisfy this relation, they are said to be *congruent* matrices.

Note that congruence is distinct from similarity in that, if τ is a bilinear form on V and T is a linear operator on V , it might be the case that τ and T have the same matrix in some specific basis of V , but they do not necessarily have the same matrix in any other basis of V .

The *rank* of a bilinear form τ is the rank of its matrix \mathbf{A} .

A vector $\mathbf{v} \in K^n$ is zero if and only if $\mathbf{v}^\top \mathbf{w} = \mathbf{0}$ for all vectors $\mathbf{w} \in K^n$. Since,

$$\tau(\mathbf{v}, \mathbf{w}) = \mathbf{v}^\top \mathbf{A} \mathbf{w}$$

the kernel of \mathbf{A} is equal to,

$$\text{span}\{\mathbf{v} \in V : \tau(\mathbf{w}, \mathbf{v}) = 0 \forall \mathbf{w} \in V\}$$

which is also called the *right radical* of τ , and the kernel of \mathbf{A}^\top is equal to,

$$\text{span}\{\mathbf{v} \in V : \tau(\mathbf{v}, \mathbf{w}) = 0 \forall \mathbf{w} \in V\}$$

which is also called the *left radical* of τ .

Since \mathbf{A} and \mathbf{A}^\top have the same rank, the left and right radicals both have dimension $n - r$ where r is the rank of τ . In particular, the rank of τ is n if and only if the left and right radicals have dimension 0, and we say that τ is nondegenerate. That is, τ is nondegenerate if and only if its matrix (in any basis) is nonsingular.

A bilinear form τ on V is *symmetric* if $\tau(\mathbf{w}, \mathbf{v}) = \tau(\mathbf{v}, \mathbf{w})$ for all $\mathbf{v}, \mathbf{w} \in V$. τ is *antisymmetric* or *alternating* if $\tau(\mathbf{v}, \mathbf{v}) = 0$ for all $\mathbf{v} \in V$.

The antisymmetry condition implies that for all $\mathbf{v}, \mathbf{w} \in V$,

$$\tau(\mathbf{v} + \mathbf{w}, \mathbf{v} + \mathbf{w}) = \tau(\mathbf{v}, \mathbf{w}) + \tau(\mathbf{w}, \mathbf{v}) = 0$$

and hence,

$$\tau(\mathbf{v}, \mathbf{w}) = -\tau(\mathbf{w}, \mathbf{v})$$

If $2 \neq 0$ in K , then the converse of this result holds: that is, $\tau(\mathbf{v}, \mathbf{w}) = -\tau(\mathbf{w}, \mathbf{v})$ for all $\mathbf{v} \in V$ implies that $\tau(\mathbf{v}, \mathbf{v}) = 0$ for all $\mathbf{v} \in V$.

An $n \times n$ matrix \mathbf{A} is *symmetric* if $\mathbf{A}^\top = \mathbf{A}$, and *antisymmetric* if $\mathbf{A}^\top = -\mathbf{A}$ and \mathbf{A} has zeros along the diagonal.

Theorem 4.3. *The bilinear form τ is symmetric (resp. antisymmetric) if and only if its matrix (with respect to any basis) is symmetric (resp. antisymmetric).*

One example of a symmetric bilinear form is when $V = \mathbb{R}^n$ and τ is defined by,

$$\begin{aligned}\tau(\mathbf{v}, \mathbf{w}) &= \sum_{i=1}^n v_i w_i \\ &= \langle \mathbf{v}, \mathbf{w} \rangle \\ &= \mathbf{v} \cdot \mathbf{w}\end{aligned}$$

also called the *dot product* or *scalar product*. This bilinear form has matrix form equal to the identity matrix \mathbf{I}_n with respect to the standard basis of \mathbb{R}^n .

Theorem 4.4. *Suppose that $2 \neq 0$ in K . Then, any bilinear form τ can be written uniquely as $\tau_1 + \tau_2$, where τ_1 is symmetric and τ_2 is antisymmetric.*

Proof. For existence, take $\tau_1(\mathbf{v}, \mathbf{w}) = \frac{1}{2}(\tau(\mathbf{v}, \mathbf{w}) + \tau(\mathbf{w}, \mathbf{v}))$ and $\tau_2(\mathbf{v}, \mathbf{w}) = \frac{1}{2}(\tau(\mathbf{v}, \mathbf{w}) - \tau(\mathbf{w}, \mathbf{v}))$.

For uniqueness, suppose τ also decomposes into $\tau'_1 + \tau'_2$, with τ'_1 symmetric and τ'_2 antisymmetric. Then, by symmetry and antisymmetry,

$$\begin{aligned}\tau_1(\mathbf{v}, \mathbf{w}) &= \frac{1}{2}(\tau'_1(\mathbf{v}, \mathbf{w}) + \tau'_1(\mathbf{w}, \mathbf{v}) + \tau'_2(\mathbf{v}, \mathbf{w}) + \tau'_2(\mathbf{w}, \mathbf{v})) \\ &= \frac{1}{2}(\tau'_1(\mathbf{v}, \mathbf{w}) + \tau'_1(\mathbf{v}, \mathbf{w}) + \tau'_2(\mathbf{v}, \mathbf{w}) - \tau'_2(\mathbf{v}, \mathbf{w})) \\ &= \frac{1}{2}(\tau'_1(\mathbf{v}, \mathbf{w}) + \tau'_1(\mathbf{v}, \mathbf{w})) \\ &= \tau'_1(\mathbf{v}, \mathbf{w})\end{aligned}$$

so $\tau_1 = \tau'_1$, and hence,

$$\begin{aligned}\tau_2 &= \tau - \tau_1 \\ &= \tau - \tau'_1 \\ &= \tau_2\end{aligned}$$

so the decomposition is unique.

Note that $\frac{1}{2}$ has to exist in K for the first chain of equations to be meaningful, so we require that $2 \neq 0$ in K . ■

4.2 Quadratic Forms

Let V be a vector space over a field K . A *quadratic form* on V is a function $q : V \rightarrow K$ such that

$$q(\lambda \mathbf{v}) = \lambda^2 q(\mathbf{v})$$

for all $\mathbf{v} \in V$ and $\lambda \in K$, and the function $\tau_q : V \times V \rightarrow K$ defined by,

$$\tau_q(\mathbf{v}, \mathbf{w}) := q(\mathbf{v} + \mathbf{w}) - q(\mathbf{v}) - q(\mathbf{w})$$

is a symmetric bilinear form on V .

Given a symmetric bilinear form τ , we can also define a quadratic form by,

$$q_\tau(\mathbf{v}) := \tau(\mathbf{v}, \mathbf{v})$$

These processes are almost inverse to each other, in that, given a quadratic form q and a bilinear form τ , we have,

$$q_{\tau_q} = 2q, \quad \tau_{q_\tau} = 2\tau$$

so, if $2 \neq 0$ in K , there is a bijection between quadratic forms and symmetric bilinear forms given by,

$$q \mapsto \frac{1}{2}\tau_q, \quad \tau \mapsto q_\tau$$

If $2 = 0$ in K , then this correspondence does not hold, and indeed there exist quadratic forms that are not of the form $\tau(-, -)$ for any symmetric bilinear form τ on V . In general, the standard forms of quadratic and bilinear forms on vector spaces where $2 = 0$ in the underlying field is quite different from the case where $2 \neq 0$.

From this point onwards, we will assume that $2 = 1 + 1 \neq 0$ in the field K .

4.3 Bases for Quadratic Forms

Let $(\mathbf{e}_i)_{i=1}^n$ be a basis of V , and let $\mathbf{A} = (\alpha_{i,j})$ be the matrix of a symmetric bilinear form τ with respect to this basis. \mathbf{A} is then also said to be the matrix of the quadratic form $q := q_\tau$ with respect to this basis.

Note that \mathbf{A} is symmetric, as τ is symmetric. Then,

$$q(\mathbf{v}) = \mathbf{v}^\top \mathbf{A} \mathbf{v}$$

just like in the case for bilinear maps, so we can easily write out the matrix

Example. Write down the matrix corresponding to the quadratic form q defined by

$$q([x, y, z]) := 8x^2 - 7y^2 + 8z^2 + 8xy - 2xz + 8yz$$

As we did for bilinear forms, we again expand $q(\mathbf{v}) = \mathbf{v}^\top \mathbf{A} \mathbf{v}$ for a general matrix \mathbf{A} and vector \mathbf{v} , then compare coefficients:

$$\begin{aligned} [x \quad y \quad z] \begin{bmatrix} a & d & e \\ d & b & f \\ e & f & c \end{bmatrix} \begin{bmatrix} x \\ y \\ z \end{bmatrix} &= ax^2 + by^2 + cz^2 + 2dxy + 2exz + 2fyz \\ &= 8x^2 - 7y^2 + 8z^2 + 8xy - 2xz + 8yz \\ \mathbf{A} &= \begin{bmatrix} 8 & 4 & -1 \\ 4 & -7 & 4 \\ -1 & 4 & 8 \end{bmatrix} \end{aligned}$$

Theorem 4.5. *Let V be a vector space of dimension n equipped with a symmetric bilinear form τ (or equivalently, with a quadratic form q).*

Then, there exist a basis $(\mathbf{b}_i)_{i=1}^n$ of V and constants $(\beta_i)_{i=1}^n$ such that,

$$\tau(\mathbf{b}_i, \mathbf{b}_j) = \begin{cases} \beta_i & i = j \\ 0 & i \neq j \end{cases}$$

Equivalently,

- *Give any quadratic form q on V , there exist a basis $(\mathbf{b}_i)_{i=1}^n$ of V and constants $(\beta_i)_{i=1}^n$ such that,*

$$q\left(\sum_{i=1}^n x_i \mathbf{b}_i\right) = \sum_{i=1}^n \beta_i x_i^2$$

- Any symmetric matrix \mathbf{A} is congruent to a diagonal matrix. That is, there exists an invertible matrix \mathbf{P} such that $\mathbf{A} = \mathbf{P}^T \mathbf{D} \mathbf{P}$, where \mathbf{D} is a diagonal matrix.

We give an algorithm to find the matrix \mathbf{P} .

Algorithm 4 Orthogonal Diagonalisation

- 1: Determine the eigenvalues of \mathbf{A} .
 - 2: For each eigenvalue λ_i , find the corresponding eigenspace, $\ker(\mathbf{A} - \lambda_i \mathbf{I}_n) = \text{span}((\mathbf{v}_j)_{j=1}^k)$.
 - 3: Write out a normal diagonalisation of $\mathbf{A} = \mathbf{P} \mathbf{D} \mathbf{P}^{-1}$, recalling that \mathbf{D} is the matrix with the eigenvalues of \mathbf{A} along the diagonal, and \mathbf{P} is the matrix with the corresponding eigenvectors as columns.
 - 4: Check if the columns of \mathbf{P} are orthogonal by checking if their scalar product is zero or not. If all columns are orthogonal, we are done.
 - 5: Otherwise, apply the Gram-Schmidt process (§4.4) to the columns of \mathbf{P} .
-

Theorem 4.6. For a symmetric real matrix, eigenvectors with distinct eigenvalues are always orthogonal.

This allows us a slight shortcut in the algorithm above: we only need to check the scalar product of eigenvectors that share the same eigenvalue.

Theorem 4.7. A quadratic form q over \mathbb{C} has the form

$$q(\mathbf{v}) = \sum_{i=1}^r x_i^2$$

with respect to a suitable basis, where $r = \text{rank}(q)$.

Equivalently, given a symmetric matrix $\mathbf{A} \in \mathbb{C}^{n \times n}$, there is an invertible matrix $\mathbf{P} \in \mathbb{C}^{n \times n}$ such that $\mathbf{P}^T \mathbf{A} \mathbf{P} = \mathbf{B}$, where $\mathbf{B} = (\beta_{i,j})$ is a diagonal matrix with,

$$\beta_{i,i} = \begin{cases} 1 & i \in [1, r] \\ 0 & i \in (r, n] \end{cases}$$

where $r = \text{rank}(\mathbf{A})$.

In particular, up to a change of basis, a quadratic form on \mathbb{C}^n is uniquely determined by its rank, and we say that the rank is the only *invariant* of a quadratic form over \mathbb{C} .

Theorem 4.8 (Sylvester's Theorem). A quadratic form q over \mathbb{R} has the form

$$q(\mathbf{v}) = \sum_{i=1}^t x_i^2 - \sum_{i=1}^u x_{t+i}^2$$

with respect to a suitable basis, where $t + u = \text{rank}(q)$.

Equivalently, given a symmetric matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$, there is an invertible matrix $\mathbf{P} \in \mathbb{R}^{n \times n}$ such that $\mathbf{A} = \mathbf{P} \mathbf{B} \mathbf{P}^T$, where $\mathbf{B} = (\beta_{i,j})$ is a diagonal matrix with,

$$\beta_{i,i} = \begin{cases} 1 & i \in [1, t] \\ -1 & i \in (t, t + u] \\ 0 & i \in (t + u, n] \end{cases}$$

where $t + u = \text{rank}(\mathbf{A})$.

The numbers t and u of positive and negative terms are invariants of q , and the pair of integers (t, u) is the *signature* of q .

Theorem 4.9 (Sylvester's Law of Inertia). *Suppose that q is a quadratic form on the vector space V over \mathbb{R} , and that $(\mathbf{e}_i)_{i=1}^n$ and $(\mathbf{e}'_i)_{i=1}^n$ are two bases of V such that,*

$$q\left(\sum_{i=1}^n x_i \mathbf{e}_i\right) = \sum_{i=1}^t x_i^2 - \sum_{i=1}^u x_{t+i}^2$$

and

$$q\left(\sum_{i=1}^n x_i \mathbf{e}'_i\right) = \sum_{i=1}^{t'} x_i^2 - \sum_{i=1}^{u'} x_{t'+i}^2$$

Then, $t = t'$ and $u = u'$.

4.4 The Gram-Schmidt Process

In this section, we will take $K = \mathbb{R}$.

Let V be a vector space over K of dimension n , and let q be a quadratic form on V , with associated symmetric bilinear form τ .

The quadratic form q is *positive definite* if $q(\mathbf{v}) > 0$ for all non-zero $\mathbf{v} \in V$. τ is also called positive definite if q is positive definite.

A quadratic form q is positive definite if and only if $t = n$ and $u = 0$ in Sylvester's theorem. That is, if q has signature $(n, 0)$.

A vector space V over \mathbb{R} equipped with a positive definite symmetric bilinear form τ is called a *Euclidean space*. In this case, Sylvester's theorem just states that there is a basis $(\mathbf{e}_i)_{i=1}^n$ of V with respect to which the matrix of q is the identity matrix \mathbf{I}_n .

In other words, the basis vectors are all unit vectors, and they are all orthogonal to each other. Such a basis is called an *orthonormal basis*, and more generally, any set of vectors such that the vectors are all unit length and orthogonal to each other is called *orthonormal*.

If V is a Euclidean space with an orthonormal basis, then any positive definite symmetric bilinear form τ is equivalent to the dot product, and we will write $\mathbf{v} \cdot \mathbf{w}$ for $\tau(\mathbf{v}, \mathbf{w})$.

Given a (finite) set of linearly independent vectors $S = (\mathbf{v}_i)_{i=1}^k$, the *Gram-Schmidt process* generates an orthogonal set $S' = (\mathbf{u}_i)_{i=1}^k$ that spans the same k -dimensional subspace of V as S .

Algorithm 5 Gram-Schmidt Process

1: Define the projection operator by,

$$\text{proj}_{\mathbf{u}}(\mathbf{v}) := \frac{\mathbf{v} \cdot \mathbf{u}}{\mathbf{u} \cdot \mathbf{u}} \mathbf{u}$$

And if $\mathbf{u} = \mathbf{0}$, then we define $\text{proj}_{\mathbf{u}}(\mathbf{v}) = \mathbf{0}$. This operator projects \mathbf{v} orthogonally onto the line spanned by \mathbf{u} .

2: Recursively calculate the sequence of orthogonal vectors $(\mathbf{u}_i)_{i=1}^n$ using,

$$\begin{aligned} \mathbf{u}_1 &= \mathbf{v}_1 \\ \mathbf{u}_2 &= \mathbf{v}_2 - \text{proj}_{\mathbf{u}_1}(\mathbf{v}_2) \\ \mathbf{u}_3 &= \mathbf{v}_3 - \text{proj}_{\mathbf{u}_1}(\mathbf{v}_3) - \text{proj}_{\mathbf{u}_2}(\mathbf{v}_3) \\ \mathbf{u}_4 &= \mathbf{v}_4 - \text{proj}_{\mathbf{u}_1}(\mathbf{v}_4) - \text{proj}_{\mathbf{u}_2}(\mathbf{v}_4) - \text{proj}_{\mathbf{u}_3}(\mathbf{v}_4) \\ &\vdots \\ \mathbf{u}_n &= \mathbf{v}_n - \sum_{i=1}^{n-1} \text{proj}_{\mathbf{u}_i}(\mathbf{v}_n) \end{aligned}$$

3: Normalise each vector to obtain the orthonormal sequence $(\mathbf{e}_i)_{i=1}^n$:

$$\mathbf{e}_i = \frac{\mathbf{u}_i}{\|\mathbf{u}_i\|}$$

The calculation of the orthogonal sequence $(\mathbf{u}_i)_{i=1}^n$ is more specifically called *Gram-Schmidt orthogonalisation*, while the total calculation of the orthonormal sequence $(\mathbf{e}_i)_{i=1}^n$ is called *Gram-Schmidt orthonormalisation* as the vectors are normalised.

Geometrically, to compute the next vector \mathbf{u}_k , we project \mathbf{v}_k onto the subspace $U = \text{span}((\mathbf{u}_i)_{i=1}^{k-1})$ spanned by the previous orthogonal vectors, which by construction, is the same as the subspace spanned by the first $k-1$ original vectors, $\text{span}((\mathbf{v}_i)_{i=1}^{k-1})$. The vector \mathbf{u}_k is then defined to be the difference between \mathbf{v}_k and this projection, guaranteed to be orthogonal to all the vectors in U . Because of this projection action, the vectors $(\mathbf{u}_i)_{i=1}^n$ are sometimes also denoted $(\mathbf{v}_i^\perp)_{i=1}^n$.

Example. In the last example, we found that the matrix corresponding to the quadratic form q defined by

$$q([x, y, z]) := 8x^2 - 7y^2 + 8z^2 + 8xy - 2xz + 8yz$$

is given by

$$\mathbf{A} = \begin{bmatrix} 8 & 4 & -1 \\ 4 & -7 & 4 \\ -1 & 4 & 8 \end{bmatrix}$$

Now, find an orthonormal basis of V such that the matrix of q is diagonal. You may use that $c_{\mathbf{A}}(x) = -(x-9)^2(x+9)$ without proof.

We first find an ordinary diagonalisation of \mathbf{A} . From the characteristic equation, \mathbf{A} has eigenvalues $\lambda_1 = 9$ and $\lambda_2 = -9$, so,

$$(\mathbf{A} - \lambda_1 \mathbf{I}_3)\mathbf{v} = \mathbf{0} \xrightarrow{\text{row reduce}} \begin{bmatrix} 1 & -4 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} a \\ b \\ c \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}$$

$$a = 4b - c$$

$$\begin{bmatrix} a \\ b \\ c \end{bmatrix} = b \begin{bmatrix} 4 \\ 1 \\ 0 \end{bmatrix} + c \begin{bmatrix} -1 \\ 0 \\ 1 \end{bmatrix} \rightarrow \ker(\mathbf{A} - \lambda_1 \mathbf{I}_3) = \text{span} \left(\begin{bmatrix} 4 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} -1 \\ 0 \\ 1 \end{bmatrix} \right)$$

$$(\mathbf{A} - \lambda_2 \mathbf{I}_3) \mathbf{v} = \mathbf{0} \xrightarrow{\text{row reduce}} \begin{bmatrix} 1 & 0 & -1 \\ 0 & 1 & 4 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} a \\ b \\ c \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}$$

$$a = c$$

$$b = -4c$$

$$\begin{bmatrix} a \\ b \\ c \end{bmatrix} = c \begin{bmatrix} 1 \\ -4 \\ 1 \end{bmatrix} \rightarrow \ker(\mathbf{A} - \lambda_2 \mathbf{I}_3) = \text{span} \left(\begin{bmatrix} 1 \\ -4 \\ 1 \end{bmatrix} \right)$$

So, $\mathbf{A} = \mathbf{PDP}^{-1}$ with

$$\mathbf{D} = \begin{bmatrix} 9 & 0 & 0 \\ 0 & 9 & 0 \\ 0 & 0 & -9 \end{bmatrix}$$

$$\mathbf{P} = \begin{bmatrix} -1 & 4 & 1 \\ 0 & 1 & -4 \\ 1 & 0 & 1 \end{bmatrix}$$

$$\begin{bmatrix} -1 \\ 0 \\ 1 \end{bmatrix} \cdot \begin{bmatrix} 4 \\ 1 \\ 0 \end{bmatrix} = 4$$

$$\neq 0$$

so the first two columns of \mathbf{P} are not orthogonal. The last column must be orthogonal to the first two by Theorem 4.6

$$\mathbf{v}_1 = \mathbf{p}_1$$

$$\mathbf{v}_2 = \mathbf{p}_2 - \frac{\mathbf{p}_2 \cdot \mathbf{v}_1}{\mathbf{v}_1 \cdot \mathbf{v}_1} \mathbf{v}_1$$

$$= \mathbf{p}_2 - \frac{-4}{2} \mathbf{v}_1$$

$$= \begin{bmatrix} 2 \\ 1 \\ 2 \end{bmatrix}$$

$$\mathbf{P} = \begin{bmatrix} -1 & 2 & 1 \\ 0 & 1 & -4 \\ 1 & 2 & 1 \end{bmatrix}$$

Normalising:

$$\mathbf{P} = \begin{bmatrix} -\frac{1}{\sqrt{2}} & \frac{2}{3} & \frac{1}{3\sqrt{2}} \\ 0 & \frac{1}{3} & -\frac{4}{3\sqrt{2}} \\ \frac{1}{\sqrt{2}} & \frac{2}{3} & \frac{1}{3\sqrt{2}} \end{bmatrix}$$

So an orthonormal basis of V with q diagonal is given by,

$$\left\{ \begin{bmatrix} -\frac{1}{\sqrt{2}} \\ 0 \\ \frac{1}{\sqrt{2}} \end{bmatrix}, \begin{bmatrix} \frac{2}{3} \\ \frac{1}{3} \\ \frac{1}{3} \end{bmatrix}, \begin{bmatrix} \frac{1}{3\sqrt{2}} \\ -\frac{4}{3\sqrt{2}} \\ \frac{1}{3\sqrt{2}} \end{bmatrix} \right\}$$

4.5 Orthogonal Transformations

In a Euclidean space V , the scalar product gives us a notion of length of a vector and angles between vectors, so we might be interested in what kind of transformations preserve these quantities.

A linear map $T : V \rightarrow V$ is *orthogonal* if it preserves the scalar product on V . That is, $T(\mathbf{v}) \cdot T(\mathbf{w}) = \mathbf{v} \cdot \mathbf{w}$ for all $\mathbf{v}, \mathbf{w} \in V$.

A $n \times n$ matrix \mathbf{A} is *orthogonal* if $\mathbf{A}^\top \mathbf{A} = \mathbf{A} \mathbf{A}^\top = \mathbf{I}_n$, or equivalently, $\mathbf{A}^\top = \mathbf{A}^{-1}$.

Lemma 4.10. *A linear map $T : V \rightarrow V$ is orthogonal if and only if its matrix \mathbf{A} is orthogonal.*

Theorem 4.11. $\det(\mathbf{A}) = \pm 1$ for any orthogonal matrix \mathbf{A} .

Proof.

$$\begin{aligned} \det(\mathbf{A})^2 &= \det(\mathbf{A}) \det(\mathbf{A}^\top) \\ &= \det(\mathbf{A}^\top \mathbf{A}) \\ &= \det(\mathbf{I}_n) \\ &= 1 \end{aligned}$$

so $\det(\mathbf{A}) = \sqrt{1} = \pm 1$. ■

Theorem 4.12. *A linear map $T : V \rightarrow V$ is orthogonal if and only if $(T(\mathbf{e}_i))_{i=1}^n$ is an orthonormal basis of V .*

Theorem 4.13 (QR Decomposition). *Any invertible matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ can be written as $\mathbf{A} = \mathbf{QR}$ where \mathbf{Q} is orthogonal and \mathbf{R} is upper-triangular.*

Algorithm 6 QR Decomposition

- 1: Perform the Gram-Schmidt process on the columns of $\mathbf{A} = [\mathbf{a}_1|\mathbf{a}_2|\dots|\mathbf{a}_n]$ to obtain the orthonormalised set $(\mathbf{e}_i)_{i=1}^n$.
- 2: Express the columns of \mathbf{A} in the orthonormal basis:

$$\begin{aligned}\mathbf{a}_1 &= \langle \mathbf{e}_1, \mathbf{a}_1 \rangle \mathbf{e}_1 \\ \mathbf{a}_2 &= \langle \mathbf{e}_1, \mathbf{a}_2 \rangle \mathbf{e}_1 + \langle \mathbf{e}_2, \mathbf{a}_2 \rangle \mathbf{e}_2 \\ \mathbf{a}_3 &= \langle \mathbf{e}_1, \mathbf{a}_3 \rangle \mathbf{e}_1 + \langle \mathbf{e}_2, \mathbf{a}_3 \rangle \mathbf{e}_2 + \langle \mathbf{e}_3, \mathbf{a}_3 \rangle \mathbf{e}_3 \\ \mathbf{a}_4 &= \langle \mathbf{e}_1, \mathbf{a}_4 \rangle \mathbf{e}_1 + \langle \mathbf{e}_2, \mathbf{a}_4 \rangle \mathbf{e}_2 + \langle \mathbf{e}_3, \mathbf{a}_4 \rangle \mathbf{e}_3 + \langle \mathbf{e}_4, \mathbf{a}_4 \rangle \mathbf{e}_4 \\ &\vdots \\ \mathbf{a}_k &= \sum_{i=1}^k \langle \mathbf{e}_i, \mathbf{a}_k \rangle \mathbf{e}_i\end{aligned}$$

noting that $\langle \mathbf{e}_i, \mathbf{a}_i \rangle = \|\mathbf{v}_i^\perp\|$ (which you had to calculate before during the normalisation step).

- 3: the above set of equations can be packaged into matrix form,

$$\mathbf{A} = \mathbf{QR}$$

where,

$$\mathbf{Q} = [\mathbf{e}_1|\mathbf{e}_2|\dots|\mathbf{e}_n]$$

$$\mathbf{R} = \begin{bmatrix} \|\mathbf{v}_1^\perp\| & \langle \mathbf{e}_1, \mathbf{a}_2 \rangle & \langle \mathbf{e}_1, \mathbf{a}_3 \rangle & \langle \mathbf{e}_1, \mathbf{a}_4 \rangle & \cdots & \langle \mathbf{e}_1, \mathbf{a}_n \rangle \\ 0 & \|\mathbf{v}_2^\perp\| & \langle \mathbf{e}_2, \mathbf{a}_3 \rangle & \langle \mathbf{e}_2, \mathbf{a}_4 \rangle & \cdots & \langle \mathbf{e}_2, \mathbf{a}_n \rangle \\ 0 & 0 & \|\mathbf{v}_3^\perp\| & \langle \mathbf{e}_3, \mathbf{a}_4 \rangle & \cdots & \langle \mathbf{e}_3, \mathbf{a}_n \rangle \\ 0 & 0 & 0 & \|\mathbf{v}_4^\perp\| & \cdots & \langle \mathbf{e}_4, \mathbf{a}_n \rangle \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & \|\mathbf{v}_n^\perp\| \end{bmatrix}$$

Example. Find a QR decomposition of,

$$\mathbf{B} = \begin{bmatrix} 2 & 1 & -1 \\ 1 & 0 & 2 \\ 2 & -1 & 3 \end{bmatrix}$$

We perform the Gram-Schmidt process on the columns of $\mathbf{B} = [\mathbf{b}_1|\mathbf{b}_2|\mathbf{b}_3]$.

$$\begin{aligned}\mathbf{e}_1 &= \frac{\mathbf{b}_1}{\|\mathbf{b}_1\|} \\ &= \frac{1}{3} \begin{bmatrix} 2 \\ 1 \\ 2 \end{bmatrix} \\ &= \begin{bmatrix} \frac{2}{3} \\ \frac{1}{3} \\ \frac{2}{3} \end{bmatrix} \\ \mathbf{b}_2^\perp &= \mathbf{b}_2 - \langle \mathbf{b}_2, \mathbf{e}_1 \rangle \mathbf{e}_1 \\ &= \mathbf{b}_2 - 0\mathbf{e}_1\end{aligned}$$

$$\begin{aligned}
&= \begin{bmatrix} 1 \\ 0 \\ -1 \end{bmatrix} \\
\mathbf{e}_2 &= \frac{\mathbf{b}_2^\perp}{\|\mathbf{b}_2^\perp\|} \\
&= \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 0 \\ -1 \end{bmatrix} \\
&= \begin{bmatrix} \frac{1}{\sqrt{2}} \\ 0 \\ -\frac{1}{\sqrt{2}} \end{bmatrix} \\
\mathbf{b}_3^\perp &= \mathbf{b}_3 - \langle \mathbf{b}_3, \mathbf{e}_1 \rangle \mathbf{e}_1 - \langle \mathbf{b}_3, \mathbf{e}_2 \rangle \mathbf{e}_2 \\
&= \mathbf{b}_3 - 2\mathbf{q}_1 - (-2\sqrt{2})\mathbf{e}_2 \\
&= \begin{bmatrix} -\frac{1}{3} \\ \frac{4}{3} \\ -\frac{1}{3} \end{bmatrix} \\
\mathbf{e}_3 &= \frac{\mathbf{b}_3^\perp}{\|\mathbf{b}_3^\perp\|} \\
&= \frac{1}{\sqrt{2}} \begin{bmatrix} -\frac{1}{3} \\ \frac{4}{3} \\ -\frac{1}{3} \end{bmatrix} \\
&= \begin{bmatrix} -\frac{\sqrt{2}}{6} \\ \frac{2\sqrt{2}}{3} \\ -\frac{\sqrt{2}}{6} \end{bmatrix} \\
\mathbf{Q} &= [\mathbf{e}_1 | \mathbf{e}_2 | \mathbf{e}_3] \\
&= \begin{bmatrix} \frac{2}{3} & \frac{1}{\sqrt{2}} & -\frac{\sqrt{2}}{6} \\ \frac{1}{3} & 0 & \frac{2\sqrt{2}}{3} \\ \frac{2}{3} & -\frac{1}{\sqrt{2}} & -\frac{\sqrt{2}}{6} \end{bmatrix} \\
\mathbf{R} &= \begin{bmatrix} 3 & 0 & 2 \\ 0 & \sqrt{2} & -2\sqrt{2} \\ 0 & 0 & \sqrt{2} \end{bmatrix} \\
\mathbf{B} &= \begin{bmatrix} \frac{2}{3} & \frac{1}{\sqrt{2}} & -\frac{\sqrt{2}}{6} \\ \frac{1}{3} & 0 & \frac{2\sqrt{2}}{3} \\ \frac{2}{3} & -\frac{1}{\sqrt{2}} & -\frac{\sqrt{2}}{6} \end{bmatrix} \begin{bmatrix} 3 & 0 & 2 \\ 0 & \sqrt{2} & -2\sqrt{2} \\ 0 & 0 & \sqrt{2} \end{bmatrix}
\end{aligned}$$

4.6 Orthonormal Bases for Bilinear Forms

Suppose we have a Euclidean space V , and a linear operator $T : V \rightarrow V$ or a quadratic form q on V (not necessarily the same quadratic form as the one providing the Euclidean structure). Is it always possible to find an orthonormal basis of V such that the matrix of q has a simple form? Notice that we are now handling two different matrices simultaneously: we're trying to optimise the matrix of q , while still keeping the matrix of the original quadratic form as the identity.

It turns out that this is also a question about linear operators: given any bilinear form τ on V , we have

$$\tau(\mathbf{v}, \mathbf{w}) = \mathbf{v}^\top \mathbf{A} \mathbf{w}$$

and we can interpret matrix-vector multiplication as a linear map, so,

$$= \mathbf{v} \cdot T(\mathbf{w})$$

So, every bilinear form τ on V uniquely determines a linear operator T on V such that,

$$\tau(\mathbf{v}, \mathbf{w}) = \mathbf{v} \cdot T(\mathbf{w})$$

where T is the linear operator corresponding to the matrix \mathbf{A} of τ with entries $\mathbf{A}_{i,j} = \tau(\mathbf{e}_i, \mathbf{e}_j)$ for the standard basis $(\mathbf{e}_i)_{i=1}^n$ of V . Conversely, any linear operator T similarly determines a bilinear form τ , where bilinearity follows from the bilinearity of the scalar product and linearity of T .

So, once we have a fixed bilinear form providing the Euclidean structure (i.e., the scalar product), any other bilinear form τ on V can be obtained from applying a linear transformation to one of the arguments of the scalar product, so there is a bijection between bilinear forms and linear operators.

In particular, if T is any linear operator, then $(\mathbf{v}, \mathbf{w}) \mapsto (T\mathbf{v}) \cdot \mathbf{w}$ is certainly a bilinear form, so there exists a unique linear operator S such that,

$$(T\mathbf{v}) \cdot \mathbf{w} = \mathbf{v} \cdot (S\mathbf{w})$$

for all $\mathbf{v}, \mathbf{w} \in V$. Such a linear operator S is called the *adjoint* of T , and is alternatively denoted T^* .

If we have chosen an orthonormal basis, then the matrix of T^* is the transpose of the matrix of T . It follows that a linear operator is orthogonal if and only if $T^* = T^{-1}$.

A linear operator T is *selfadjoint* if $T^* = T$, or equivalently, if the bilinear form $\tau(\mathbf{v}, \mathbf{w}) = \mathbf{v} \cdot (T\mathbf{w})$ is symmetric.

So, if V is a Euclidean space of dimension n , then the following problems are all equivalent:

- Given a quadratic form q on V , find an orthonormal basis of V that makes the matrix of q as simple as possible;
- Given a selfadjoint linear operator T on V , find an orthonormal basis of V that makes the matrix of T as simple as possible;
- Given an $n \times n$ symmetric real matrix \mathbf{A} , find an orthogonal matrix P such that $\mathbf{P}^T \mathbf{A} \mathbf{P}$ is as simple as possible.

Lemma 4.14. *Let \mathbf{A} be a $n \times n$ symmetric real matrix. Then, \mathbf{A} has an eigenvalue in \mathbb{R} , and moreover, all complex eigenvalues of \mathbf{A} lie in \mathbb{R} .*

Theorem 4.15 (Spectral Theorem). *Let V be a Euclidean space of dimension n . Then,*

- *Given any quadratic form q on V , there is an orthonormal basis $(\mathbf{f}_i)_{i=1}^n$ of V and constants $(\alpha_i)_{i=1}^n$ uniquely determined up to reordering, such that,*

$$q\left(\sum_{i=1}^n x_i \mathbf{f}_i\right) = \sum_{k=1}^n \alpha_k (x_k)^2$$

for all $x_1, \dots, x_n \in \mathbb{R}$.

- *Given any selfadjoint linear operator $T : V \rightarrow V$, there is an orthonormal basis $(\mathbf{f}_i)_{i=1}^n$ of V consisting of eigenvectors of T .*
- *Given any $n \times n$ symmetric real matrix \mathbf{A} , there is an orthogonal matrix \mathbf{P} such that $\mathbf{P}^T \mathbf{A} \mathbf{P} = \mathbf{P}^{-1} \mathbf{A} \mathbf{P}$ is a diagonal matrix.*

Example. TO DO

4.7 Reduction of Second Degree Polynomial Equations

The general equation of a second degree polynomial in n variables $(x_i)_{i=1}^n$ is given by,

$$\sum_{i=1}^n \alpha_{i,i} x_i^2 + \sum_{i=1}^n \sum_{j=1}^{i-1} \alpha_{i,j} x_i x_j + \sum_{i=1}^n \beta_i x_i + \gamma = 0$$

for an $n \times n$ lower triangular matrix $\mathbf{A} = (\alpha_{i,j})$ of constants and n -dimensional vector $\mathbf{b} = (\beta_i)$ of constants. That is, there is a term for every variable squared, a term for every product of a pair of variables, a term for every variable alone, and a constant term. For any set of fixed coefficients, this equation describes a *quadratic* (hyper)surface in n -dimensional Euclidean space.

For example, for the $n = 3$ case, the general polynomial is given by,

$$Ax^2 + By^2 + Cz^2 + Dxy + Exz + Fyz + Gx + Hy + Iz + J = 0$$

We can simplify these equation by applying various isometries to the coordinate basis. By the spectral theorem, we can apply orthogonal basis changes to eliminate the quadratic terms in mixed variables (the terms with coefficients D , E , and F in the example above). We do this by completing the square in the

$$\sum_{i=1}^n \alpha_{i,i} x_i^2 + \sum_{i=1}^n \sum_{j=1}^{i-1} \alpha_{i,j} x_i x_j$$

term to see what coordinate changes we should effect. For instance, suppose we have the equation,

$$x^2 + xy + y^2 + x = 0$$

We complete the square on $x^2 + xy + y^2$ to obtain,

$$\left(x + \frac{1}{2}y\right)^2 + \frac{3}{4}y^2 + x = 0$$

Now, we apply the linear transformation,

$$\begin{bmatrix} x \\ y \end{bmatrix} \mapsto \begin{bmatrix} x + \frac{1}{2}y \\ y \end{bmatrix}$$

giving,

$$x^2 + \frac{3}{4}y^2 + x - \frac{1}{2}y = 0$$

and we no longer have any mixed quadratic terms.

Now, whenever $\alpha_{i,i} \neq 0$, we can perform the translation isometry,

$$x_i \mapsto x_i - \frac{\beta_i}{2\alpha_{i,i}}$$

thus eliminating the term $\beta_i x_i$.

From example above, we would have,

$$\begin{bmatrix} x \\ y \end{bmatrix} \mapsto \begin{bmatrix} x - \frac{1}{2 \cdot 1} \\ y - \frac{-\frac{1}{2}}{2 \cdot \frac{3}{4}} \end{bmatrix} = \begin{bmatrix} x - \frac{1}{2} \\ y + \frac{1}{3} \end{bmatrix}$$

giving,

$$\begin{aligned} \left(x - \frac{1}{2}\right)^2 + \frac{3}{4}\left(y + \frac{1}{3}\right)^2 + \left(x - \frac{1}{2}\right) - \frac{1}{2}\left(y + \frac{1}{3}\right) &= 0 \\ \left(x^2 - x + \frac{1}{4}\right) + \frac{3}{4}\left(y^2 + \frac{2}{3}y + \frac{1}{9}\right) + \left(x - \frac{1}{2}\right) - \left(\frac{1}{2}y + \frac{1}{6}\right) &= 0 \\ x^2 - x + \frac{1}{4} + \frac{3}{4}y^2 + \frac{1}{2}y + \frac{1}{12} + x - \frac{1}{2} - \frac{1}{2}y - \frac{1}{6} &= 0 \\ x^2 + \frac{3}{4}y^2 &= \frac{1}{3} \end{aligned}$$

and we can now see that the original equation $x^2 + xy + y^2 + x = 0$ describes an ellipse, a fact that may not be obvious from the original expression.

However, if $\alpha_{i,i} = 0$ for some i , then we cannot eliminate the term $\beta_i x_i$. That is, we can only eliminate linear terms if there is a corresponding quadratic term. Instead, we permute the coordinates such that $\alpha_{i,i} \neq 0$ for $1 \leq i \leq r$ and $\beta_i \neq 0$ for $r < i \leq r + s$.

If $s > 1$, then we don't change x_i for $1 \leq i \leq r$, but replace $\sum_{i=1}^s \beta_{r+i} x_{r+i}$ by βx_{r+1} . To show that this transformation is orthogonal, suppose our orthonormal basis is $(e_i)_{i=1}^n$. Then, we can extend,

$$e_1, \dots, e_r, \frac{1}{\sqrt{\sum_{i=1}^s \beta_{r+i}^2}} \sum_{i=1}^s \beta_{r+i} e_{r+i}$$

to an orthonormal basis of our Euclidean space using the Gram-Schmidt process. Note that the $(r + 1)$ th vector of the basis is chosen such that our equation will have just the term $(\sqrt{\sum_{i=1}^s \beta_{r+i}^2}) x_{r+1}$, so the equation has at most one non-zero β_i ; either there are no linear terms at all, or there is just β_{r+1} .

Finally, if there is a linear term, then $\beta_{r+1} \neq 0$, and, by dividing through by a constant, we can force $\beta_{r+1} = -1$, then perform the translation,

$$x_{r+1} \mapsto x_{r+1} - \frac{\gamma}{\beta_{r+1}}$$

to eliminate the constant γ . If there is no linear term, then we again divide the equation through by a constant to force $\gamma = 0$ or $\gamma = -1$, and move it to the right side in the latter case.

We have proved:

Theorem 4.16. *Any second degree polynomial equation can be transformed through isometries of Euclidean space into an equation with one of the following forms:*

- $\sum_{i=1}^r \alpha_i x_i^2 = 0$
- $\sum_{i=1}^r \alpha_i x_i^2 = 1$
- $\sum_{i=1}^r \alpha_i x_i^2 - x_{r+1} = 0$

where $0 \leq r \leq n$ and $(\alpha_i)_{i=1}^r$ are non-zero constants, and in the third case, $r < n$.

The sets of solutions defined by the first two cases are called *central* quadrics, as they have central symmetry; that is, if a vector \mathbf{v} satisfies the equation, then so does $-\mathbf{v}$.

4.8 Singular Value Decomposition

In this section, we will study linear maps $T : V \rightarrow W$ between Euclidean spaces V and W . Again, we wish to find bases of V and W such that the matrix of T is as simple as possible.

From row and column operations, we know it is always possible to choose bases of V and W such that the matrix of T has Smith normal form,

$$\left[\begin{array}{c|c} \mathbf{I}_r & \mathbf{0} \\ \hline \mathbf{0} & \mathbf{0} \end{array} \right]$$

where r is the rank of T . However, while simple, this form isn't very useful as it does not respect the Euclidean structure of V and W . The problem now is to choose *orthonormal* bases of V and W such that the matrix of T has a simple form.

Theorem 4.17 (Singular Value Decomposition for Linear Maps). *Suppose $T : V \rightarrow W$ is a linear map of rank r between Euclidean spaces V and W . Then, there exist unique positive numbers $\gamma_1 \geq \gamma_2 \geq \dots \geq \gamma_r > 0$ called the singular values of T , and orthonormal bases of V and W such that the matrix of T with respect to these bases is*

$$\Sigma = \left[\begin{array}{c|c} \mathbf{D} & \mathbf{0} \\ \hline \mathbf{0} & \mathbf{0} \end{array} \right]$$

where $\mathbf{D} = \text{diag}(\gamma_1, \dots, \gamma_r)$.

Corollary 4.17.1 (Singular Value Decomposition for Matrices). *Given any real $m \times n$ matrix \mathbf{A} of rank $r \leq \min\{m, n\}$, there exist unique singular values $\gamma_1 \geq \gamma_2 \geq \dots \geq \gamma_r > 0$ and (non-unique) orthogonal matrices \mathbf{P} and \mathbf{Q} such that,*

$$\Sigma = \left[\begin{array}{c|c} \mathbf{D} & \mathbf{0} \\ \hline \mathbf{0} & \mathbf{0} \end{array} \right] = \mathbf{P}^\top \mathbf{A} \mathbf{Q}$$

Equivalently, we say that the SVD of \mathbf{A} is,

$$\mathbf{A} = \mathbf{P} \Sigma \mathbf{Q}^\top$$

Theorem 4.18. *The matrices \mathbf{A} and \mathbf{A}^\top share the same singular values.*

Proof.

$$\begin{aligned} \mathbf{A} &= \mathbf{P} \Sigma \mathbf{Q}^\top \\ \mathbf{A}^\top &= (\mathbf{P} \Sigma \mathbf{Q}^\top)^\top \\ &= \mathbf{Q}^\top \Sigma^\top \mathbf{P} \end{aligned}$$

■

We present two algorithms for computing the singular value decomposition of a real $m \times n$ matrix \mathbf{A} .

Algorithm 7 Singular Value Decomposition

- 1: Compute the matrices $\mathbf{A} \mathbf{A}^\top$ and $\mathbf{A}^\top \mathbf{A}$.
 - 2: The eigenvectors of $\mathbf{A}^\top \mathbf{A}$ form the columns of the $n \times n$ orthogonal matrix \mathbf{Q} , and the eigenvectors of $\mathbf{A} \mathbf{A}^\top$ form the columns of the $m \times m$ orthogonal matrix \mathbf{P} .
 - 3: Perform the Gram-Schmidt process on the columns of \mathbf{Q} and \mathbf{P} if necessary.
 - 4: The square roots of the eigenvalues of either matrix form the singular values.
-

This requires finding eigenvectors for a pair of matrices, namely, $\mathbf{A}\mathbf{A}^\top$ and $\mathbf{A}^\top\mathbf{A}$. But, because $\mathbf{A} = \mathbf{P}\Sigma\mathbf{Q}$, once we have one of \mathbf{P} or \mathbf{Q} , it is possible to compute the other by multiplying the columns of the matrix we have by \mathbf{A} .

Algorithm 8 Singular Value Decomposition Shortcut

- 1: Let A be an $m \times n$ matrix.
- 2: Compute whichever of $\mathbf{A}\mathbf{A}^\top$ and $\mathbf{A}^\top\mathbf{A}$ has higher dimensions (if \mathbf{A} is a “tall” matrix, compute the former; if \mathbf{A} is a “wide” matrix, compute the latter).
- 3: Order the (orthonormalised) eigenvectors $(\mathbf{q}_i)_{i=1}^n$ of this matrix such that the corresponding eigenvalues are in decreasing order, $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_r > 0$, and $\lambda_i = 0$ for $i > r$, where r is the rank of \mathbf{A} .
- 4: The square roots of the eigenvalues form the singular values.
- 5: The $n \times n$ orthogonal matrix \mathbf{Q} is given by

$$\mathbf{Q} = [\mathbf{q}_1 | \mathbf{q}_2 | \dots | \mathbf{q}_n]$$

- 6: Define the sequence of vectors $(\mathbf{p}_i)_{i=1}^r$ by multiplying the corresponding \mathbf{q}_i by \mathbf{A} , then normalising:

$$\mathbf{p}_i = \frac{1}{\|\mathbf{A}\mathbf{q}_i\|} \mathbf{A}\mathbf{q}_i$$

for $1 \leq i \leq r$. Then, $(\mathbf{p}_i)_{i=1}^r$ is an orthonormal set. Using Gram-Schmidt or otherwise, extend this set to an orthonormal basis $(\mathbf{p}_i)_{i=1}^m$ of \mathbb{R}^m .

- 7: The $m \times m$ orthogonal matrix \mathbf{P} is given by

$$\mathbf{P} = [\mathbf{p}_1 | \mathbf{p}_2 | \dots | \mathbf{p}_m]$$

Example. Find a SVD decomposition of,

$$\mathbf{A} = \begin{bmatrix} 1 & 0 & 1 \\ -1 & 1 & 0 \end{bmatrix}$$

We compute the eigenvectors of $\mathbf{A}\mathbf{A}^\top$:

$$\mathbf{A}^\top\mathbf{A} = \begin{bmatrix} 2 & -1 & 1 \\ -1 & 2 & 0 \\ 1 & 0 & 1 \end{bmatrix}$$

$$c_{\mathbf{A}^\top\mathbf{A}}(x) = (x-3)(x-1)x$$

so we have eigenvalues (in descending order) $\lambda_1 = 3$, $\lambda_2 = 1$, and $\lambda_3 = 0$, giving singular values $\gamma_1 = \sqrt{\lambda_1} = \sqrt{3}$ and $\gamma_2 = \sqrt{\lambda_2} = 1$.

$$\Sigma = \begin{bmatrix} \sqrt{3} & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}$$

Next, we find the eigenvectors:

$$(\mathbf{A}^\top\mathbf{A} - \lambda_1\mathbf{I}_3)\mathbf{v} = \mathbf{0} \xrightarrow{\text{row reduce}} \begin{bmatrix} 1 & 0 & -2 \\ 0 & 1 & 1 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} a \\ b \\ c \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}$$

$$a = 2c$$

$$b = -c$$

$$\begin{bmatrix} a \\ b \\ c \end{bmatrix} = c \begin{bmatrix} 2 \\ -1 \\ 1 \end{bmatrix} \longrightarrow \ker(\mathbf{A} - \lambda_2 \mathbf{I}_3) = \text{span} \left(\underbrace{\begin{bmatrix} 2 \\ -1 \\ 1 \end{bmatrix}}_{\mathbf{v}_1} \right)$$

$$(\mathbf{A}^\top \mathbf{A} - \lambda_2 \mathbf{I}_3) \mathbf{v} = \mathbf{0} \xrightarrow{\text{row reduce}} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & -1 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} a \\ b \\ c \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}$$

$$b = -c$$

$$\begin{bmatrix} a \\ b \\ c \end{bmatrix} = c \begin{bmatrix} 0 \\ -1 \\ 1 \end{bmatrix} \longrightarrow \ker(\mathbf{A} - \lambda_2 \mathbf{I}_3) = \text{span} \left(\underbrace{\begin{bmatrix} 0 \\ -1 \\ 1 \end{bmatrix}}_{\mathbf{v}_2} \right)$$

$$(\mathbf{A}^\top \mathbf{A} - \lambda_3 \mathbf{I}_3) \mathbf{v} = \mathbf{0} \xrightarrow{\text{row reduce}} \begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} a \\ b \\ c \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}$$

$$a = -c$$

$$b = -c$$

$$\begin{bmatrix} a \\ b \\ c \end{bmatrix} = c \begin{bmatrix} -1 \\ -1 \\ 1 \end{bmatrix} \longrightarrow \ker(\mathbf{A} - \lambda_2 \mathbf{I}_3) = \text{span} \left(\underbrace{\begin{bmatrix} -1 \\ -1 \\ 1 \end{bmatrix}}_{\mathbf{v}_3} \right)$$

Normalise the eigenvectors,

$$\mathbf{q}_1 = \frac{1}{\sqrt{6}} \mathbf{v}_1, \quad \mathbf{q}_2 = \frac{1}{\sqrt{2}} \mathbf{v}_2, \quad \mathbf{q}_3 = \frac{1}{\sqrt{3}} \mathbf{v}_3$$

These eigenvectors have distinct eigenvalues, so they are orthogonal by Theorem 4.6, and we do not have to perform the Gram-Schmidt process. So, we have,

$$\mathbf{Q} = [\mathbf{q}_1 | \mathbf{q}_2 | \mathbf{q}_3]$$

$$= \begin{bmatrix} \frac{2}{\sqrt{6}} & 0 & -\frac{1}{\sqrt{3}} \\ -\frac{1}{\sqrt{6}} & \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{3}} \\ \frac{1}{\sqrt{6}} & \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{3}} \end{bmatrix}$$

To find \mathbb{P} , the first algorithm just repeats all of these calculations again on $\mathbf{A}\mathbf{A}^\top$. For the second algorithm, we calculate,

$$\mathbf{A}\mathbf{q}_1 = \frac{\sqrt{6}}{2} \begin{bmatrix} 1 \\ -1 \end{bmatrix}, \quad \mathbf{A}\mathbf{q}_2 = \frac{\sqrt{2}}{2} \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \quad \mathbf{A}\mathbf{q}_3 = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$

Ignore the zero vector and normalise $\mathbf{A}\mathbf{q}_1$ and $\mathbf{A}\mathbf{q}_2$:

$$\mathbf{p}_1 = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ -1 \end{bmatrix}, \quad \mathbf{p}_2 = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 1 \end{bmatrix}$$

Here, \mathbf{p}_1 and \mathbf{p}_2 already form a basis of \mathbb{R}^2 , so we skip the Gram-Schmidt process, and we have,

$$\begin{aligned}\mathbf{P} &= [\mathbf{p}_1 | \mathbf{p}_2] \\ &= \begin{bmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ -\frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{bmatrix}\end{aligned}$$

So the SVD of \mathbf{A} is given by

$$\begin{aligned}\mathbf{A} &= \mathbf{P}\mathbf{\Sigma}\mathbf{Q}^\top \\ &= \begin{bmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ -\frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{bmatrix} \begin{bmatrix} \sqrt{3} & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} \frac{2}{\sqrt{6}} & -\frac{1}{\sqrt{6}} & \frac{1}{\sqrt{6}} \\ 0 & \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ -\frac{1}{\sqrt{3}} & -\frac{1}{\sqrt{3}} & \frac{1}{\sqrt{3}} \end{bmatrix}\end{aligned}$$

5 Sesquilinear Forms

Not examinable.

6 Operators on Hilbert Spaces

Not examinable.

7 Finitely Generated Abelian Groups

So far, we have considered vector spaces over *fields*. In this section, we will consider a generalisation of a vector space called a *module* in which this field of scalars is replaced by a ring. Modules also generalise the notion of abelian groups, since the abelian groups are exactly the modules over the ring of integers.

7.1 Review

We quickly review some basic group theory. The definitions and theorems here will be stated in terms of abelian groups, but most will hold for general groups.

A *group*, $(G, *)$ is a set, G , equipped with a binary operation, $*$, that obeys the following axioms:

- $\forall a, b \in G, a * b \in G$ (closure);
- $\forall a, b, c \in G, a * (b * c) = (a * b) * c$ (associativity);
- $\exists e \in G$ such that $a * e = e * a = a \forall a \in G$ (existence of identity);
- $\forall a \in G, \exists (a^{-1}) \in G$ such that $a * (a^{-1}) = (a^{-1}) * a = e$ (existence of inverses).

Furthermore, if the operation is also commutative, the group is *abelian*. The identity, e , is also written as id_G or 0_G , the latter being used mainly for abelian groups.

For abelian groups, it is common to use additive notation where the binary operation is written as $+$, and we will continue with this notation from this point onwards.

A group G is *cyclic* if there exists an element $x \in G$ such that every element of G is of the form nx for some $n \in \mathbb{Z}$.

Let $(G, +)$ and $(H, *)$ be groups. A function $\phi : G \rightarrow H$ is a *group homomorphism* between G and H if

$$\phi(a + b) = \phi(a) * \phi(b)$$

for all $a, b \in G$. Note that this necessarily requires that,

$$\phi(\text{id}_G) = \text{id}_H$$

and

$$\phi(-a) = -\phi(a)$$

An injective homomorphism is called a *monomorphism* and a surjective homomorphism is called a *epimorphism*. If a homomorphism ϕ has an inverse, or equivalently, if ϕ is bijective, then ϕ is furthermore a *group isomorphism* and we write $G \cong H$ if such an isomorphism exists.

Theorem 7.1. *Every cyclic group is isomorphic to either $(\mathbb{Z}, +)$ or to $\mathbb{Z}_n = \mathbb{Z}/n\mathbb{Z}$ for some $n > 0$.*

The *order* of an element $g \in G$, denoted $|g|$, is the smallest natural n such that $ng = \text{id}_G$. If $ng \neq \text{id}_G$ for all $n \in \mathbb{N}$, then we say that g has *infinite order*.

Theorem 7.2. *If $\phi : G \rightarrow H$ is an isomorphism, then $|g| = |\phi(g)|$ for all $g \in G$.*

A group G is *generated* or *spanned* by a subset $X \subseteq G$ if every $g \in G$ can be written as a finite sum,

$$\sum_{i=1}^{|X|} m_i x_i$$

with $m_i \in \mathbb{Z}$ and $x_i \in X$, and we write $G = \langle X \rangle$, or $\langle x_1, \dots, x_n \rangle$ if X has finite cardinality. In this latter case, we say that G is *finitely generated*.

In multiplicative notation, $G = \langle X \rangle$ if and only if

$$\prod_{i=1}^{|X|} x_i^{m_i}$$

so every element can be “factored” into elements of X .

A group is cyclic if and only if X is a singleton set.

The *direct sum* of a set of abelian groups $(G_i)_{i=1}^n$ is defined to be the set,

$$\{(g_i)_{i=1}^n : g_i \in G_i\}$$

with component-wise addition

$$(g_i)_{i=1}^n + (h_i)_{i=1}^n = (g_i + h_i)_{i=1}^n$$

This forms a group with identity $(\text{id}_{G_i})_{i=1}^n$ and $-(g_i)_{i=1}^n = (-g_i)_{i=1}^n$.

For non-abelian groups, this is more commonly called the *direct product* of groups.

A subset $H \subseteq G$ of a group $(G, +)$ is a *subgroup* of G if $(H, +)$ is also a group, and we write $H \leq G$ to denote this relation. The group itself, G , and the trivial group $\{\text{id}_G\}$ are always subgroups of G . Any subgroup H not equal to G is a *proper* subgroup, and we write $H < G$ to denote this relation. Any subgroup not equal to $\{\text{id}_G\}$ is a *non-trivial* subgroup.

Lemma 7.3. *If $H \leq G$, then $\text{id}_H = \text{id}_G$.*

Theorem 7.4. *Let $H \subseteq G$. Then, the following statements are equivalent:*

- (i) $H \leq G$;
- (ii) (a) $H \neq \emptyset$;

$$(b) a, b \in H \rightarrow a + b \in H;$$

$$(c) a \in H \rightarrow -a \in H.$$

$$(iii) (a) H \neq \emptyset;$$

$$(b) a, b \in H \rightarrow a - b \in H.$$

(ii) and (iii) are the *two step* and *one step* subgroup tests (so called because H is often assumed to be non-empty, and hence doesn't count as a step).

Let G be a group, $H \leq G$ and $g \in G$. The set $g + H = \{g + h : h \in H\}$ is a *left coset* of H in G , and $H + g = \{h + g : h \in H\}$ is a *right coset* of H in G . For abelian groups, left and right cosets coincide, and we just say *coset* alone.

Theorem 7.5. *The following statements are equivalent for any $x, g \in G$:*

- $x \in H + g$;
- $H + g = H + x$;
- $x - g \in H$.

Corollary 7.5.1. *Two cosets $H + a$ and $H + b$ are either equal or disjoint.*

Corollary 7.5.2. *The cosets of H in G partition G .*

Theorem 7.6. *If H is finite, then all cosets of H in G have $|H|$ elements.*

The number of distinct left (or right) cosets of H in G is called the *index* of H in G , and is written as $[G : H]$

Theorem 7.7 (Lagrange's Theorem). *If $H \leq G$, then,*

$$|G| = [G : H]|H|$$

Corollary 7.7.1. *If $H \leq G$, then the order of H divides the order of G .*

Corollary 7.7.2. *For any $g \in G$, $|g|$ divides $|G|$.*

Theorem 7.8. *Let G be a group of prime order p . Then, G is cyclic and $G \cong \mathbb{Z}_p$.*

If A and B are subsets of a group G , then we define their *sum* by,

$$A + B := \{a + b : a \in A, b \in B\}$$

Lemma 7.9. *If H is a subgroup of an abelian group G , and $H + a, H + b$ are cosets of H in G , then,*

$$(H + g) + (H + k) = H + (g + k)$$

Theorem 7.10. *Let H be a subgroup of an abelian group G . Then, the set G/H of cosets $H + g$ of H in G forms a group under addition of sets as defined above.*

Such a group is called a *quotient group* or *factor group* of G by H . Note that if G is finite, then $|G/H| = [G : H] = |G|/|H|$.

Let $\phi : G \rightarrow H$ be a group homomorphism. Then, the *kernel* $\ker(\phi)$ of ϕ is defined to be the set of elements of G mapped to the identity id_H . That is,

$$\ker(\phi) = \{g \in G : \phi(g) = \text{id}_H\}$$

Note that $\ker(\phi)$ always contains id_H as group homomorphisms must map identities to identities. If id_H is the only element of $\ker(\phi)$, then the kernel is *trivial*.

The *image* of ϕ is then defined as,

$$\text{im}(\phi) = \{\phi(g) : g \in G\}$$

Theorem 7.11. *Let $\phi : G \rightarrow H$ be a group homomorphism. Then, ϕ is a monomorphism (injection) if and only if the kernel $\ker(\phi)$ is trivial.*

Proof. Since $\text{id}_G \in \ker(\phi)$, if ϕ is injective, then we must have $\ker(\phi) = \{\text{id}_G\}$, completing the forward implication.

Conversely, suppose $\ker(\phi) = \{\text{id}_G\}$, and let $a, b \in G$ with $\phi(a) = \phi(b)$. Then, $\phi(a - b) = \phi(a) - \phi(b) = \text{id}_H$, so $a - b \in \ker(\phi)$. But then, $a - b = \text{id}_G$, and hence $a = b$, so ϕ is injective, completing the reverse implication. ■

Theorem 7.12. *Let $\phi : G \rightarrow H$ be a group homomorphism. Then, $\ker(\phi)$ is a subgroup of G , and $\text{im}(\phi)$ is a subgroup of H .*

Furthermore, if K is a subgroup of G , then the map $\phi : G \rightarrow G/K$ defined by $\phi(g) = K + g$ is an epimorphism (surjection) with kernel K .

Proof. The first statement follows from the two step subgroup test. For the second, it is clear that ϕ is surjective, and $\phi(g) = \text{id}_{G/K} \leftrightarrow K + g = K + \text{id}_G \leftrightarrow g \in K$, so $\ker(\phi) = K$. ■

Theorem 7.13 (First Isomorphism Theorem). *Let $\phi : G \rightarrow H$ be a group homomorphism with kernel K . Then, $G/K \cong \text{im}(\phi)$. More precisely, there is an isomorphism $\bar{\phi} : G/K \rightarrow \text{im}(\phi)$ defined by $\bar{\phi}(K + g) = \phi(g)$ for $g \in G$.*

7.2 Free Abelian Groups

The direct sum $\mathbb{Z}^n := \underbrace{\mathbb{Z} \oplus \mathbb{Z} \oplus \dots \oplus \mathbb{Z}}_n$ of n copies of \mathbb{Z} is called a (finitely generated) *free abelian group* of rank n .

More generally, a finitely generated abelian group is *free abelian* if it is isomorphic to \mathbb{Z}^n for some $n \geq 0$, with the free abelian group \mathbb{Z}^0 of rank 0 defined to be the trivial group.

The free abelian groups have many properties in common with vector spaces like \mathbb{R}^n , but we would expect some differences, as \mathbb{Z} is not a field. Similarly to vector spaces, we will write elements of \mathbb{Z}^n as column vectors.

We then define the *standard basis* of \mathbb{Z}^n exactly as for the vector space \mathbb{R}^n ; that is, a set of vectors $(\mathbf{x}_i)_{i=1}^n$ such that \mathbf{x}_i is 0 everywhere apart from the i th component, which has a 1. This basis has the same properties as a basis of a vector space: the vectors are linearly independent, and they span (generate) \mathbb{Z}^n , for a modified definition of linear independence and span.

Elements $(x_i)_{i=1}^n$ of an abelian group G are called *linearly independent* if the equation

$$\sum_{i=1}^n \alpha_i x_i = \text{id}_G$$

with integer coefficients $\alpha_i \in \mathbb{Z}$ holds only if $\alpha_i = 0$ for all $1 \leq i \leq n$.

Elements $S = \{x_i\}_{i=1}^n$ of an abelian group G form a *free basis* or *integral basis* of G if and only if they are linearly independent and span (generate) G . That is,

$$G = \langle (x_i)_{i=1}^n \rangle$$

or equivalently, every $g \in G$ can be written as a *unique* linear integer combination of elements in S :

$$g = \sum_{i=1}^n \alpha_i x_i$$

where all the $\alpha_i \in \mathbb{Z}$.

Note that a set of elements in \mathbb{Z}^n that form a basis of \mathbb{Q}^n or \mathbb{R}^n need not be a free basis of \mathbb{Z}^n . For instance, the set,

$$\left\{ \begin{bmatrix} 2 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 2 \end{bmatrix} \right\}$$

is a basis of \mathbb{Q}^2 and \mathbb{R}^2 , and are linearly independent in \mathbb{Z}^2 , but not of \mathbb{Z}^2 , as we are only allowed integer coefficients in \mathbb{Z}^n : there is no way to write, say,

$$\begin{bmatrix} 1 \\ 0 \end{bmatrix} = \alpha_1 \begin{bmatrix} 2 \\ 0 \end{bmatrix} + \alpha_2 \begin{bmatrix} 0 \\ 2 \end{bmatrix}$$

with α_1 and α_2 as integers. This also shows that a set of n linearly independent elements of \mathbb{Z}^n does not necessarily form a free basis.

Theorem 7.14. *For any set of elements $(g_i)_{i=1}^n$ of an abelian group G , it is possible to extend the assignment $\mathbf{x}_i \mapsto g_i$ to a group homomorphism $\phi: \mathbb{Z}^n \rightarrow G$. We define*

$$\phi \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \vdots \\ \alpha_n \end{bmatrix} := \sum_{i=1}^n \alpha_i g_i$$

Then,

- ϕ is a group homomorphism;
- The set (g_i) is linearly independent if and only if ϕ is an monomorphism (injection);
- The set (g_i) span G if and only if ϕ is an epimorphism (surjection);
- The set (g_i) form a free basis of G if and only if ϕ is an isomorphism.

Theorem 7.15 (Universal Property of the Free Abelian Group). *Let G be a free abelian group with a free basis $(g_i)_{i=1}^n$. Let H be an abelian group, and $(h_i)_{i=1}^n$ be elements of H . Then, there exists a unique group homomorphism $\phi: G \rightarrow H$ such that $\phi(g_i) = h_i$.*

As for finite dimensional vector spaces, we have yet to prove that any two free bases of a free abelian group have the same size. Let $(\mathbf{x}_i)_{i=1}^n$ be the standard free basis of \mathbb{Z}^n , and let $(\mathbf{y}_i)_{i=1}^m$ be another free basis of \mathbb{Z}^n (expressed in terms of the standard basis). As in linear algebra, we define the associated change of basis matrix \mathbf{P} with respect to the original basis (\mathbf{x}_i) and target basis (\mathbf{y}_i) by,

$$\mathbf{P} = [\mathbf{y}_1 | \mathbf{y}_2 | \dots | \mathbf{y}_m]$$

That is, if \mathbf{x} and \mathbf{y} are column vectors expressed in terms of the standard basis $(\mathbf{x}_i)_{i=1}^n$ and free basis $(\mathbf{y}_i)_{i=1}^m$, respectively, then,

$$\mathbf{x} = \mathbf{P}\mathbf{y}$$

Theorem 7.16. *Let $(\mathbf{y}_i)_{i=1}^m \subset \mathbb{Z}^n$. Then, the following statements are equivalent:*

- $(\mathbf{y}_i)_{i=1}^m$ is a free basis of \mathbb{Z}^n ;

- $n = m$ and the change of basis matrix $\mathbf{P} \in \mathbb{Z}^{n \times n}$ has an inverse $\mathbf{P}^{-1} \in \mathbb{Z}^{n \times n}$ (that is, the inverse of \mathbf{P} has integer entries);
- $n = m$ and $\det(\mathbf{P}) = \pm 1$.

A square matrix with integer entries and determinant ± 1 is called *unimodular*.

For example, if $n = 1$, and we have elements,

$$\mathbf{y}_1 = \begin{bmatrix} 2 \\ 7 \end{bmatrix}, \quad \mathbf{y}_2 = \begin{bmatrix} 1 \\ 4 \end{bmatrix}$$

then we have

$$\mathbf{P} = \begin{bmatrix} 2 & 1 \\ 7 & 4 \end{bmatrix}$$

with $\det(\mathbf{P}) = 2 \cdot 4 - 1 \cdot 7 = 1$, so $\{\mathbf{y}_1, \mathbf{y}_2\}$ is a free basis of \mathbb{Z}^2 .

In contrast, take our previous example of

$$\mathbf{y}_1 = \begin{bmatrix} 2 \\ 0 \end{bmatrix}, \quad \mathbf{y}_2 = \begin{bmatrix} 0 \\ 2 \end{bmatrix}$$

Then, $\det(\mathbf{P}) = 4 \neq \pm 1$, so this set is not a free basis of \mathbb{Z}^2 .

7.3 Unimodular Smith Normal Form

Recall that, in linear algebra, we may use elementary row and column operations to reduce an $m \times n$ matrix \mathbf{A} of rank r to a matrix,

$$\left[\begin{array}{c|c} \mathbf{I}_r & \mathbf{0} \\ \hline \mathbf{0} & \mathbf{0} \end{array} \right]$$

called the *Smith normal form* of \mathbf{A} .

For matrices over \mathbb{Z} , we can similarly reduce a matrix to a Smith normal form, but now, the non-zero entries will not necessarily be equal to 1.

For matrices over \mathbb{Z} , we use *unimodular* row and column operations instead:

- (UC1) Replace a column \mathbf{c}_i by $\mathbf{c}_i + \lambda \mathbf{c}_j$, $\lambda \in \mathbb{Z}$.
- (UC2) Interchange two columns \mathbf{c}_i and \mathbf{c}_j .
- (UC3) Replace a column \mathbf{c}_i with $-\mathbf{c}_i$.
- (UR1) Replace a row \mathbf{r}_i by $\mathbf{r}_i + \lambda \mathbf{r}_j$, $\lambda \in \mathbb{Z}$.
- (UR2) Interchange two rows \mathbf{r}_i and \mathbf{r}_j .
- (UR3) Replace a row \mathbf{r}_i with $-\mathbf{r}_i$.

Elementary row and column operations on a matrix \mathbf{A} correspond to multiplying \mathbf{A} on the left or right, respectively, by an elementary matrix. These matrices have determinant ± 1 , and are hence also unimodular matrices. From this, unimodular row and column operations correspond to the following change of bases, where $(\mathbf{e}_i)_{i=1}^n$ is a free basis for \mathbb{Z}^n (the domain of the linear map \mathbf{A} represents) and $(\mathbf{f}_i)_{i=1}^m$ is a free basis of \mathbb{Z}^m (the codomain).

- (UC1) $\mathbf{e}_i \mapsto \mathbf{e}_i + \lambda \mathbf{e}_j$;
- (UC2) $\mathbf{e}_i \leftrightarrow \mathbf{e}_j$;

(UC3) $\mathbf{e}_i \mapsto -\mathbf{e}_i$;

(UR1) $\mathbf{f}_j \mapsto \mathbf{f}_j - \lambda \mathbf{f}_i$ (note the sign change and the reversal of indices);

(UR2) $\mathbf{f}_i \leftrightarrow \mathbf{f}_j$;

(UR3) $\mathbf{f}_i \mapsto -\mathbf{f}_i$;

Theorem 7.17. *Let \mathbf{A} be an $m \times n$ matrix over \mathbb{Z} with rank r . Then, by using a sequence of unimodular elementary row and column operations, we can reduce \mathbf{A} to a matrix*

$$\mathbf{S} = \left[\begin{array}{c|c} \mathbf{D} & \mathbf{0} \\ \hline \mathbf{0} & \mathbf{0} \end{array} \right]$$

with $\mathbf{D} = \text{diag}(d_1, d_2, \dots, d_r)$, where $(d_i)_{i=1}^r$ are positive integers satisfying $d_i | d_{i+1}$ for $1 \leq i < r$.

Subject to these conditions, the d_i are uniquely determined by the matrix \mathbf{A} .

The matrix \mathbf{S} is then called the *unimodular Smith normal form* of \mathbf{A} , or the \mathbb{Z} SNF of \mathbf{A} .

Lemma 7.18. *Let $\mathbf{A} \in \mathbb{Z}^{m \times n}$ have unimodular Smith normal form \mathbf{S} with r non-zero diagonal entries $(d_i)_{i=1}^r$. Then, the greatest common divisor of all the entries of \mathbf{A} is d_1 .*

Algorithm 9 Smith Normal Form Decomposition

- 1: Compute the greatest common divisor x of the entries in the first column, and the greatest common divisor y of the entries in the first row. Without loss of generality, suppose $x < y$ (simply reverse row and columns in the following, if otherwise).
 - 2: Using the Euclidean algorithm, we can form a row whose first element is x .
 - 3: Move this row to the first row such that x is in the first entry. This is the *pivot* element (and row/column).
 - 4: Subtract multiples of the pivot row from every other row until the pivot column has 0s everywhere else (possible since x divides everything in this column).
 - 5: Repeat this process for the rows. This will likely undo some of our work on the pivot column, but just repeat this process again. This process is guaranteed to terminate since the greatest common divisor is reduced with each iteration.
 - 6: Eventually, the pivot column and row will be zero every outside the pivot element. Iterate this process on each column/row until the matrix is diagonal.
 - 7: The diagonal entries may not satisfy the divisibility requirements of the unimodular Smith normal form, so we again use the Euclidean algorithm to obtain the greatest common divisor of the diagonal elements, then add or subtract this value until the divisibility requirements are met.
-

Rows and columns can also be interchanged, if it makes the Bézout coefficients smaller or if the divisor is easier to obtain. If it is easy to spot or calculate the greatest common divisor of all the entries of the matrix, we can shortcut the first few steps somewhat.

Example. Find a \mathbb{Z} SNF decomposition of,

$$\mathbf{A} = \begin{bmatrix} -18 & -18 & -18 & 90 \\ 54 & 12 & 45 & 48 \\ 9 & -6 & 6 & 63 \\ 18 & 6 & 15 & 12 \end{bmatrix}$$

It is easy to see that every entry of \mathbf{A} is divisible by 3, but for the sake of illustration, we will not use this shortcut.

Instead, $\gcd(-18, 54, 9, 18) = 9$ and $\gcd(-18, -18, -18, 90) = 9$, so we can choose to work on columns or rows. We already have a 9 in the first column, we will work on the first column:

$$\begin{bmatrix} -18 & -18 & -18 & 90 \\ 54 & 12 & 45 & 48 \\ 9 & -6 & 6 & 63 \\ 18 & 6 & 15 & 12 \end{bmatrix} \xrightarrow{\mathbf{r}_3 \leftrightarrow \mathbf{r}_1} \begin{bmatrix} 9 & -6 & 6 & 63 \\ 54 & 12 & 45 & 48 \\ -18 & -18 & -18 & 90 \\ 18 & 6 & 15 & 12 \end{bmatrix}$$

Now, clear out the rest of the pivot column,

$$\begin{bmatrix} 9 & -6 & 6 & 63 \\ 54 & 12 & 45 & 48 \\ -18 & -18 & -18 & 90 \\ 18 & 6 & 15 & 12 \end{bmatrix} \xrightarrow{\begin{array}{l} \mathbf{r}_2 \mapsto \mathbf{r}_2 - 6\mathbf{r}_1 \\ \mathbf{r}_3 \mapsto \mathbf{r}_3 + 2\mathbf{r}_1 \\ \mathbf{r}_4 \mapsto \mathbf{r}_4 - 2\mathbf{r}_1 \end{array}} \begin{bmatrix} 9 & -6 & 6 & 63 \\ 0 & 48 & 9 & -330 \\ 0 & -30 & -6 & 216 \\ 0 & 18 & 3 & -114 \end{bmatrix}$$

$\gcd(9, -6, 6, 63) = 3$, so, we make a 3 in the pivot row.

$$\begin{bmatrix} 9 & -6 & 6 & 63 \\ 0 & 48 & 9 & -330 \\ 0 & -30 & -6 & 216 \\ 0 & 18 & 3 & -114 \end{bmatrix} \xrightarrow{\mathbf{c}_1 \mapsto \mathbf{c}_1 - \mathbf{c}_3} \begin{bmatrix} 3 & -6 & 6 & 63 \\ -9 & 48 & 9 & -330 \\ 6 & -30 & -6 & 216 \\ -3 & 18 & 3 & -114 \end{bmatrix}$$

Clear the pivot row,

$$\begin{bmatrix} 3 & -6 & 6 & 63 \\ -9 & 48 & 9 & -330 \\ 6 & -30 & -6 & 216 \\ -3 & 18 & 3 & -114 \end{bmatrix} \xrightarrow{\begin{array}{l} \mathbf{c}_2 \mapsto \mathbf{c}_2 + 2\mathbf{c}_1 \\ \mathbf{c}_3 \mapsto \mathbf{c}_3 - 2\mathbf{c}_1 \\ \mathbf{c}_4 \mapsto \mathbf{c}_4 - 21\mathbf{c}_1 \end{array}} \begin{bmatrix} 3 & 0 & 0 & 0 \\ -9 & 30 & 27 & -141 \\ 6 & -18 & -18 & 90 \\ -3 & 12 & 9 & -51 \end{bmatrix}$$

Clear the pivot column again,

$$\begin{bmatrix} 3 & 0 & 0 & 0 \\ -9 & 30 & 27 & -141 \\ 6 & -18 & -18 & 90 \\ -3 & 12 & 9 & -51 \end{bmatrix} \xrightarrow{\begin{array}{l} \mathbf{r}_2 \mapsto \mathbf{r}_2 + 3\mathbf{r}_1 \\ \mathbf{r}_3 \mapsto \mathbf{r}_3 - 2\mathbf{r}_1 \\ \mathbf{r}_4 \mapsto \mathbf{r}_4 + \mathbf{r}_1 \end{array}} \begin{bmatrix} 3 & 0 & 0 & 0 \\ 0 & 30 & 27 & -141 \\ 0 & -18 & -18 & 90 \\ 0 & 12 & 9 & -51 \end{bmatrix}$$

$\gcd(30, -18, 12) = 6$, and $\gcd(30, 27, -141) = 3$, so we will work on the row first this time.

$$\begin{bmatrix} 3 & 0 & 0 & 0 \\ 0 & 30 & 27 & -141 \\ 0 & -18 & -18 & 90 \\ 0 & 12 & 9 & -51 \end{bmatrix} \xrightarrow{\mathbf{c}_2 \mapsto \mathbf{c}_2 - \mathbf{c}_3} \begin{bmatrix} 3 & 0 & 0 & 0 \\ 0 & 3 & 27 & -141 \\ 0 & 0 & -18 & 90 \\ 0 & 3 & 9 & -51 \end{bmatrix}$$

Clear the row out,

$$\begin{bmatrix} 3 & 0 & 0 & 0 \\ 0 & 3 & 27 & -141 \\ 0 & 0 & -18 & 90 \\ 0 & 3 & 9 & -51 \end{bmatrix} \xrightarrow{\begin{array}{l} \mathbf{c}_3 \mapsto \mathbf{c}_3 - 9\mathbf{c}_1 \\ \mathbf{c}_4 \mapsto \mathbf{c}_4 + 47\mathbf{c}_1 \end{array}} \begin{bmatrix} 3 & 0 & 0 & 0 \\ 0 & 3 & 0 & 0 \\ 0 & 0 & -18 & 90 \\ 0 & 3 & -18 & 90 \end{bmatrix}$$

$\gcd(3, 0, 3) = 3$, so our pivot is already the correct divisor. Clear the rest of the pivot column,

$$\begin{bmatrix} 3 & 0 & 0 & 0 \\ 0 & 3 & 0 & 0 \\ 0 & 0 & -18 & 90 \\ 0 & 3 & -18 & 90 \end{bmatrix} \xrightarrow{\mathbf{r}_4 \mapsto \mathbf{r}_4 - \mathbf{r}_1} \begin{bmatrix} 3 & 0 & 0 & 0 \\ 0 & 3 & 0 & 0 \\ 0 & 0 & -18 & 90 \\ 0 & 0 & -18 & 90 \end{bmatrix}$$

From this point, the algorithm would compute $\gcd(-18, -18) = 18$, and $\gcd(-18, 90) = 9$, so we would then work on the row, but at this point, the matrix is small enough that we can obviously just clear the last column, then row:

$$\begin{array}{ccc}
 \begin{bmatrix} 3 & 0 & 0 & 0 \\ 0 & 3 & 0 & 0 \\ 0 & 0 & -18 & 90 \\ 0 & 0 & -18 & 90 \end{bmatrix} & \xrightarrow{\mathbf{c}_4 \mapsto \mathbf{c}_4 - 5\mathbf{c}_3} & \begin{bmatrix} 3 & 0 & 0 & 0 \\ 0 & 3 & 0 & 0 \\ 0 & 0 & -18 & 0 \\ 0 & 0 & -18 & 0 \end{bmatrix} \\
 \begin{bmatrix} 3 & 0 & 0 & 0 \\ 0 & 3 & 0 & 0 \\ 0 & 0 & -18 & 0 \\ 0 & 0 & -18 & 0 \end{bmatrix} & \xrightarrow{\mathbf{r}_4 \mapsto \mathbf{r}_4 - \mathbf{r}_3} & \begin{bmatrix} 3 & 0 & 0 & 0 \\ 0 & 3 & 0 & 0 \\ 0 & 0 & -18 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \\
 \begin{bmatrix} 3 & 0 & 0 & 0 \\ 0 & 3 & 0 & 0 \\ 0 & 0 & -18 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} & \xrightarrow{\mathbf{r}_3 \mapsto -\mathbf{r}_3} & \begin{bmatrix} 3 & 0 & 0 & 0 \\ 0 & 3 & 0 & 0 \\ 0 & 0 & 18 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}
 \end{array}$$

In this case, we were lucky in that, the diagonal entries already satisfy the divisibility requirements, so we are done.

7.4 Subgroups of Free Abelian Groups

Theorem 7.19. *Any subgroup of a finitely generated abelian group is finitely generated.*

Let H be a subgroup of the free abelian group \mathbb{Z}^n , and suppose that $H = \langle \mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_m \rangle$. Then, H can be represented by a $n \times m$ matrix \mathbf{A} defined by

$$\mathbf{A} = [\mathbf{v}_1 | \mathbf{v}_2 | \dots | \mathbf{v}_m]$$

For example, if $n = 3$ and H is generated by,

$$\mathbf{v}_1 = \begin{bmatrix} 1 \\ 3 \\ -1 \end{bmatrix}, \quad \mathbf{v}_2 = \begin{bmatrix} 2 \\ 0 \\ 1 \end{bmatrix}$$

then,

$$\mathbf{A} = \begin{bmatrix} 1 & 2 \\ 3 & 0 \\ -1 & 1 \end{bmatrix}$$

If a different free basis $(\mathbf{y}_i)_{i=1}^n$ of \mathbb{Z}^n with change of basis matrix \mathbf{P} is used, then each column \mathbf{v}_i of \mathbf{A} is replaced by $\mathbf{P}^{-1}\mathbf{v}_i$, and hence \mathbf{A} itself is replaced by $\mathbf{P}^{-1}\mathbf{A}$.

For example, if we have the basis

$$\mathbf{y}_1 = \begin{bmatrix} 0 \\ -1 \\ 0 \end{bmatrix}, \quad \mathbf{y}_2 = \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}, \quad \mathbf{y}_3 = \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}$$

of \mathbb{Z}^3 , then,

$$\mathbf{P} = \begin{bmatrix} 0 & 1 & 1 \\ -1 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix}, \quad \mathbf{P}^{-1} = \begin{bmatrix} 1 & -1 & -1 \\ 0 & 0 & 1 \\ 1 & 0 & -1 \end{bmatrix}, \quad \mathbf{P}^{-1}\mathbf{A} = \begin{bmatrix} -1 & 1 \\ -1 & 1 \\ 2 & 1 \end{bmatrix}$$

Theorem 7.20. *Suppose that a subgroup H of \mathbb{Z}^n is represented by the matrix $\mathbf{A} \in \mathbb{Z}^{n \times m}$. Then, if the matrix $\mathbf{B} \in \mathbb{Z}^{n \times m}$ can be obtained by applying unimodular row and column operations on \mathbf{A} , then \mathbf{B} represents the same subgroup H of \mathbb{Z}^n using a (possibly) different free basis of \mathbb{Z}^n .*

In particular, we can transform \mathbf{A} to its unimodular Smith normal form, \mathbf{S} . So, if \mathbf{S} represents the subgroup H with free basis $(\mathbf{y}_i)_{i=1}^r$ of \mathbb{Z}^n , then the r non-zero columns of \mathbf{S} correspond to the elements $(d_i \mathbf{y}_i)_{i=1}^r$. So,

Theorem 7.21. *Let H be a subgroup of \mathbb{Z}^n . Then, there exists a free basis $(\mathbf{y}_i)_{i=1}^r$ of \mathbb{Z}^n such that $H = \langle (d_i \mathbf{y}_i)_{i=1}^r \rangle$, where each $d_i > 0$ and $d_i | d_{i+1}$ for $1 \leq i < r$.*

By keeping track of the row operations used, we can find the free basis of \mathbb{Z}^n such that the matrix of H has a simple form by applying the operations to the basis vectors.

7.5 General Finitely Generated Abelian Groups

A *presentation* is one method of specifying a group. A presentation of a group G is composed of a set S of generators, and set R of relations among those generators. We then say that G has presentation

$$\langle S \mid R \rangle$$

For instance, a cyclic group of order n has the presentation,

$$\langle a \mid a^n = \text{id} \rangle$$

This is also sometimes written as

$$\langle a \mid a^n \rangle$$

under the convention that any terms without an equals is assumed to be equal to the identity element. For instance, the dihedral group D_n has presentation,

$$\langle r, f \mid r^n, f^2, (rf)^2 \rangle$$

where r is a rotation and f a reflection; \mathbb{Z}^2 has presentation,

$$\langle x, y \mid xy = yx \rangle$$

and the free group $F(S)$ on a set S has presentation,

$$\langle S \mid \emptyset \rangle$$

Note that the presentation of a group is not unique.

Let G be a finitely generated abelian group. If G has n generators $(x_i)_{i=1}^n$, then Theorem 7.14 gives us a way to define a surjective homomorphism $\phi : \mathbb{Z}^n \rightarrow G$, and by the first isomorphism theorem, we can deduce that $G \cong \mathbb{Z}^n / K$, where $K = \ker(\phi)$, so we have proved that every finitely generated abelian group is isomorphic to a quotient group of a free abelian group.

From the definition of ϕ , we see that K is given by,

$$\begin{aligned} K &= \{ \mathbf{v} \in \mathbb{Z}^n : \phi(\mathbf{v}) = \text{id}_G \} \\ &= \left\{ [v_1, \dots, v_n]^\top \in \mathbb{Z}^n : \sum_{i=1}^n v_i x_i = \text{id}_G \right\} \end{aligned}$$

and, because K is a subgroup of G , which is finitely generated, K is also finitely generated by elements $(\mathbf{v}_i)_{i=1}^m$ of \mathbb{Z}^n . The quotient group \mathbb{Z}^n / K then has presentation,

$$\langle \mathbf{x}_1, \dots, \mathbf{x}_n \mid \mathbf{v}_1, \dots, \mathbf{v}_m \rangle$$

(where $(\mathbf{x}_i)_{i=1}^n$ is the standard basis of \mathbb{Z}^n) and as before, this group is isomorphic to G ,

$$G \cong \langle \mathbf{x}_1, \dots, \mathbf{x}_n \mid \mathbf{v}_1, \dots, \mathbf{v}_m \rangle$$

Now, we can find the unimodular Smith normal form of the matrix of K to find a free basis $(\mathbf{y}_i)_{i=1}^n$ of \mathbb{Z}^n such that $K = \langle (d_i \mathbf{y}_i)_{i=1}^r \rangle$ for some $r \leq n$ and $d_i > 0$ and $d_i \mid d_{i+1}$ for $1 \leq i < r$, giving,

$$G \cong \langle \mathbf{y}_1, \dots, \mathbf{y}_n \mid d_1 \mathbf{y}_1, \dots, d_r \mathbf{y}_r \rangle$$

Theorem 7.22. *The group,*

$$\langle \mathbf{y}_1, \dots, \mathbf{y}_n \mid d_1 \mathbf{y}_1, \dots, d_r \mathbf{y}_r \rangle$$

is isomorphic to the direct sum of cyclic groups,

$$\mathbb{Z}_{d_1} \oplus \mathbb{Z}_{d_2} \oplus \dots \oplus \mathbb{Z}_{d_r} \oplus \mathbb{Z}^{n-r}$$

Putting these results together, we obtain,

Theorem 7.23 (Fundamental Theorem of Finitely Generated Abelian Groups). *If G is a finitely generated abelian group, then G is isomorphic to a direct sum of cyclic groups. More precisely, if G is generated by n elements, then, for some r with $0 \leq r \leq n$, there exist integers $(d_i)_{i=1}^r$ with $d_i > 0$ and $d_i \mid d_{i+1}$ for $1 \leq i < r$, such that,*

$$G \cong \mathbb{Z}_{d_1} \oplus \mathbb{Z}_{d_2} \oplus \dots \oplus \mathbb{Z}_{d_r} \oplus \mathbb{Z}^{n-r}$$

or more compactly,

$$\left(\bigoplus_{i=1}^r \mathbb{Z}_{d_i} \right) \oplus \mathbb{Z}^{n-r}$$

That is, G is isomorphic to the direct sum of r finite cyclic groups of orders d_1, \dots, d_r , and $n - r$ infinite cyclic groups.

There may be some factors $\mathbb{Z}_1 = \{\text{id}\}$, which can be omitted from the direct sum (unless it is the only factor and $G \cong \mathbb{Z}_1$ is trivial). It could be the case that $n - r = 0$, which occurs if and only if G is finite. We can also have that $d_i = 1$ for all i , which occurs if and only if G is free abelian.

Example. From before, we found that the matrix

$$\mathbf{A} = \begin{bmatrix} -18 & -18 & -18 & 90 \\ 54 & 12 & 45 & 48 \\ 9 & -6 & 6 & 63 \\ 18 & 6 & 15 & 12 \end{bmatrix}$$

has unimodular smith normal form

$$\begin{bmatrix} 3 & 0 & 0 & 0 \\ 0 & 3 & 0 & 0 \\ 0 & 0 & 18 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

This means that the group defined by \mathbf{A} , which has presentation,

$$\left\langle \mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, \mathbf{x}_4 \mid \begin{array}{l} -18\mathbf{x}_1 + 54\mathbf{x}_2 + 9\mathbf{x}_3 + 18\mathbf{x}_4, \quad -18\mathbf{x}_1 + 12\mathbf{x}_2 - 6\mathbf{x}_3 + 6\mathbf{x}_4, \\ -18\mathbf{x}_1 + 45\mathbf{x}_2 + 6\mathbf{x}_3 + 15\mathbf{x}_4, \quad 90\mathbf{x}_1 + 48\mathbf{x}_2 + 63\mathbf{x}_3 + 12\mathbf{x}_4 \end{array} \right\rangle$$

is isomorphic to,

$$\mathbb{Z}_3 \oplus \mathbb{Z}_3 \oplus \mathbb{Z}_{18} \oplus \mathbb{Z}^1$$

which has a maximal finite subgroup of order $3 \times 3 \times 18 = 162$.

7.6 Finite Abelian Groups

For any finite abelian group G , we now have $G \cong \bigoplus_{i=1}^r \mathbb{Z}_{d_i}$, where $d_i | d_{i+1}$ for $1 \leq i < r$ and $|G| = d_1 d_2 \cdots d_r$. Because the unimodular Smith normal form is unique, this implies that this decomposition is also unique, and so, the isomorphism classes of finite abelian groups of order $n > 0$ are in bijection with the factorisations of n ,

$$n = \prod_{i=1}^r d_i$$

for which $d_i | d_{i+1}$ for $1 \leq i < r$. This allows us to classify isomorphism classes of finite abelian groups.

Example.

- $n = 4$ – the valid decompositions are 2 and 2×2 , so every group of order 4 is isomorphic to \mathbb{Z}_4 or $\mathbb{Z}_2 \oplus \mathbb{Z}_2$.
- $n = 15$ – the only valid decomposition is 15, so every group of order 15 is isomorphic to \mathbb{Z}_{15} and is hence necessarily cyclic.
- $n = 36$ – we have 36 , 2×18 , 3×12 , and 6×6 , so groups of order 36 are isomorphic to \mathbb{Z}_{36} , $\mathbb{Z}_2 \oplus \mathbb{Z}_{18}$, $\mathbb{Z}_3 \oplus \mathbb{Z}_{12}$, or $\mathbb{Z}_6 \oplus \mathbb{Z}_6$.

Lemma 7.24. *Let $G = \bigoplus_{i=1}^n G_i$ be a finite abelian group. Then, the order of $g = (g_1, g_2, \dots, g_n)$ is the least common multiple of the orders $|g_i|$ of the components of g .*