Abstract and Numerical Linear Algebra

Amir Kadivar

Winter 2010, 2011
Montréal

Preface

This is the report of a personal reading on abstract and numerical linear algebra during the winter of 2011. The first chapter is a review of (maybe unnecessary) Abstract Algebra concepts. The only part of this study that called for such an introduction was the study of the rational canonical form (discussed in section 8.3) which uses extensive use of evaluating polynomial functions of linear transformations which, as far as I understand, can not be justified without the background of section 1.4. But if the curiosity of the reader is not stimulated by "whether feeding linear transformation to polynomials is legitimate", and "why at times some of the trivial properties of polynomials do not hold when we do so", she can ignore this section as well as any reference I have given thereto through the passage.

In chapters 2 and 3 I have started with the basic properties of vector spaces and linear transformations on vector spaces, and have tried to raise, even with heavy simplifications, all the issues we try to tackle in later chapters, just to mention the raison-d'être of the subject matter of those chapters. I have given the notion of *duality* fairly rigorous attention, that made the analysis of *bilinear forms* a lot easier. From the properties of symmetric bilinear forms, all we need to know about *inner products* just follow, and therefore I have spent relatively short time on that subject. The most interesting part of this study is the study of canonical forms, embedded in which lies the *spectral theory*. As a prerequisite to this, I have given a derivation of the determinant, which borrows its main blueprint from [HK71]. I specifically felt the need to "derive" and prove the uniqueness (given the simplest and most reasonable of requirements) of what we know as "the" determinant, and not just accept it as some sort of "factually known to be useful" tool. Then I have gone through the study of the rational canonical form, which is the culminating result of an attempt to maximally breakdown the behavior of endomorphisms to undecomposable invariant direct sum decompositions. From there the *Jordan normal form* for the special case of algebraically closed carrier fields follows almost immediately.

My main guideline through this study was the amazing text by Hoffman and Kunze [HK71]. For the study of Jordan normal form, I have mainly used the text by Finkbeiner [Fin66]. The study of rational canonical form, due to its intimate relationship with polynomials, arises through heavy abstract algebraic notions (and even a lot of notions and terms are borrowed directly from the study of integral domains), and for the lay reader that I am in that field, Hungerford's text on abstract algebra [Hun74] and Curtis' neat text on linear algebra [Cur84] together opened the way to a sufficient grasp of simple abstract algebra to get to the point of relevance. Finally at all times I got lost in notions or lengthy proofs, John Armstrong's mathematical weblog[1] came to my rescue with useful intuitive explanations and in one occasion, that is indicated in the text by his name, with a short neat proof to replace the wordy proofs of textbooks.

[1] The Unapologetic Mathematician: `http://unapologetic.wordpress.com/`

Contents

Chapter 1

Abstract Algebra overview

We start to build up algebraic notions to get to the notion of a vector space. Given a set S, all algebraic constructs rely in some sort on binary operation[1], namely $* : S \times S \to T$, and are named according to their properties. The following are the possible collective behaviors of S and $*$ together, that we are most concerned about:

(i) S is *closed* under $*$ is for all $x, y \in S$, $x * y$ also lies in S.

(ii) $*$ is *associative* if for all $x, y, z \in S$ we have: $x * (y * z) = (x * y) * z$.

(iii) $*$ is *commutative* if for all $x, y \in S$ we have: $x * y = y * x$.

(iv) $x \in S$ is a left (right) *identity* with respect to $*$ if for all $y \in S$ we have: $y * x = y$ ($x * y = y$). Such element might not exist. Also if such an element exists such that it is both a right and left identity, we drop the direction prefix and just call it an identity (and that is specifically what we mean by an element being identity: it is both a right and left identity). One can easily prove that if $(S, *)$ has an identity, it must be unique.

(v) If $(S, *)$ has an identity, namely 1_S, then for any $x \in S$ an element $y \in S$ is the left (right) *inverse* of x, if $x * y = 1_S$ ($y * x = 1_S$). Some elements of S might not have a right or left inverse, but it can easily be shown that if they both exist, they have to be equal. If x is two-sided invertible, we call its inverse x^{-1}.

(vi) if $*$ and $\bullet$ are two binary operations under which S is closed, we say $\bullet$ is *distributive* over $*$ if for all $x, y, z \in S$ we have:

$$x \bullet (y * z) = (x \bullet y) * (x \bullet z)$$

$$(y * z) \bullet x = (y \bullet x) * (z \bullet x)$$

In the special case where $\bullet$ is commutative, any of the above implies the other one. But in general $\bullet$ has distribute over $*$ both from right and left to meet our definition.

[1] While referring to binary functions on algebraic constructs we prefer the operation notation (x, y) or $x * y$ to the $f(\cdot, \cdot)$ notation.

1.1 Elementary algebraic constructs

When equipping a set with a binary operation, we always assume that the set is closed under the operation. Here we define two most basic binary operations, addition and multiplication (not to be confused with their respective meaning in arithmetic), based on those of the above criteria that they satisfy. The simplest case we consider is when $*$ is just associative, and imposes an identity on S. Elements of S do not necessarily have inverses, and commutativity might not hold:

Definition 1.1.1. $(M, *)$ is a **monoid** if $*$ is associative and has an identity on M. If $*$ is commutative too, we call $(M, *)$ a commutative monoid.

As an example of a monoid, one can look at the set of all functions over a fixed set $M = \{f : A \to A\}$, with $*$ being function composition: $f * g = f \circ g$. Associativity obviously holds and there exists an identity (the identity function). This monoid is obviously not commutative (since function composition is not commutative), and there exist elements of M that do not have inverses. If a monoid allows for inverses all over its carrier set, the more special construct we get is a group:

Definition 1.1.2. $(G, *)$ is a **group** if $*$ is associative, not necessarily commutative, has an identity on G, and allows all members of G to have inverses. In this case we call $*$ an *addition* and denote its unique identity by 0_G. A commutative group $(G, *)$ (one whose $*$ is commutative too) is often called an **Abelian group**.

As a classic example one can look at the set of all rotation operations in $\mathbb{R}^3$ with function composition as the operator. This group obviously allows for inverses (any rotation has an inverse) but is not commutative. As the most obvious example of abelian groups, one can think of the set of integers $\mathbb{Z}$ equipped with arithmetic addition.
As mentioned before in a group the inverse of any element and the identity can easily be proven to be unique (note that in the definition just their existence is assumed). Thus one can introduce a new well-defined binary operation $x \div y = x * y^{-1}$. If we agree on calling $*$ an addition, $\div$ can be intuitively called a *subtraction*. If we, as some do, call $*$ a multiplication, then $\div$ is intuitively called a division.

Now we turn to the algebraic constructs with two binary operations:

Definition 1.1.3. $(R, \bullet, *)$ is a **ring** if (i) $(R, *)$ is an abelian group (this is called the additive group of the ring), (ii) $(R, \bullet)$ is a monoid, and finally (iii) that $\bullet$ distributes from right and left over $*$. We call $*$ the addition of the ring (hence we prefer to replace it by $+$), and $\bullet$ the *multiplication* of the ring (and hence we prefer to replace it by $\cdot$). We generally denote a ring by $(R, \cdot, +)$. We call the multiplicative identity of R the **one**, and denote it by 1_R, and as before we denote the additive identity by 0_R, both of which are obviously unique.

We notice that in a ring the addition admits all the desirable properties of a binary operation, and the limitation is due to the multiplication which in the most general case is just associative and admits an identity. If the multiplication is commutative too, we call $(R, \cdot, +)$ a *commutative ring*. One can simply show the following for a (not necessarily commutative) ring:

(i) The additive identity (zero) multiplied by anything (either from right or left) is again zero: let $x \cdot 0_R = y$ then $x = x.1_R = x.(1_R + 0_R) = x.1 + x \cdot 0_R = x + y$. Since the addition operation admits an inverse y has to be 0_R. The same argument can be made for left multiplication of zero.

(ii) As an immediate result of the above, the additive identity (the zero) and the multiplicative identity (the one) have to be distinct ($1_R \neq 0_R$).

An example of a non-commutative ring is $(R_1, \circ, +)$ where R_1 is the set $\{f : S \to S\}$ for some abelian group $(S, +)$, with function composition as multiplication, and the addition of $(R_1, +)$ being (roughly) inherited from $(S, +)$. For an example of a commutative ring one can simply look at $(R_2, \times, +)$ where R_2 is the set $\{f : \mathbb{R} \to \mathbb{R}\}$, with usual addition and arithmetic multiplication operators. Note that while $(R_2, \times, +)$ is a commutative ring, it does not admit an inverse over its multiplication.

We saw that 0_R times anything is 0_R again. From our arithmetic intuition, an interesting question would be that: "Are there elements $x, y \in R$ both distinct from 0_R such that $x \cdot y = 0_R$?". In general this is possible, and we call such non-zero members of R that can make a multiplication with other non-zero elements vanish, *zero divisors*. For example in the above exarmple of $(R_2, \times, +)$, one can easily see that the unique zero $0_R = f$ of R is the all-zero function: $f(\cdot) = 0$. One can simply think of two non-zero functions f_1 and f_2 such that their multiplication is zero.
Another interesting question one can ask is that: "in a commutative ring does $x \cdot y = x \cdot z$ imply $y = z$?". The answer in general is obviously no. But one might think the answer to the question is yes, only if the multiplication of the field admits an inverse. This is not true! One can easily prove that this propositions fails to be true if and only if $(R, \cdot)$ admits *zero divisors*. Thus it would make sense to give special attention to the case of commutative rings that do not allow zero divisors (but still do not admit multiplicative inverse). The next example makes us even more determined in this regard:

An important commutative ring is $(\mathbb{Z}, \times, +)$ with ordinary arithmetic addition and multiplication. As we mentioned earlier, if a binary operation admits inverses, then one can neatly define the inverse of the operation itself (subtraction for addition and division for multiplication). As a result one can not form a corresponding *division* on a ring in general. Despite this, there do exist elements in a commutative ring that divide each other, but if x and y are such that there exists a q such that $x = q \cdot y$, is such q unique? The answer can easily be seen to be "yes if and only if the multiplication operator of the ring does not allow zero divisors". Now if we assume the latter is true, for arbitrary x, y for which there might not be any element q such that $x = q \cdot y$, we could try to find the *closest* we can get to $x \div y$, i.e a $q \in R$ such that $x - q \cdot y$ is the smallest possible. Notice here that we have not explicitly talked about any notion of *order* (comparison function) yet. After these two motivating examples we define the notion of integral domains:

Definition 1.1.4. $(I, \cdot, +)$ is an **integral domain** if (i) $(I, \cdot, +)$ is a commutative ring, and (ii) the operator $\cdot$ does not allow for zero divisors.

The example above $(\mathbb{Z}, \times, +)$ is the classic example of integral domains (hence the name *integral*). The interesting treatise of the division algorithm and immediately thereafter the story of *factorization* and *prime* members of R makes complete sense now when we are dealing with integral

domains. We will go through such a treatise for a special family of rings known as polynomial rings in section 1.4.

1.2 Fields and Vector spaces

Getting back to where we left the rings, and in their special case integral domains, we notice that if the multiplication admits an inverse, all we said about the division algorithm and factorization would lose their meaning, since a well-defined division operator can be defined and there is no remainder to any division. And thus here comes the next elementary algebraic constructs we need:

Definition 1.2.1. $(F, \cdot, +)$ is a **field** if (i) $(F, \cdot, +)$ is a commutative ring, and (ii) the operator $\cdot$ admits inverses all over F except for the zero of the addition.

In other words, both operations of a field are commutative, associative, and admit identity and inverses (except for multiplicative inverse for the additive identity), and the multiplication distributes over addition. The following can easily be proved for a field $(F, \cdot, +)$:

(i) Since $(F, \cdot, +)$ is before anything a ring, we know that 0_F times anything is 0_F again, and also that $1_F \neq 0_F$.

(ii) There are no zero divisors: If x, y are both non-zero, they both have inverses, and hence $x \cdot y = 0$ yields to $1_F = y^{-1} \cdot x^{-1} \cdot x \cdot y = y^{-1} \cdot x^{-1} \cdot 0 = 0$ which is a contradiction!

Generalizing $(\mathbb{Z}, \times, +)$ one immediately can see that $(\mathbb{Q}, \times, +)$ is a field, with $\mathbb{Q}$ being the set of rational numbers. However, the most common fields we are involved with are $\mathbb{R}$ and $\mathbb{C}$, the set of real and complex numbers with ordinary multiplication and addition, respectively. Although the study of *finite fields* (a finite cardinality F) is of radical importance in field theory, we do not need this distinction for the case of this study (except for one theorem which we will meet later on).
Finally, we are ready to grasp the most important construct we need from abstract algebra for the study of linear transformations: vector spaces.

Let $\mathbb{F} = (F, \cdot, \oplus_f)$ be a field. We are already familiar with tying a bunch of members of F together and referring to them as vectors. The only operation we are used to expect from these new objects is a commutative addition over themselves (hence an abelian group). But we do not need to confine ourselves to this notion of vectors. We will here build an algebraic construct, a field and an abelian group on top of that. Specifically let $\mathbb{X} = (X, \oplus_x)$ be an abelian group (notice that we are defining the *vector*s independent of the field), and let $*\colon F \times X \to X$ be a new binary operation relating $\mathbb{F}$ and $\mathbb{X}$ to each other. We could think of $*$ as being a multiplication in the sense we will describe in a second. Furthermore we can think of $*$ being consistent in some sense with the multiplication operation that $\mathbb{F}$ is already armed with. Specifically, the sense in which we mean the last two properties to hold is the following:

(i) The new multiplication $*$ inherits the identity of $\cdot$, the natural multiplication of $\mathbb{F}$:

$$\forall u \in \mathbb{X},\ u \neq 0_X :\ 1_F * u = u$$

(ii) The new multiplication *distributes* over $\oplus_x$:

$$\forall \alpha \in \mathbb{F},\ \forall u, v \in \mathbb{X}:\ \alpha * (u \oplus_x v) = (\alpha * u) \oplus_x (\alpha * v).$$

(iii) The two multiplications $\cdot$ and $*$ *associate* together:

$$\forall \alpha, \beta \in \mathbb{F},\ \forall u \in \mathbb{X},\ (\alpha \cdot \beta) * u = \alpha * (\beta * u)$$

(iv) The new multiplication *distributes* on $\oplus_x$ and $\oplus_f$:

$$\forall \alpha, \beta \in \mathbb{F},\ \forall u \in \mathbb{X},\ (\alpha \oplus_f \beta) * u = (\alpha * u) \oplus_x (\beta * u)$$

Definition 1.2.2. We define $\mathscr{V} = (\mathbb{X}, \mathbb{F}, *)$ to be a **vector space** defined over the field $\mathbb{F}$, if $\mathbb{F} = (F, \cdot, \oplus_f)$ is a field, $\mathbb{X} = (X, \oplus_x)$ is an abelian group, and $*$ satisfies properties (i) to (iv) above. We specifically call $\mathbb{F}$ the field of **scalars** of $\mathscr{V}$.

From the above definition, the following properties can be shown to hold:

(i) 0_F times any thing is 0_X: Let $0_F * u = v$, then we have $u = (1_F \oplus_f 0_F) * u = u \oplus_x v$. Since $\oplus_x$ admits an inverse, v has to be 0_X.

(ii) Anything times 0_X is again 0_X: Let $\alpha * 0_X = u$, then we have

$$\alpha * v = \alpha * (0_X \oplus_x v) = (\alpha * 0_X) \oplus_x (\alpha * v) = u + \alpha * v$$

which can not hold unless $u = 0_X$.

(iii) The new multiplication does not allow for zero divisors, in the sense that $\alpha * u = 0_X$ happens only if (and by the (1) and (2) either $\alpha = 0_F$ or $u = 0_X$: Let $\alpha * u = 0$ and let $\alpha \neq 0$, then we have:

$$u = (\alpha^{-1} \cdot \alpha) * u = \alpha^{-1} * (\alpha * u) = \alpha^{-1} * 0_X = 0_X$$

(iv) The identity of $*$ is unique: let $u \neq 0_X$ and $\alpha * u = u$ then we have $\alpha * u = 1_F * u \Rightarrow (\alpha \oplus_f (1_F)^{-\oplus}) * u = 0_X$, where by $(1_F)^{-\oplus}$ we mean the additive inverse of 1_F. Now by (3) and the fact that $u \neq 0_X$ we get $\alpha = 1_F$.

Now that the operations $\oplus_f$, $\oplus_x$ and the multiplications $\cdot$ and $*$ are compatible to this extent, the first thing we need to do is to get rid of this messy notation! First of all we drop the X and F subscripts for additions and use $+$ for both, since the one we mean can always be inferred from the context, but we should not forget that they are in essence different operations, no matter how compatible they are together. Since the two multiplications are complete consistent with each other, we will drop both of them and use the common notation of $\alpha\beta = \alpha \cdot \alpha$, and $\alpha u = \alpha * u$. We also will denote the shared multiplicative identity of both multiplications 1_F by 1, with no subscript. We also drop the subscript from the additive identity of $\mathbb{F}$ and $\mathbb{X}$, and denote the two *different* elements 0_F and 0_X, by just 0. This will not cause any confusion as long as we are aware that 0 is now representing two elements, one in $\mathbb{X}$ and one in $\mathbb{F}$. And finally, just for convenience instead of differentiating between the algebraic construct $\mathscr{V}$ and the set X, we use the first one to

represent both, i.e. we will at times use the notation $u \in \mathscr{V}$ which is not exact since $u \in X$, and $\mathscr{V}$ is the algebraic construct $(\mathbb{X}, \mathbb{F}, *)$.
We have seen up to now some kinds of algebraic constructs, each of which is a special subset of one another. Here we review all the constructs we care about, and present examples. For the case of constructs with one binary operations we have the following in roughly decreasing order of generality:

type and notation	*property*	*example*
non-commutative monoid $(M, *)$	just has an identity	$(\{f\colon A \to A\}, \circ)$
non-commutative group $(G, *)$	monoid with inverse	$(R_3, \circ)$
commutative monoid $(M, *)$	$*$ is commutative	$(\mathbb{N}, \times)$
abelian group $(G, *)$	commutative group	$(\mathbb{Z}, +)$

Notice that in the example for a non-commutative group, R_3 is the set of all rotations in $\mathbb{R}^3$. For the case of constructs with two binary operations, we have the following examples. In all the cases we assume the usual addition and multiplication operations:

type and notation	*property*	*example*
non-commutative ring $(R, \cdot, +)$	$(R, +)$ an abelian group $(R, \cdot)$ just a monoid	$(\mathbb{R}^{n\times n}, \cdot, +)$ $\mathbb{R}^{n\times n}$ is the set of $n \times n$ matrices
commutative ring $(R, \cdot, +)$	$(R, \cdot)$ a commutative monoid	$(\mathbb{Z}[x], \cdot, +)$ $\mathbb{Z}[x]$: polynomials with coefficients in $\mathbb{Z}$
integral domain $(I, \cdot, +)$	$(I, \cdot, +)$ commutative ring $(I, \cdot)$ does not allow divisors of + identity	$(\mathbb{Z}, \times, +)$
field $(F, \cdot, +)$	$(F, \cdot, +)$ commutative ring $(F, \cdot)$ admits inverse except for 0_F	$(\mathbb{Q}, \times, +)$
vector space $(F, X, \cdot, *, \oplus_f, \oplus_x)$ $*\colon F \times X \to X$	$(F, \cdot, \oplus_f)$ a field $(X, \oplus_x)$ an abelian group $*$ is a multiplication *in some sense* $\oplus_x$, $\oplus_f$,$\cdot$, and $*$ are consistent	$(\mathbb{R}, \mathbb{R}^n, \cdot, +)$ scalars are real numbers vectors are real n-tuples notation for operations merged

1.3 Homomorphisms

We have seen that an algebraic construct is nothing but one or two carrier sets and one or more binary operations satisfying some properties over the carrier sets. These properties could be thought of as the *algebraic* structure of each construct. For example two fields have the same *type* of algebraic structure. They are not the same, since they have possibly different carrier sets, and possibly different operations. But counterpart operations in each field satisfy the same kind of properties on their respective carrier set. Thinking about mappings between same types algebraic constructs (for example a mapping from the ring of all 2×2 real matrices to the ring of all polynomials with integer coefficients), a desirable property would be that the mapping *respects the structure* (or as some say, *preserves the shape*) of the underlying algebraic construct (which in our example is "the ring structure"). In specific let $f\colon \mathbb{R}_1 \to \mathbb{R}_2$ be a mapping between the two aforementioned rings. The rings are the following:

$$\mathbb{R}_1 = (\mathbb{R}^{2\times 2}, \odot_1, \oplus_1)$$

$$\mathbb{R}_2 = (\mathbb{Z}[x], \odot_2, \oplus_2)$$

both equipped with common types of multiplication and addition operations we are familiar with on their respective carrier sets. For example let f be:

$$\begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \xrightarrow{f} \lfloor a_{11} \rfloor - \lfloor a_{12} \rfloor x + \lfloor a_{21} \rfloor x^2 - \lfloor a_{22} \rfloor x^3$$

Now the question of respecting the algebraic structure, breaks down to respecting the properties that make a set with two operations be a ring. In the case of properties of addition the question is whether following holds for any two 2×2 real matrices $x, y \in \mathbb{R}_1$:

$$f(x \oplus_1 y) = f(x) \oplus_2 f(y)$$

Similarly for the multiplications the question is whether the following holds for all two 2×2 real matrices :

$$f(x \odot_1 y) = f(x) \odot_2 f(y)$$

Obviously in our example f does not satisfy any of these, and hence is not a *shape preserving* mapping. The special kind of mappings that do respect the algebraic structure of the underlying construct, are called homomorphisms:

Definition 1.3.1. Let $\mathbb{A}$ and $\mathbb{B}$ be two instances of an algebraic construct (two groups, two rings, etc.). Any mapping $f : \mathbb{A} \to \mathbb{B}$ is a [construct name] **homomorphism** (e.g. a group homomorphism, a ring homomorphism, etc.) if for all the operations $*$ that are included in the definition of the construct, f respects $*$. Specifically if we denote the corresponding operations of $\mathbb{A}$ and $\mathbb{B}$ by $*_{\mathbb{A}}$ and $*_{\mathbb{B}}$ respectively, we expect that for all $x, y \in \mathbb{A}$ we have:

$$f(x *_{\mathbb{A}} y) = f(x) *_{\mathbb{B}} f(y)$$

and this should hold for all such operations $*$ of the construct. If f is bijective we call it an **isomorphism** and if $\mathbb{A} = \mathbb{B}$ we call it an **endomorphism**.

An immediate observation about homomorphisms is that they preserve any sort of identity that $\mathbb{A}$ might have. In the case of groups, just the additive identity, and for rings and fields, both additive and multiplicative identities. Specifically, if $a \in \mathbb{A}$ is the identity with respect to some operation $*$, then $f(a)$ is the same in $\mathbb{B}$. We will get back to homomorphisms when we start to treat vector spaces, and see how homomorphisms on vector spaces are in fact what we call linear transformations in linear algebra terminology.

1.4 Polynomial Rings

The last chunk of abstract algebra we need is some knowledge of polynomials defined over a ring. Getting back to the properties of rings, we can see that regardless of a ring being commutative or not, α^n, where n is non-negative integer, is completely well defined all over any arbitrary ring (with the understanding that $\alpha^0 = 1_R$ for all α). As soon as natural exponentiation is defined,

immediately one can define polynomial functions on a set. We here build an algebraic construct that corresponds to the notion of polynomials, but are *not* exactly the same as polynomial functions. We later provide means to relate polynomials as algebraic objects to polynomial functions. Any polynomial with coefficient in some set, can be regarded as a sequence of members of the set, with the property that after some point all the entries of the sequence are zero.

Definition 1.4.1. Let $(R, \cdot, +)$ be an arbitrary ring. A **polynomial** on R is a sequence of its members: $\{\alpha_0, \alpha_1, \alpha_2, \ldots\}$ such that all but at most a finite number of α_is are zero. We define $R[x]$ to be the set of all polynomials on R (x is just a matter of notation, not to be confused with the notion of *variable* in polynomial functions).

Remark. Although most authors define polynomials over fields, we have seen (and will see more) that not only the underlying construct being a ring suffices for defining the notion of a polynomial, but also the ring does not even need to be commutative for this notion to be well defined.

Again gaining intuition from what we know of polynomial functions, we equip $R[x]$ with an addition and a multiplication. Let $a = \{\alpha_i\}_{i=0}^{\infty}$ and $b = \{\beta_i\}_{i=0}^{\infty}$ be two members of $R[x]$. We define the following addition operation on $R[x]$:

$$a + b = \{\alpha_0 + \beta_0, \alpha_1 + \beta_1, \alpha_2 + \beta_2, \ldots\}$$

and the following *convolution* operator as the multiplication for $R[x]$:

$$\begin{aligned} a \cdot b &= \{\gamma_0, \gamma_1, \gamma_2, \ldots\} \\ \gamma_0 &= \alpha_0 \cdot \beta_0 \\ \gamma_1 &= \alpha_0 \cdot \beta_1 + \alpha_1 \cdot \beta_0 \\ \gamma_2 &= \alpha_0 \cdot \beta_2 + \alpha_1 \cdot \beta_1 + \alpha_2 \cdot \beta_0 \\ &\vdots \\ \gamma_n &= \sum_{i=0}^{n} \alpha_i \beta_{n-i} \end{aligned}$$

We can easily show that $R[x]$ together with the natural extension of the addition of R to polynomials and convolution as multiplication, is in fact a ring:

Proposition 1.4.1. *For any arbitrary ring $(R, \cdot, +)$, the set $R[x]$ together with the addition and multiplications defined above is a ring, hence we refer to $R[x]$ by the **polynomial ring** of R. Furthermore one can think as the above operations as generalizations of $\cdot$ and $+$ of R (and hence safely denote by $(R[x], \cdot, +)$ the polynomial ring), since the subring $R_0[x]$ of $R[x]$ containing all degree zero members of $R[x]$ is isomorphic to R.*

Definition 1.4.2. For any non-zero polynomial $p = \{\alpha_0, \alpha_1, \alpha_2, \ldots\}$ in $R[x]$, (almost) by definition there exists a unique $m \in \mathbb{N}$ such that:

$$\alpha_m \neq 0$$

$$\forall n > m : \ \alpha_n = 0$$

We define the **degree** of this polynomial to be m and write $\deg(p) = m$. We do not assign a degree to the polynomial $\{0, 0, 0, \ldots\}$.

Proposition 1.4.2. *Let R be a zero-divisor free ring, and $R[x]$ be its polynomial ring. For any two polynomials p, q in $R[x]$ the following hold:*

(i) $\deg(p+q) \leq \max\{\deg(p), \deg(q)\}$

(ii) $\deg(p \cdot q) = \deg(p) + \deg(q)$

Remark. Specific *caution* has to be made about the above properties of degrees of polynomials. These happen to hold if only if R is *free of zero-divisors*, which is not necessarily the case.

Getting back to the algebraic properties of $R[x]$, we can easily prove that $R[x]$ inherits important properties of R too:

Proposition 1.4.3. *For any arbitrary ring $(R, \cdot, +)$ and if we denote its polynomial ring by $R[x]$, if any of the following holds for R, so do they for $R[x]$:*

(i) *$(R, \cdot, +)$ is a commutative ring,*

(ii) *$(R, \cdot, +)$ does not allow for zero divisors,*

(iii) *$(R, \cdot, +)$ is an integral domain.*

Furthermore since R is isomorphic to a subring of $R[x]$, any of the above properties hold for $R[x]$ only if so do they for R.

Remark. It is important to notice that the polynomial ring does *not* inherit a multiplicative inverse from its carrier ring. Specifically the polynomial ring $\mathbb{F}[x]$ for any field $\mathbb{F}$ is *not* a field, but an *integral domain.*

Here comes the interesting part about relating polynomials as abstract algebraic objects, to polynomial functions. First of all we choose a notation to get closer in this regard. We use the alternative below notation for a polynomial:

$$p(x) = \alpha_0 + \alpha_1 x + \alpha_2 x^2 + \ldots + \alpha_n x^n = \sum_{i=0}^{n} \alpha_i x^n$$

with the understanding that $\alpha_0 x^0 = \alpha_0$. We notice that x is just a piece of notation we are using. Some authors regard x as a new element (distinct from members of R), that we use to build $R[x]$, and some just define x to be the following member of $R[x]$:

$$x = \{0, 1_R, 0, 0, \ldots\}$$

In any way, we now introduce the *evaluation* map:

Definition 1.4.3. Let $(R, \odot_r, +)$ and $(S, \odot_s, +)$ be two arbitrary (possibly identical) rings. And let $\varphi : R \to S$ be a homomorphism. An **evaluation map** with respect to φ is a mapping $\mathrm{ev}_\varphi : R[x] \times S \to S$. If to $p(x) \in R[x]$:

$$p(x) = \alpha_0 + \alpha_1 \odot_r x + \alpha_2 \odot_r x^2 + \ldots + \alpha_n \odot_r x^n = \sum_{i=0}^{n} \alpha_i \odot_r x^n$$

and $s \in S$, ev_φ assigns the following:

$$\text{ev}_\varphi\Big(p(x), s\Big) = \varphi(\alpha_0) + \varphi(\alpha_1) \odot_s s + \varphi(\alpha_2) \odot_s s^2 + \ldots + \varphi(\alpha_n) \odot_s s^n = \sum_{i=0}^{n} \varphi(\alpha_i) \odot_s s^i \in S$$

We call $\text{ev}_\varphi\Big(p(x), s\Big)$ the **evaluation of** $p(x)$ **at** s.

We have differentiated between the multiplication operations of R and S to emphasize the fact that by evaluating $p(x)$ on s we leave the ring on which our polynomials are defined (of which the coefficients are members), and land on a member of the domain ring S. But since φ is a homomorphism, we can drop this differentiation and simply write:

$$\text{ev}_\varphi\Big(p(x), s\Big) = \varphi(\alpha_0) + \varphi(\alpha_1)s + \varphi(\alpha_2)s^2 + \ldots + \varphi(\alpha_n)s^n = \sum_{i=0}^{n} \varphi(\alpha_i)s^i$$

We notice that all these definitions do not require any of R and S to be free of zero divisors or even commutative, the latter since the direction of multiplications are not changed in any of the operations. But some interesting observations can be made about these properties. We can see that by fixing a specific member of $R[x]$ the evaluation map induces a function $p_\varphi \colon S \to S$ such that:

$$p_\varphi(s) = \text{ev}_\varphi\Big(p(x), s\Big)$$

which is the common way we tend to look at polynomial functions. But a more interesting case is the one in which we fix an element of S. By doing so the evaluation map induces another function $s_\varphi \colon R[x] \to S$, which assigns to any polynomial in $R[x]$ its evaluation on s (for all polynomials and a fixed s). This intuition is quite useful and we will meet again the same notion in the study of bilinear forms. The new mapping s_φ is a mapping from $R[x]$ to S, so one might wonder if s_φ is a homomorphism. Obviously s_φ respects the addition of both rings:

$$s_\varphi(p(x) \underset{R[x]}{\oplus} q(x)) = \sum_{i=0}^{m} \varphi(\alpha_i \oplus_r \beta_i)s^m = \sum_{i=0}^{m} \varphi(\alpha_i)s^i \oplus_s \sum_{i=0}^{m} \varphi(\beta_i)s^i$$

and the last term is in fact $s_\varphi(p) \oplus_s s_\varphi(q)$. But what about multiplication:

$$s_\varphi(p(x) \underset{R[x]}{\odot} q(x)) = \text{ev}_\varphi(p(x) \underset{R[x]}{\odot} q(x), s) \stackrel{?}{=} p_\varphi(s)q_\varphi(s)$$

Relaxing the terminology rigor, the question is the following:

$$(pq)_\varphi(s) \stackrel{?}{=} p_\varphi(s)q_\varphi(s)$$

which is considered to be brutally trivial when we are talking about polynomial functions with real coefficients and operating over $\mathbb{R}$. Unfortunately this is not always the case. By expanding the above we get:

$$\sum_{i=0}^{m} \varphi(\gamma_i)s^i \stackrel{?}{=} \left(\sum_{i=0}^{m} \varphi(\alpha_i)s^i\right)\left(\sum_{i=0}^{m} \varphi(\beta_i)s^i\right)$$

where:

$$\gamma_i = \sum_{j=0}^{i} \alpha_i \beta_{i-j}$$

Since φ is a homomorphism we can write:

$$\varphi(\gamma_i) = \varphi\left(\sum_{j=0}^{i} \alpha_i \beta_{i-j}\right) = \sum_{j=0}^{i} \varphi(\alpha_i)\varphi(\beta_{i-j})$$

Which is the coefficient of the i-th power of s on the left hand side, where the corresponding term on the right hand side does not have a neat i-th power of s, and instead is the following

$$\sum_{j=0}^{i} \varphi(\alpha_i) s^j \varphi(\beta_{i-j}) s^{j-i}$$

Notice that all the above multiplications are applications of $\odot_s$, and a *sufficient condition* for s_φ to be a homomorphism is that the multiplication of S is commutative, in which case we would havve:

$$\sum_{j=0}^{i} \varphi(\alpha_i) s^j \varphi(\beta_{i-j}) s^{j-i} = \sum_{j=0}^{i} \varphi(\alpha_i)\varphi(\beta_{i-j}) s^i = \gamma_i s^i$$

Since there exists a homomorphism from R to S, S being a commutative ring implies that the multiplication of R is commutative as well. We have proved the following:

Proposition 1.4.4. *Let $R[x]$ be the polynomial ring of an arbitrary ring R, and S be another arbitrary ring, to which from R we are armed with a homomorphism φ. A sufficient condition for the induced function $s_\varphi : R[x] \to S$ by the evaluation mapping being a homomorphism for all members of S, is that the multiplication of R and S both be commutative. In more relaxed terms, the sufficient condition for*

$$\forall s \in S,\ \forall p(x), q(x) \in R[x] :\ (pq)_\varphi(s) = p_\varphi(s) q_\varphi(s)$$

to hold, is that R and S are both commutative rings.

Although the above condition is not necessary, a perfect example of the above not holding is where R and S are both the non-commutative ring of all 2×2 real matrices (with φ being the identity homomorphism), with ordinary matrix addition and multiplication. Let I be the identity matrix and M be some 2×2 matrix, and let:

$$p = \{M, I, 0, 0, \ldots\} = M + Ix$$

$$q = \{-M, I, 0, 0, \ldots\} = -M + Ix$$

be two members of $R[x]$. Their multiplication in $R[x]$ is the convolution of the above sequences:

$$pq = \{M \cdot (-M),\ M \cdot I - I \cdot M,\ I \cdot I,\ 0, 0, \ldots\}$$

which simplifies to:

$$pq = \{-M^2, 0, I, 0, 0, \ldots\} = -M^2 + Ix^2$$

For some member $s \in S$ which in its own turn is a 2×2 real matrix we have:

$$pq(s) = s^2 - M^2$$

$$p(s) = s + M$$

$$q(s) = s - M$$

And there exist endless real matrices for which the equation does not hold, i.e such that:

$$s^2 - M^2 \neq (s - M)(s + M)$$

In fact any two matrices that do not commute in matrix multiplication make the proposition fail to be true, for instance take:

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

and

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

At the end of this part, we look at what a *root* means in our algebraic formulation of polynomials. We saw that evaluation, in its most general sense, happens when a ring homomorphic to the carrier ring of the polynomial ring exists. In a special case, if we use the same carrier ring as the domain of evaluation, the immediate choice for a homomorphism $\varphi \colon R \to R$ is the identity mapping I. In this case the evaluation map would be one that does not travel between different rings, and thus for any polynomial $p \in R[x]$, and that is what we expect to be understood when we simply write $p(\xi)$ for some $\xi \in R$. In this sense the evaluation of p acts as a function on R:

$$\xi \longmapsto p(\xi) := \mathrm{ev}_{\mathrm{I}}(p, \xi) = \sum_{n=0}^{\deg(p)} \alpha_n \xi^n$$

In this sense those elements of R that make the above mapping vanish to 0_R, the additive identity of R, are what we are used to call *roots* of the polynomial $p(x) \in R[x]$.

1.5 Factorization in Integral Domains

As we mentioned before in a remark to proposition 1.4.3, the polynomial ring is an integral domain if (and only if) the carrier ring itself is so, but never goes further than that, i.e the polynomial ring of a field is still just an integral domain. Thus one can think of an analogous division algorithm (and hence the notion of prime elements, and factorization to prime elements) for the integral domain of polynomials. We also noted when we first mentioned the division algorithm, that we need some sort of order over the integral domain on which we want to divide elements by each other. Since we have assigned a non-negative integer label to all polynomials (their degree), that is the handiest way to go. Furthermore we will see that in fact *small* polynomials in that sense, in fact induce

way simpler polynomials. From now on we will use $\prec$ to denote the order imposed by degrees of polynomials:

$$p(x) \prec q(x) \Leftrightarrow \deg(p) < \deg(q)$$

But in order for us to be able to define the division neatly and in correspondence to what we are already familiar with factorization of polynomial factorization, we have to assume that the carrier ring admits a multiplicative inverse. Also as we mentioned in proposition 1.4.2 the nice properties of the degree of addition and multiplication of polynomials, holds when the carrier ring is free of zero divisors. Putting all these together we will here assume that the coefficients are coming from a field.

Here we mention a series of propositions about polynomial rings of *fields* (the polynomial ring is thus an integral domains) that are all easy to prove, and are analogous to similar results in elementary number theory. From now on we fix the carrier field $\mathbb{F}$ to be any field, and all the polynomials we will be talking about to be members of $\mathbf{F}[x]$. Also by the *zero polynomial* we mean the additive identity of the ring $\mathbb{F}[x]$ which is $0_{\mathbf{F}[x]} = \{0, 0, 0, \ldots\}$, and by the *one polynomial* we mean the multiplicative identity of the ring $\mathbb{F}[x]$ which is $1_{\mathbb{F}[x]} = \{1_{\mathbb{F}}, 0, 0, \ldots\}$. We also define a polynomial to be **monic** if it has its leading coefficient to be $1_{\mathbf{F}}$.

Remark. For all the discussions in this section we always have an assumption the the polynomials we claim statements about are monic. However this is not a heavy assumption since in case they are not, all the coefficients be simultaneously divided by the leading coefficient (which is not zero), and almost nothing would change in the course of arguments, except for *uniqueness* arguments.

Proposition 1.5.1 (Division theorem)**.** *For any polynomial $a(x)$ and any nonzero monic polynomial $b(x)$, one can find unique polynomials q and r such that:*

$$a = q \cdot b + r \ , \qquad r \prec b$$

*We will refer to $q(x)$ and $r(x)$ by the **quotient** and **remainder** of dividing $a(x)$ by $b(x)$. We say $b(x)$ **divides** $a(x)$, or $b(x)$ is a factor of $a(x)$ and write $b(x)|a(x)$ if $r(x)$ happens to be the zero polynomial.*

Remark. If the dividend $b(x) = \{\beta_0, \beta_1, \ldots, \beta_{\deg(b)}, 0, 0, \ldots\}$ is not monic, the division theorem would be used for $b'(x)$ which is the polynomial obtained by dividing all coefficients of $b(x)$ by its leading coefficient. It is easy to show that if $a = b' \cdot q + r$ satisfies the criterion of the division algorithm, we have $r \prec b$ and also $a = b \cdot q' + r$ where $q'(x)$ is the polynomial whose coefficients are all $\beta_{\deg(b)}$ times the coefficients of $q(x)$.

Remark. It is easy to show that any non-zero polynomial has at least two divisors, itself and the one polynomial. We call these two the trivial divisors.

Definition 1.5.1. A *monic* non-zero polynomial $p(x)$ is called to be **prime** if the only *monic divisors* of $p(x)$ are its trivial divisors.

Remark. By using proposition 1.4.2, it can easily be seen that regardless of the structure of the underlying field $\mathbb{F}$, all degree one polynomials in $\mathbb{F}[x]$ are prime.

Definition 1.5.2. For any set of non-zero polynomials $f_1, f_2, \ldots, f_n$ there exist a unique monic polynomial g which we call the **greatest common divisor** (or simply the **g.c.d**) of $f_1, f_2, \ldots, f_n$, that satisfies the following:

(i) g divides all f_i.

(ii) g has the largest degree among all the common divisors of $f_1, f_2, \ldots, f_n$.

We also use the notation $(f_1, f_2, \ldots, f_n)$ to refer to the greatest common divisor. We also say that $f_1, f_2, \ldots, f_n$ are **relatively prime** if their g.c.d is the one polynomial.

Remark. It can easily be proved that any common divisor of $f_1, f_2, \ldots, f_n$ also divides their greatest common divisor:

$$b|f_1,\ b|f_2,\ \ldots\ , b|f_n\ \Rightarrow b|(f_1, f_2, \ldots, f_n)$$

Proposition 1.5.2. *If a prime polynomial p divides the multiplication $f_1 f_2 \ldots f_n$ there should exist some f_i which p divides:*

$$p|f_1 f_2 \ldots f_n \Rightarrow \exists i :\ p|f_i$$

Proposition 1.5.3 (Bézout's Theorem)**.** *For any n non-zero polynomials $f_1, f_2, \ldots, f_n$, with their their g.c.d being named g, there exist polynomials $t_1, t_2, \ldots, t_n$ such that:*

$$g = f_1 t_1 + f_2 t_2 + \ldots + f_n t_n$$

Proposition 1.5.4 (Unique Factorization Theorem)**.** *For any non-zero monic polynomial f there exist prime polynomials $p_1, p_2, \ldots, p_n$ and positive integers $e_1, e_2, \ldots, e_n$ such that:*

$$f = p_1^{e_1} p_2^{e_2} \ldots p_n^{e_n}$$

Furthermore this prime decomposition of f is unique, up to a reordering of $p_i^{e_i}$.

Here it is worthwhile mentioning two important observations regarding the relationship between the factors of a polynomial, and its *roots*. We call a polynomial who has degree 1, a **linear factor**.

Proposition 1.5.5. *Let $f \in \mathbb{F}[x]$ be a non-zero polynomial, and let $\xi \in \mathbb{F}$ be a root of f, in the sense defined in 1.4, then the following polynomial divides p:*

$$b = \{-\xi, 1_R, 0, 0, \ldots\}$$

In other words:

$$f(\xi) = 0_{\mathbb{F}} \quad \Rightarrow (x - \xi)|f(x)$$

And finally:

Proposition 1.5.6. *[Fundamental Theorem of Algebra] Referring to the field of complex numbers by $\mathbb{C}$, the only prime polynomials in $\mathbb{C}[x]$ are monic degree one polynomials (which are all always prime regardless of the underlying field).*

Proposition 1.5.7. *If the underlying field of the polynomial ring is the set of real of numbers $\mathbb{R}$, the set of prime polynomials of $\mathbb{R}[x]$ consists of all monic degree one polynomials and those monic degree two polynomials that for some $\alpha, \beta \in \mathbb{R}$ satisfying $\alpha^2 < \beta$, have the following form:*

$$x^2 + 2\alpha x + \beta$$

Chapter 2

Basics of Vector Spaces

First we sum up the definition of a vector space we offered in chapter 1 in relaxed terminology. Given a field of scalars $\mathbb{F}$, the set $\mathscr{V}$ with an associative and commutative addition with inverse and identity[1] under which it is closed, is a vector space over $\mathbb{F}$, if (i) scalar multiplication distributes over both vector addition and scalar addition, (ii) $\mathscr{V}$ is closed under scalar multiplication, (iii) the additive zero of $\mathscr{V}$ is scalar multiplicative zero, and (iv) additive zero of $\mathbb{F}$ is a scalar multiplicative zero:

(1) $u + v \in \mathscr{V}$ $\quad (\forall u, v \in \mathscr{V})$

(2) $u + v = v + u$ $\quad (\forall u, v \in \mathscr{V})$

(3) $u + (v + w) = (u + v) + w$ $\quad (\forall u, v, w \in \mathscr{V})$

(4) $u + 0_{\mathscr{V}} = u$ $\quad (\exists!\ 0_{\mathscr{V}} \in \mathscr{V}, \forall u \in \mathscr{V})$

(5) $u + (\text{-}u) = 0_{\mathscr{V}}$ $\quad (\forall u \in \mathscr{V}, \exists!\ (\text{-}u) \in \mathscr{V})$

(6) $\alpha u \in \mathscr{V}$ $\quad (\forall \alpha \in \mathbb{F}, \forall u \in \mathscr{V})$

(7) $0_{\mathbb{F}} u = 0_{\mathscr{V}}$ $\quad (\forall \alpha \in \mathbb{F}, \forall u \in \mathscr{V})$

(8) $\alpha(u + v) = \alpha u + \alpha v$ $\quad (\forall \alpha \in \mathbb{F}, \forall u, v \in \mathscr{V})$

(9) $(\alpha + \beta)u = \alpha u + \beta u$ $\quad (\forall \alpha, \beta \in \mathbb{F}, \forall u \in \mathscr{V})$

Definition 2.0.3. If $\mathscr{V}$ is a vector space over $\mathbb{F}$ and $\mathscr{W} \subseteq \mathscr{V}$. We say that $\mathscr{W}$ is a **subspace** of $\mathscr{V}$ if $\mathscr{W}$ qualifies as a vector space over $\mathbb{F}$ too. Also since $\mathscr{W}$ shares its carrier field with $\mathscr{V}$, it is a vector space if and only if it is closed under vector addition and scalar multiplication (all the other necessities of being a vector space are automatically inherited from $\mathscr{V}$).

Corollary 2.0.1. If $\mathscr{W}$ and $\mathscr{Y}$ are subspaces of the vector space $\mathscr{V}$, then so are $\mathscr{W} \cap \mathscr{Y}$ and $\mathscr{W} \cup \mathscr{Y}$.

[1] additive zero $0_{\mathscr{V}}$ which can easily be proved to be unique given its existence.

2.1 Different meanings of linear (in)dependence

Over the course of years, a source of confusion for me has been the different meanings in which "linear (in)dependence" is used in the literature. These different meanings do not translate to one another easily, and if not taken care of, give rise to nasty hidden bugs. I here classify these different meanings and define them separately, and later on prove some relationships each has with the other two. The different meanings are the following, in increasing order of complexity:

(i) A single vector is linearly dependent or independent *of* a set of vectors (a vector vs. set relationship).

(ii) Two sets of vectors are linearly dependent or independent *of* each other (a set vs. set relationship).

(iii) A set of vectors is linearly dependent or independent (a property of a set as a whole).

There exist a couple of equivalent ways defining each of the above, but the most consistent way, I think, is to define all the above using the notion of the linear hull (span) of a set of vectors:

Definition 2.1.1. Let $\mathscr{V}$ be a vector space over the field $\mathbb{F}$. Let $C = \{u_1, u_2, \ldots, u_k\}$ be an ordered set of vectors in $\mathscr{V}$. For any ordered set of scalars $\{\alpha_1, \alpha_2, \ldots, \alpha_k\}$, we call the following the corresponding **linear combination** of $u_1, u_2, \ldots, u_k$:

$$v = \sum_{i=1}^{k} \alpha_i u_i$$

Of course, by definition[2], All linear combinations of any set of vectors in $\mathscr{V}$ lie again in $\mathscr{V}$. We define the **linear hull** or the **span** of C to be the set of all linear combinations of members of C, and use the following notation:

$$\operatorname{span}(C) = \operatorname{span}\{u_i\}_{i=1}^{k} = \{ \sum_{i=1}^{k} \alpha_k u_k : \forall i, \alpha_i \in \mathbb{F}\}$$

Definition 2.1.2 (*vector-set independence*)**.** Let $\mathscr{V}$ be a vector space over the field $\mathbb{F}$. Let $C = \{u_1, u_2, \ldots, u_k\}$ be an ordered set of vectors in $\mathscr{V}$. For any vector $v \in \mathscr{V}$, we say that v is linearly dependent on C if v lies in $\operatorname{span}(C)$ and linearly independent of C otherwise. Also if $v = \sum_{i=1}^{k} \alpha_i u_i$ we say that the ordered set $\{\alpha_1, \alpha_2, \ldots, \alpha_k\}$ is **a representation** of v in terms of C. Also if a vector $v \in \operatorname{span}(C)$ has two or more distinct representations in terms of members of C, we say that v is a **duplicate member** of $\operatorname{span}(C)$.

Remark. Obviously for any set of vectors C, we have $\operatorname{span}(C) \subseteq \mathscr{V}$.

Proposition 2.1.1. *Let $\mathscr{V}$ be a vector space over the field $\mathbb{F}$. The span of any ordered set C of vectors in $\mathscr{V}$ is again a vector space over $\mathbb{F}$ and hence a subspace of $\mathscr{V}$.*

[2] Any vector space is closed under scalar multiplication and vector addition.

Proof. Let C be $\{u_1, u_2, \ldots, u_k\}$. All we need to prove is that span(C) is closed under scalar multiplication and vector addition. The other properties of a vector space obviously are inherited from $\mathscr{V}$. For any $v, w \in \text{span}(C)$ and any scalar γ we have:

$$v = \sum_{i=1}^{k} \alpha_i u_i \Rightarrow \gamma v = \sum_{i=1}^{k} (\gamma\alpha_i) u_i \in \text{span}(C)$$

$$w = \sum_{i=1}^{k} \beta_i u_i \Rightarrow v + w = \sum_{i=1}^{k} (\alpha_i + \beta_i) u_i \in \text{span}(C)$$

□

Corollary 2.1.1. If $\mathscr{W}$ is a subspace of $\mathscr{V}$, its span is itself: span$(\mathscr{W}) = \mathscr{W}$.

Corollary 2.1.2. Let C_1 and C_2 be two sets of vectors in $\mathscr{V}$, such that $C_1 \subseteq C_2$. Obviously any linear combination of members of C_1 is also a linear combination of members of C_2. As a result the span of C_1 is a subspace of C_2:

$$C_1 \subseteq C_2 \Rightarrow \text{span}(C_1) \subseteq \text{span}(C_2)$$

Definition 2.1.3. Any subspace of $\mathscr{V}$ contains the zero vector of $\mathscr{V}$. Thus any two subspaces have a trivial intersection at the zero vector. We say two subspaces of $\mathscr{V}$ **intersect non-trivially** if they share more than the zero vector.

No we define the second notion of linear (in)dependence that could be used:

Definition 2.1.4 (*set-set independence*)**.** Let C_1 and C_2 be two sets of non-zero vectors in $\mathscr{V}$. We say that C_1 and C_2 are linearly independent *of* each other if the subspaces span(C_1) and span(C_2) intersect only trivially, and linearly dependent, otherwise.

Corollary 2.1.3. If two *set*s are linearly independent of each other, it immediately follows that all the members of each set (and all the members of their linear hulls) are individually linearly independent of the other set.

Definition 2.1.5 (*innate set independence*)**.** Let $\mathscr{V}$ be a vector space over the field $\mathbb{F}$. We say that $u_1, u_2, \ldots, u_k$, all non-zero vectors in $\mathscr{V}$, are **linearly independent** if span$\{u_i\}_{i=1}^{k}$ does not contain any duplicate vector (i.e. all the vectors in span$\{u_i\}_{i=1}^{k}$ have unique representations), and linearly dependent otherwise.

Proposition 2.1.2. *Let $\mathscr{V}$ be a vector space over the field $\mathbb{F}$. Let $C = \{u_1, u_2, \ldots, u_k\}$ be an ordered set of non-zero vectors in $\mathscr{V}$. The span of C contains non-zero duplicate members, if and only if 0 is a duplicate member of* span(C).

Proof. Let 0 be a duplicate member of span(C). It already has a trivial representation in terms of C (an all zero sequence). Now let's assume there exists a sequence such that:

$$\sum_{i=1}^{k} \alpha_i u_i = 0$$

such that not all α_i are zero. Let t be the smallest index for which $\alpha_t \neq 0$:

$$\sum_{i=t}^{k} \alpha_i u_i = 0 \Rightarrow \alpha_t u_t = \sum_{i=t+1}^{k} \alpha_i u_i$$

since α_t is not zero we have:

$$u_t = \sum_{i=t+1}^{k} \frac{\alpha_i}{\alpha_t} u_i$$

Thus u_t is a non-zero duplicate member of span(C).
Now let v be any non-zero duplicate member of span(C).

$$v = \sum_{i=1}^{k} \alpha_i u_i = \sum_{i=1}^{k} \beta_i u_i$$

thus:

$$\sum_{i=1}^{k} (\alpha_i - \beta_i) u_i = 0$$

since not all α_i and β_i are equal, the left hand side is a representation of 0, different from its trivial representation, and hence 0 is a duplicate member of span(C). □

Corollary 2.1.4. For any ordered set of vectors $C = \{u_1, u_2, \ldots, u_k\}$ in $\mathscr{V}$ the following are equivalent:

(i) C is linearly dependent.

(ii) 0 is a duplicate member of span(C).

(iii) There exists a sequence of scalars $\{\alpha_1, \alpha_2, \ldots, \alpha_k\}$, not all zero, such that $\sum_{i=1}^{k} \alpha_i u_i = 0$.

(iv) There exists $u_t \in S$ such that $u_t \in \text{span}(C - \{u_t\})$(i.e u_t is a linear combination of other members of C and hence has two distinct representation in terms of C).

(v) span(C) has non-zero duplicate members.

Corollary 2.1.5. Analogously for any ordered set of vectors $C = \{u_1, u_2, \ldots, u_k\}$ in $\mathscr{V}$ the following are equivalent:

(i) C is linearly independent.

(ii) The only set of scalars $\{\alpha_1, \alpha_2, \ldots, \alpha_k\}$ that satisfy $\sum_{i=1}^{k} \alpha_i u_i = 0$, is $\alpha_1 = \alpha_2 = \ldots = \alpha_k = 0$.

(iii) For all $u_i \in S$ we have: $u_i \notin \text{span}\{u_j\}_{j \neq i}$.

(iv) All the vectors in span(C) have unique representations in terms of C.

Notice that there is a lot of distinction between the three definitions we have mentioned. Linear independence, with no relative reference, is an innate property of a *set* of vectors, and can not be easily broken down to linear (in)dependence of its subsets (however, look at proposition 2.1.5). Two sets can be linearly independent *of* each other without being linearly independent themselves. In fact, hand being linearly (in)dependent *of* a set is the property of a vector or a set with respect to a set. The three notions we mentioned can be tied together as follows:

Proposition 2.1.3. *Let $C = \{u_1, u_2, \ldots, u_k\}$ be a set of vectors in $\mathscr{V}$. The following are equivalent:*

(i) *C is a linearly independent set of vectors.*

(ii) *Each member of C is linearly independent of the rest (all u_i is linearly independent of $S - \{u_i\}$).*

(iii) *Linear hulls of any two disjoint subsets of C intersect only trivially (i.e. Any two disjoint subsets of C are linearly independent of each other).*

As we mentioned there is no way one could neatly break down the requirements for a set to be linearly independent, by splitting the set into multiple subsets. In fact, this usually gives rise to nasty traps. Here we mention some traps, about which we should be cautious.

Proposition 2.1.4 (Linear independence arguments traps)**.** *Let C_1 and C_2 be two ordered sets of vectors. If all the members of C_1 are individually linearly independent **of** C_2, then:*

(i) *It does **not** follow in general that all the members of C_2 are also individually linearly independent of C_1.*

(ii) *It does **not** follow in general that C_1 and C_2 are, in the set vs set meaning, linearly independent of each other (i.e.* $\text{span}(C_1)$ *and* $\text{span}(C_2)$ *might intersect non-trivially).*

(iii) *Adding the assumption that C_1 and C_2 are both linearly independent sets of vectors, it does **not** follow in general that $C_1 \cup C_2$ is a linearly independent set of vectors.*

Proof. Take the following example for $\mathscr{V} = \mathbb{R}^3$: Let C_2 be two linearly independent vectors in the $x - y$ plane, and C_1 be three linearly independent vectors anywhere but in the $x - y$ plane. All the vectors in C_1 are linearly independent *of* C_2, and C_1 and C_2 are individually linearly independent. But none of the above hold. □

In fact no more general independence conclusion can be made from all members of a set of vectors being *individually* linearly independent of another set. As one can see, the most straightforward to look at linear independence is through linear hulls. We will deliberate this idea more when we are discussing direct sum decompositions.

Proposition 2.1.5 (Adding vectors to sets of linearly independent vectors)**.** *Let C be a set of linearly independent vectors in $\mathscr{V}$. For any vector $v \in \mathscr{V}$, not in C, the set $S \cup \{v\}$ is linearly independent if v does not belong to* $\text{span}(C)$ *(i.e v is linearly independent of C).*

2.2 Dimensionality

We now turn our attention to the notion of dimensionality:

Definition 2.2.1. A set of linearly independent vectors $u_1, u_2, \ldots, u_n$ in a vector space $\mathscr{V}$, is a **maximal** set of linearly independent vectors, if one can not add any other vector v of $\mathscr{V}$ to the set that would not violate linear independence. Obviously for such a set we have:

$$\mathscr{V} = \text{span}\{u_i\}_{i=1}^{k}$$

Proposition 2.2.1. *Let $\mathscr{V}$ be a vector space over $\mathbb{F}$. If $C_1 = \{u_1, u_2, \ldots, u_k\}$ is a maximal linearly independent sets of vectors in $\mathscr{V}$, then for any linearly independent set of vectors $C_2 = \{v_1, v_2, \ldots, v_t\}$, we have $t \leq k$. And equality holds if only if C_2 is maximal as well.*

Proof. To prove the inequality we perform induction on k. For the case where $k = 1$ the result is obvious. Now Let's assume $k > 1$, and that for any vector space if there exists a maximal set of $k' < k$ linearly independent vectors, no set of more then k' linearly independent vectors could exist. Since C_1 is maximal we have:

$$\mathscr{V} = \text{span}\{u_i\}_{i=1}^{k}$$

and thus all members of C_2 are linear combinations of $\{u_i\}_{i=1}^{k}$:

$$\begin{aligned} v_1 &= \alpha_{1,1}u_1 + \alpha_{1,2}u_2 + \ldots + \alpha_{1,k}u_k \\ v_2 &= \alpha_{2,1}u_1 + \alpha_{2,2}u_2 + \ldots + \alpha_{2,k}u_k \\ \vdots & \qquad\qquad \vdots \\ v_t &= \alpha_{t,1}u_1 + \alpha_{t,2}u_2 + \ldots + \alpha_{t,k}u_k \end{aligned}$$

Looking at the representation of v_1 we notice that since v_1 is non-zero (if any of the vectors of a set is zero, the set can not be linearly independent), not all $\alpha_{1,i}$ can be zero. Without loss of generality we can assume $\alpha_{1,1} \neq 0$. We now build $t-1$ new vectors:

$$\begin{aligned} w_1 &= v_2 - \tfrac{\alpha_{2,1}}{\alpha_{1,1}}v_1 \\ w_2 &= v_3 - \tfrac{\alpha_{3,1}}{\alpha_{1,1}}v_1 \\ \vdots & \qquad \vdots \\ w_{t-1} &= v_t - \tfrac{\alpha_{t,1}}{\alpha_{1,1}}v_1 \end{aligned}$$

We first notice that the new set of vectors $\{w_1, w_2, \ldots, w_{t-1}\}$ is linearly independent, (roughly) since any vanishing linear combination of w_i, would immediate result (by definition of w_i, in a vanishing linear combination of v_i. Also it is easy to see that:

$$w_1, w_2, \ldots, w_{t-1} \in \text{span}\{u_i\}_{i=2}^{k}$$

and that since C_1 is maximal in $\mathscr{V}$, the set $\{u_2, u_3, \ldots, u_k\}$ is a maximal set of $k-1$ linearly independent vectors in the subspace $\mathscr{V}' = \text{span}\{u_i\}_{i=2}^{k}$. Thus we have found a set of $t-1$ linearly

independent vectors in the vector space $\mathscr{V}'$, in which also exists a maximal set of $k-1$ linearly independent vectors. From the induction hypothesis we have $t-1 \leq k-1$ and hence the inequality is proved.
Now let's consider the case of equality. Obviously if C_2 is maximal too, the same inequality would hold the other way around and we would have $t \geq k$ and equality would hold. But is there any other situation in which equality could hold? Let's assume $t = k$,and C_2 is not maximal. As a result there exists a maximal set C_2' with $t+s$ members, such that $C_2 \subset C_2'$. Now by using what we just proved, that any two maximal sets have the same number of members, for C_1 and C_2' we get $t+s = k$, and thus $s = 0$. And this holds for any possible amendment to C_2 not violating linear independence. Thus C_2 has to be maximal as well. □

Definition 2.2.2. We just proved that all maximal sets of linearly independent vectors in any vector space $\mathscr{V}$ have the same number of elements. We call this specific number the **dimension** of $\mathscr{V}$ and denote it by $\dim(\mathscr{V})$.

Corollary 2.2.1. From proposition 2.2.1, it immediately follows that any set of linearly independent vectors in $\mathscr{V}$ can not have more than $\dim(\mathscr{V})$ members. And as a result, any set with more than $\dim(\mathscr{V})$ vectors from $\mathscr{V}$ is doubtlessly linearly dependent.

Definition 2.2.3. As we saw all *maximal* sets of linearly independent vectors have the same number of elements $\dim(\mathscr{V})$, and hence are all *maximum* in the sense of dimensionality of their linear hull. From now on we drop the term "maximal set of linearly independent vectors" for the term *basis*. We call a set $\mathscr{B}$ of vectors in $\mathscr{V}$ a **basis** for $\mathscr{V}$ if $\mathscr{B}$ (i) is linearly independent, and (ii) has $\dim(\mathscr{V})$ elements. It immediately follows from proposition 2.2.1 that $\mathscr{B}$ has to be maximal and hence spans all $\mathscr{V}$:

$$\mathscr{V} = \text{span}(\mathscr{B})$$

Remark. We adopted two new terms for "maximal set of linearly independent vectors" and for the "shared cardinality" of such sets, *basis* and *dimensionality*, respectively. However, one should not forget that the original lengthy terms contain the definition of the two new ones. At all times we we have to refer implicitly to the lengthy original terms, to prove properties of bases and dimensionalities.

Corollary 2.2.2. Let C be a set of vectors in $\mathscr{V}$. Noticing that if C is linearly independent, it is maximal in the vector space $\text{span}(C)$, it follows that in general: C is a basis for $\text{span}(C)$ if and only if C is linearly independent.

Remark. Notice that all the above results started from the *assumption* that there exists at least one maximal set of finite number of linearly independent vectors in our vector space. Such assumption is not necessarily true, and is the point of distinction between finite dimensionality and infinite dimensionality vector spaces. In this study we always assume *finite dimensionality*, although except for a small minority, all the results we will derive hold as well for vector spaces of infinite dimension.

Proposition 2.2.2. *Let $\mathscr{W}$ be a subspace of the vector space $\mathscr{V}$. For any basis $\mathscr{B}_w$ for $\mathscr{W}$, there exists a basis $\mathscr{B}_v$ for $\mathscr{V}$ such that $\mathscr{B}_w \subseteq \mathscr{B}_v$.*

Proof. $\mathscr{B}_w$ is either maximal in $\mathscr{V}$ or not. In the former case we are done. In the latter, by the essence of maximality, one can add vectors one by one to $\mathscr{B}_w$ without violating the linear independence property. At the point where one can add no more vectors, by proposition 2.2.1 the resulting set has exactly $\dim(\mathscr{V})$ members, and hence would be a basis for $\mathscr{V}$. □

Remark. Notice that what we proved in proposition 2.2.2 does not hold the other way around: Given $\mathscr{W}$ is a subspace of $\mathscr{V}$, for *any* arbitrary basis $\mathscr{B}_v$ for $\mathscr{V}$, we can *not* necessarily find a subset $\mathscr{B}_w$ of $\mathscr{B}_v$ that is a basis for $\mathscr{W}$.

Proposition 2.2.3 (Properties of dimensionality)**.** *Let $\mathscr{W}$ and $\mathscr{Y}$ be two subspaces of the vector space $\mathscr{V}$. From corollary 2.0.1 we know that $\mathscr{W} \cup \mathscr{Y}$ and $\mathscr{W} \cap \mathscr{Y}$ are also subspaces of $\mathscr{V}$. The following hold regarding their dimensionalities:*

(i) $\dim(\mathscr{W}) \leq \dim(\mathscr{V})$

(ii) $\dim(\mathscr{W}) = \dim(\mathscr{V}) \Leftrightarrow \mathscr{W} = \mathscr{V}$

(iii) $\dim(\mathscr{W} \cap \mathscr{Y}) \leq \min\big\{\dim(\mathscr{W}), \dim(\mathscr{Y})\big\}$

(iv) $\dim(\mathscr{W} \cup \mathscr{Y}) = \dim(\mathscr{W}) + \dim(\mathscr{Y}) - \dim(\mathscr{W} \cap \mathscr{Y})$

Proof.

(i) is obvious if we look at the definition of dimensionality as the shared cardinality of all maximal sets of linearly independent members.

(ii) follows immediately from proposition 2.2.2.

(iii) immediately results from (i), since $\mathscr{W} \cap \mathscr{Y}$ is a subspace of both $\mathscr{W}$ and $\mathscr{Y}$.

(iv) Let $\mathscr{B}_\cap$ be a basis for the subspace $\mathscr{W} \cap \mathscr{Y}$. From proposition 2.2.2 one can find a basis $\mathscr{B}_y$ for $\mathscr{W}$ and a basis $\mathscr{B}_y$ for $\mathscr{Y}$, such that $\mathscr{B}_\cap \subseteq \mathscr{B}_w$ and $\mathscr{B}_\cap \subseteq \mathscr{B}_y$. Furthermore it is easy to show that $\mathscr{B}_w \cap \mathscr{B}_y = \mathscr{B}_\cap$, and hence $\mathscr{B}_w \cup (\mathscr{B}_y - \mathscr{B}_\cap)$ is a basis for $\mathscr{W} \cup \mathscr{Y}$. The wanted result follows from elementary set theory.

□

In the same fashion as we proved part (iii) of the last proposition, one can prove the analogous of the inclusion exclusion principle in set theory for dimensionalty of vector spaces:

Proposition 2.2.4 (Dimensionality Inclusion and Exclusion)**.** *Let $\mathscr{W}_1, \mathscr{W}_2, \ldots, \mathscr{W}_t$ be subspaces of a vector space $\mathscr{V}$. We have:*

$$\dim(\bigcup_{i=1}^{t} \mathscr{W}_i) = \sum_{i=1}^{t} \dim(\mathscr{W}_i) - \sum_{i<j} \dim(\mathscr{W}_i \cap \mathscr{W}_j) + \sum_{i<j<k} \dim(\mathscr{W}_i \cap \mathscr{W}_j \cap \mathscr{W}_k) - \ldots + (-1)^{t+1} \dim(\bigcap_{i=1}^{t} \mathscr{W}_i)$$

or in more compact form:

$$\dim(\bigcup_{i=1}^{t} \mathscr{W}_i) = \sum_{l=1}^{k} (-1)^{l+1} \left[\sum_{1 \leq t_1 < t_2 < \ldots < t_l \leq t} \dim(\bigcap_{i=1}^{l} \mathscr{W}_{t_i}) \right]$$

Up to this point, we have been talking about the dimensionality of a vector space. A vector space can be thought of to be *linearly compact*, in the sense that all linear combinations of its members are also contained. But an arbitrary set of vectors from a vector space is not necessarily so, but it would still be useful to have a notion of dimensionality for such arbitrary sets of vectors. There are three equivalent ways for defining the dimensionality of an arbitrary set, which we first introduce, and then prove their equivalence.

Definition 2.2.4 (Dimensionality of arbitrary sets of vectors)**.** Let C be a set of vectors, possibly uncountably infinite, in the vector space $\mathscr{V}$. We define the dimensionality of C, to be any of the following equivalent values:

(i) $\dim(C) := \dim(\text{span}(C))$

(ii) $\dim(C) :=$ dimensionality of the smallest subspace of $\mathscr{V}$ that contains C:

$$\dim(C) := \min_{\substack{C \subseteq \mathscr{W} \\ \mathscr{W} \text{ subspace of } \mathscr{V}}} \dim(\mathscr{W})$$

(iii) $\dim(C) :=$ the cardinality of the largest linearly independent subset of C:

$$\dim(C) := \max_{\substack{A \subseteq C \\ A \text{ lin. indep.}}} |A|$$

To emphasize the distinction between the two notions of dimensionality (one for vector spaces, and for arbitrary sets of vectors from a vector space) we will use the notation $\dim^{\text{span}}(C)$ to refer to the latter, although the one we are referring to can always be inferred from the context, and furthermore the two are consistent (see the corollary below).

Corollary 2.2.3. From the above definition one can easily prove the following properties:

(i) By definition (i), if $|S|$ is well defined (C is finite), we have $\dim^{\text{span}}(C) \leq |S|$. Although if C is infinite, setting $|S|$ to be ∞ the inequality still holds, since the dimensionality is finite, so long as $\mathscr{V}$ has finite dimensionality.

(ii) By definition (i), and using the fact that the span of a subspace is itself, we will observe that if $\mathscr{W}$ is a *subspace* of $\mathscr{V}$, the two definitions of dimensionality are consistent:

$$\dim^{\text{span}}(\mathscr{W}) = \dim(\mathscr{W})$$

(iii) If $\mathscr{W}$ is a subspace of $\mathscr{V}$, then $\dim^{\text{span}}(\mathscr{V} - \mathscr{W})$ is either 0 or $\dim(\mathscr{V})$. In other words, if we strike all the members of the subspace $\mathscr{W}$ out of the vector space $\mathscr{V}$ they live in, no matter how big $\mathscr{W}$ is, as long as it is not equal to the whole $\mathscr{V}$, the linear hull of the remaining vectors is $\mathscr{V}$ (linear combinations of the remaining vectors will build up the whole vector space again).

Proposition 2.2.5. *The three different definitions of dimensionality of arbitrary sets of vectors given above are in fact equivalent.*

Proof. We first prove that the lowest dimensionality subspace of $\mathscr{V}$ that contains C, which we will call $\mathscr{W}$, coincides with $\text{span}(C)$, which will prove the equivalence of (i) and (ii). We first notice that since $S \subseteq \mathscr{W}$, by using two of the earlier corollaries we have: $\text{span}(C) \subseteq \text{span}(\mathscr{W}) = \mathscr{W}$. This leaves us with proving $\mathscr{W} - \text{span}(C) = \phi$. If there exists a vector in $\mathscr{W} - \text{span}(C)$, using 2.2.2 it immediately follows that $\dim(\mathscr{W}) > \text{span}(C)$, which is a contradiction, since we assumed $\mathscr{W}$ has the least dimensionality among all linear spaces containing C (of which $\text{span}(C)$ is one).

We now prove the equivalence of (i) and (iii). We first notice that since $S \subseteq \text{span}(C)$, by corollary 2.2.1, there exists no linearly independent subset of C containing more than $n = \dim(\text{span}(C))$ vectors. Now let's assume the largest linearly independent subset of C is C_1, and $|C_1| \leq n$. Since C_1 is the largest such set, it is maximal too. As a result adding any other member of C to C_1 will violate its linear independence and hence all the members of C are in $\text{span}(C_1)$ and therefore:

$$\begin{aligned} S \subseteq \text{span}(C_1) &\overset{\text{Cor. 2.1.2}}{\Longrightarrow} \text{span}(C) \subseteq \text{span}\left(\text{span}(C_1)\right) = \text{span}(C_1) \\ &\overset{\text{Prop. 2.2.3}}{\Longrightarrow} n = \dim(\text{span}(C)) \leq \dim(\text{span}(C_1)) \\ &\overset{\text{Cor. 2.2.2}}{\Longrightarrow} n \leq |C_1| \\ &\overset{|C_1| \geq n}{\Longrightarrow} |C_1| = n \end{aligned}$$

and the proof is complete. □

There exists a really important counter-intuitive property of $\dim^{\text{span}}$. Let $\mathscr{Y}$ be a subspace of $\mathscr{V}$ and $u \notin \mathscr{Y}$ be any vector. Now if we build $C = u + \mathscr{Y}$ to be a the set of all vectors resulting from adding a vector $y \in \mathscr{Y}$ to u, C would be in non-exact words, a shift of a whole subspace, and one would expect it to have the same dimensionality as $\mathscr{Y}$, which is *not* true:

Definition 2.2.5. Let $\mathscr{Y}$ be a subspace of the vector space $\mathscr{V}$ and let $u \notin \mathscr{Y}$ be any vector in $\mathscr{V}$. Define the addition $C = u + \mathscr{Y}$ (which is different from $\text{span}\{u\} + \mathscr{Y}$ which will be defined later in section 2.4) to be the following set:

$$u + \mathscr{Y} = \{y + u : y \in \mathscr{Y}\}$$

We call such set of vectors **shifted** subspaces.

Proposition 2.2.6. *A shifted subspace $C = u + \mathscr{Y}$ is* not *a subspace itself, unless $u \in \mathscr{Y}$. Furthermore, if $u \notin \mathscr{Y}$, we have:*

$$\dim^{\text{span}}(u + \mathscr{Y}) = \dim(\mathscr{Y}) + 1$$

Proposition 2.2.7. *Let $C = u + \mathscr{Y}$ be a shifted subspace in the vector space $\mathscr{V}$ whose carrier field is $\mathbb{F}$. Although C is not a subspace itself, but for any basis $\mathscr{B}$ for $\mathscr{V}$ its isomorphic set $[C]_{\mathscr{B}}$ in $\mathbb{F}^n$ is a compact set, and in case $\mathbb{F}$ is either $\mathbb{R}$ or $\mathbb{C}$, $[C]_{\mathscr{B}}$ is a convex set as well.*

Proof. Let $u_1, u_2 \in C$. For any $\lambda \in (0, 1)$ we would have:

$$\lambda u_1 + (1 - \lambda)u_2) = \lambda(u + y_1) + (1 - \lambda)(u + y_2) = u + [\lambda y_1 + (1 - \lambda)y_2] \in u + \mathscr{Y} = C$$

Since C and $[C]_{\mathscr{B}}$ are isomorphic, for any two vectors $[u_1]_{\mathscr{B}}, [u_2]_{\mathscr{B}} \in [C]_{\mathscr{B}}$, $\lambda[u_1]_{\mathscr{B}} + (1 - \lambda)[u_2]_{\mathscr{B}}$ also falls inside $[C]_{\mathscr{B}}$ which completes the proof for convexity. To see why $[C]_{\mathscr{B}}$ is closed as well, we assume the Euclidean metric over $\mathbb{F}^n$, and we notice that for any sequence of elements $\{[u_i]_{\mathscr{B}}\}_{i=1}^{\infty}$ in $[C]_{\mathscr{B}}$ converging to some $[u_\circ]_{\mathscr{B}}$ we have:

$$[u_\circ]_{\mathscr{B}} = \lim_{i \to \infty} [u_i]_{\mathscr{B}} = \lim_{i \to \infty} ([u]_{\mathscr{B}} + [y_i]_{\mathscr{B}}) = [u]_{\mathscr{B}} + [y_\circ]_{\mathscr{B}} \in [C]_{\mathscr{B}}$$

which completes the proof. □

2.3 The $\mathbb{F}^n$ vector space and Numerical identity of vectors

Since more often than not we are dealing with vector spaces over $\mathbb{R}$ or $\mathbb{C}$, we tend to think of scalars as *numbers*. But the vectors need not all have any numerical identity. They might have different numerical *representations* (as we will clarify now), but that should not be attached to their identity. Up to now we have not attached any numerical identity to vectors, and in fact, we will insist on not doing so. We, instead, prefer to think of vectors as abstract algebraic creatures, on which we have some operations with nice properties. But one special case exists, that clears up the relationship between vectors and numbers. For any field $\mathbb{F}$ one can think of the set of all n-tuples of $\mathbb{F}$ members:

$$v = \begin{bmatrix} \nu_1 \\ \nu_2 \\ \vdots \\ \nu_n \end{bmatrix}, \ \nu_i \in \mathbb{F}$$

Naturally calling the set of all such n-tuples $\mathbb{F}^n$, and equipping it with the following addition:

$$\begin{bmatrix} \nu_1 \\ \nu_2 \\ \vdots \\ \nu_n \end{bmatrix} + \begin{bmatrix} \eta_1 \\ \eta_2 \\ \vdots \\ \eta_n \end{bmatrix} = \begin{bmatrix} \nu_1 + \eta_1 \\ \nu_2 + \eta_2 \\ \vdots \\ \nu_n + \eta_n \end{bmatrix}$$

and the following scalar multiplication:

$$\alpha * \begin{bmatrix} \nu_1 \\ \nu_2 \\ \vdots \\ \nu_n \end{bmatrix} = \begin{bmatrix} \alpha\nu_1 \\ \alpha\nu_2 \\ \vdots \\ \alpha\nu_n \end{bmatrix}$$

one can easily check that $(\mathbb{F}^n, \mathbb{F}, *)$ is a vector space. We, in short, refer to this vector space by $\mathbb{F}^n$. In this case one can see that the zero of the vector space is automatically inherited from the carrier field $\mathbb{F}$:

$$0_{\mathbb{F}^n} = \begin{bmatrix} 0_{\mathbb{F}} \\ 0_{\mathbb{F}} \\ \vdots \\ 0_{\mathbb{F}} \end{bmatrix}$$

and also a *natural* basis is imposed by the structure of the vector space.

$$\mathscr{E} = \{e_1, e_2, \ldots, e_n\} = \left\{ \begin{bmatrix} 1_{\mathbb{F}} \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1_{\mathbb{F}} \\ \vdots \\ 0 \end{bmatrix}, \ldots, \begin{bmatrix} 0 \\ \vdots \\ 0 \\ 1_{\mathbb{F}} \end{bmatrix} \right\}$$

and hence $\mathbb{F}^n$ has dimensionality n. But these are all in the case when we are looking at the $\mathbb{F}^n$ vector space. Of course in this case, due to the nature of the vector space, vectors do have a natural numerical identity, and in fact, this is their only identity. But if we are talking about polynomials, functions, sets, and such mathematical objects as vectors, one would see that they might have numerical representations, but not numerical identities. We now prove that for any vector space $\mathscr{V}$ over the field $\mathbb{F}$, any basis $\mathscr{B}$ for $\mathscr{V}$ induces an isomorphism between $\mathscr{V}$ and $\mathbb{F}^{\dim(\mathscr{V})}$.

Proposition 2.3.1. *Let $\mathscr{V}$ be an n-dimensional vector space over the field $\mathbb{F}$. $\mathscr{V}$ is isomorphic to $\mathbb{F}^n$.*

Proof. Let $\mathscr{B}=\{b_1,\dots,b_n\}$ be a basis for $\mathscr{V}$. By definition, any vector $u\in\mathscr{V}$ has a *representation* in terms of $\mathscr{B}$:

$$u=\sum_{i=1}^{n}\alpha_i b_i$$

which, again by definition, is unique. Define the mapping $\varphi^{(\mathscr{B})}\colon \mathscr{V}\to\mathbb{F}^n$ to be the following:

$$u=\sum_{i=1}^{n}\alpha_i b_i \xrightarrow{\varphi^{(\mathscr{B})}} \begin{bmatrix}\alpha_1\\ \alpha_2\\ \vdots\\ \alpha_n\end{bmatrix}$$

Since $\mathscr{B}$ is a basis and thus spans all $\mathscr{V}$, this mapping is injective. One can easily observe that it is surjective as well. Also by uniqueness of representation, $\varphi^{(\mathscr{B})}$ is one-to-one. As a result if we can prove that $\varphi^{(\mathscr{B})}$ is a homomorphism, it, in fact, is an isomorphism. To prove the latter we need to show that $\varphi^{(\mathscr{B})}$ respects the algebraic structure of vector spaces. Specifically, since both $\mathscr{V}$ and $\mathbb{F}^n$ share their underlying field, we just have to show that the following hold for all $u,v\in\mathscr{V}$ and all $\alpha\in\mathbb{F}$:

$$\varphi^{(\mathscr{B})}(u+v)=\varphi^{(\mathscr{B})}(u)+\varphi^{(\mathscr{B})}(v)$$
$$\alpha\varphi^{(\mathscr{B})}(u)=\varphi^{(\mathscr{B})}(\alpha u)$$

both of which immediately follow from uniqueness of representation in terms of bases. □

Remark. It is easy to see that, for *any* basis $\mathscr{B}=\{b_1,\dots,b_n\}$ for $\mathscr{V}$, all b_i s are mapped by $\varphi^{(\mathscr{B})}$ to the *natural* basis members of $\mathbb{F}^n$:

$$[b_i]_{\mathscr{B}}=\begin{bmatrix}0\\0\\\vdots\\1\\0\\\vdots\\0\end{bmatrix}\ (i)$$

From now on, instead of using the notation $\varphi^{(\mathscr{B})}(u)$, where $\varphi^{(\mathscr{B})}$ is the isomorphism defined in the proof of proposition 2.3.1, we use $[u]_{\mathscr{B}}$ (or $[S]_{\mathscr{B}}$, or $[\mathscr{Y}]_{\mathscr{B}}$) to refer to the corresponding $\mathbb{F}^n$

vector (or set of vectors, or subspace) resulting from the representation of the vector u (or set C of vectors, or subspace $\mathscr{Y}$) in a basis $\mathscr{B}=\{b_1,\ldots,b_n\}$:

$$u=\sum_{i=1}^{n}\alpha_i b_i \Rightarrow [u]_{\mathscr{B}}=\begin{bmatrix}\alpha_1\\ \alpha_2\\ \vdots\\ \alpha_n\end{bmatrix}$$

One should however notice that this numerical representation that we attach to u (if we think of scalars as numbers), is *not* the identity of u, since it completely depends on the choice of the basis. In other words *any* vector in $\mathscr{V}$ could be represented by:

$$\begin{bmatrix}1\\ 0\\ \vdots\\ 0\end{bmatrix}$$

or any other n-tuple for that matter, if an appropriate basis is chosen.

Remark. The only element of $\mathscr{V}$ that always keeps its numerical representation is the zero vector, which is always, regardless of the choice of basis, mapped to the all 0 n-tuple.

Corollary 2.3.1. As a result of the existence of isomorphism(s) between an n-dimensional vector space $\mathscr{V}$ and $\mathbb{F}^n$, all the basis-independent properties of vectors in $\mathscr{V}$ are logically equivalent to the analogous properties for corresponding vectors in $\mathbb{F}^n$. For example for *any* basis $\mathscr{B}$ the following hold:

(i) $[\text{span}(C)]_{\mathscr{B}}=\text{span}([C]_{\mathscr{B}})$

(ii) The set $\{u_1,u_2,\ldots,u_k\}$ of vectors in $\mathscr{V}$ is linearly (in)dependent, if and only if the set

$$\{[u_1]_{\mathscr{B}},[u_2]_{\mathscr{B}},\ldots,[u_k]_{\mathscr{B}}\}$$

of vectors in $\mathbb{F}^n$ is linearly (in)dependent.

(iii) A vector v is linearly (in)dependent *of* a set C of vectors in $\mathscr{V}$, if and only if $[v]_{\mathscr{B}}$ is linearly (in)dependent of $[C]_{\mathscr{B}}$ in $\mathbb{F}^n$.

(iv) If $\mathscr{W}$ is a subspace of $\mathscr{V}$, then $[\mathscr{W}]_{\mathscr{B}}$ is a subspace of $\mathbb{F}^n$ and furthermore: $\dim(\mathscr{W})=\dim([\mathscr{W}]_{\mathscr{B}})$.

(v) The same holds for arbitrary sets of vectors: $\dim^{\text{span}}(C)=\dim^{\text{span}}([C]_{\mathscr{B}})$.

(vi) Two subspaces $\mathscr{W}_1$ and $\mathscr{W}_2$ of $\mathscr{V}$ intersect non-trivially if and only if $[\mathscr{W}_1]_{\mathscr{B}}$ and $[\mathscr{W}_2]_{\mathscr{B}}$ intersect non-trivially.

and so on and so forth ...

Let $\mathscr{B}=\{b_1,\ldots,b_n\}$ be a basis for $\mathscr{V}$, and $\mathscr{B}'=\{b'_1,\ldots,b'_n\}$ be another basis, whose members have the following representation in terms of $\mathscr{B}$:

$$\begin{aligned} b'_1 &= \beta_{1,1}b_1+\beta_{1,2}b_2+\ldots+\beta_{1,n}b_n \\ b'_2 &= \beta_{2,1}b_1+\beta_{2,2}b_2+\ldots+\beta_{2,k}b_n \\ \vdots & \qquad\qquad \vdots \\ b'_n &= \beta_{n,1}b_1+\beta_{n,2}b_2+\ldots+\beta_{t,k}b_n \end{aligned}$$

Now if an arbitrary $u\in\mathscr{V}$ has the following representation in terms of $\mathscr{B}'=\{b'_1,\ldots,b'_n\}$:

$$u=\sum_{i=1}^{n}\alpha_i b'_i$$

it follows that:

$$u=\sum_{i=1}^{n}\alpha_i b'_i=\sum_{i=1}^{n}\alpha_i\sum_{j=1}^{n}\beta_{i,j}b_j=\sum_{i=1}^{n}\sum_{j=1}^{n}\alpha_i\beta_{i,j}b_j=\sum_{j=1}^{n}\left(\sum_{i=1}^{n}\alpha_i\beta_{i,j}\right)b_j$$

which is the representation of u in terms of $\mathscr{B}$. This relationship relates the two $\mathbb{F}^n$ counterparts of u resulting from the two choices of basis to each other. One can see that the j-th entry of $[u]_{\mathscr{B}}$ is:

$$\sum_{i=1}^{n}\beta_{i,j}\alpha_i$$

Investigating this closely, one can formulate the change of basis as the multiplication by a matrix of scalars:

Proposition 2.3.2 (Change of basis)**.** *Let $\mathscr{V}$ be an n-dimensional vector space over the field $\mathbb{F}$. Let $\mathscr{B}$ and $\mathscr{B}'$ be two bases for $\mathscr{V}$. We define the the matrix $\mathrm{P}_{(\mathscr{B}'\to\mathscr{B})}$ of scalars as the following:*

$$\mathrm{P}_{(\mathscr{B}'\to\mathscr{B})}=\begin{bmatrix}[b'_1]_{\mathscr{B}} & [b'_2]_{\mathscr{B}} & \ldots & [b'_n]_{\mathscr{B}}\end{bmatrix}$$

then for any vector u in $\mathscr{V}$ we have $[u]_{\mathscr{B}}=\mathrm{P}_{(\mathscr{B}'\to\mathscr{B})}[u]_{\mathscr{B}'}$. We can more succinctly refer to $\mathrm{P}_{(\mathscr{B}'\to\mathscr{B})}$ by $[\mathscr{B}']_{\mathscr{B}}$ and write:

$$[u]_{\mathscr{B}}=[\mathscr{B}']_{\mathscr{B}}[u]_{\mathscr{B}'}$$

An easy way to look at the above rule is the following:

$$[u]_{\mathscr{B}'}=\begin{bmatrix}\alpha'_1\\ \alpha'_2\\ \vdots\\ \alpha'_n\end{bmatrix}$$

$$\begin{aligned} [\mathscr{B}']_{\mathscr{B}}[u]_{\mathscr{B}'} &= \begin{bmatrix}[b'_1]_{\mathscr{B}} & [b'_2]_{\mathscr{B}} & \ldots & [b'_n]_{\mathscr{B}}\end{bmatrix}\begin{bmatrix}\alpha'_1\\ \alpha'_2\\ \vdots\\ \alpha'_n\end{bmatrix} \\ &= \alpha'_1[b'_1]_{\mathscr{B}}+\alpha'_2[b'_2]_{\mathscr{B}}+\ldots+\alpha'_n[b'_n]_{\mathscr{B}} \\ &= [\alpha'_1b'_1+\alpha'_2b'_2+\ldots+\alpha'_nb'_n]_{\mathscr{B}}=[u]_{\mathscr{B}} \end{aligned}$$

Also we could notice that since:

$$[b_i]_{\mathscr{B}} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix} (i)$$

we will have:

$$[\mathscr{B}']_{\mathscr{B}}[b_i]_{\mathscr{B}'} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix} (i)$$

Corollary 2.3.2. As we just demonstrated, for any two bases $\mathscr{B}$ and $\mathscr{B}'$ for an n-dimensional vector space $\mathscr{V}$, and the matrices of scalars $\mathrm{P}_{(\mathscr{B}'\to\mathscr{B})}$ and $\mathrm{P}_{(\mathscr{B}\to\mathscr{B}')}$ defined as above, the following holds:

$$\mathrm{P}_{(\mathscr{B}'\to\mathscr{B})}\mathrm{P}_{(\mathscr{B}\to\mathscr{B}')} = \mathrm{P}_{(\mathscr{B}\to\mathscr{B}')}\mathrm{P}_{(\mathscr{B}'\to\mathscr{B})} = \mathrm{I}_n$$

$$[\mathscr{B}']_{\mathscr{B}}[\mathscr{B}]_{\mathscr{B}'} = [\mathscr{B}]_{\mathscr{B}'}[\mathscr{B}']_{\mathscr{B}} = \mathrm{I}_n$$

where I_n is the $n \times n$ identity matrix.

2.4 Direct Sums

We saw that a set of vectors in a vector space $\mathscr{V}$, if linearly compact (a subspace), i.e. closed under vector addition and scalar multiplication, would admit desirable properties that arbitrary sets of vectors do no admit. A subspace will have a well defined dimensionality (with no need to be embedded in a larger set), and admits the notion of bases, in the sense that there exists a bunch of vectors, such that *every* vector in the set would reduce to a *unique* linear combination of the basis vectors. We here adopt a new perspective towards the notion of representation and decomposition to basis vectors. Instead of looking at a basis $\{u_1, u_2, \ldots, u_n\}$ for the vector space $\mathscr{V}$, as a set of $\dim(\mathscr{V})$ vectors, in terms of which any vector has a unique representation of the form:

$$v = \alpha_1 u_1 + \alpha_2 u_2 + \ldots + \alpha_n u_n$$

we look at the n respective subspaces of dimensionality 1, namely $\mathscr{W}_i = \mathrm{span}(u_i)$, in terms of which any vector has a unique representation of the form:

$$v = w_1 + w_2 + \ldots + w_n \qquad s.t. \quad w_i \in \mathscr{W}_i \tag{2.4.1}$$

And of course for any v, all the w_i will be unique and all distinct since $\mathscr{W}_i$ intersect only trivially. In this new sense, of course, the notion of *decomposition* is more natural since we are breaking up a vector space $\mathscr{V}$, to objects of the same type, $\mathscr{W}_i$, and not to objects of another type, vectors. We already use the terminology "$\mathscr{B}=\{b_1,\ldots,b_n\}$ is a basis for $\mathscr{V}$" to talk about a set of *vectors*, to linear combinations of which the vector space decomposes, and we proved that any such basis, has to have exactly the same number of members. We now introduce a new terminology for the new perspective, the need for which will be obvious shortly. To denote the fact that $\mathscr{V}$ has been decomposed to vector spaces $\mathscr{W}_1,\mathscr{W}_2,\ldots,\mathscr{W}_n$ in the sense that any vector $v\in\mathscr{V}$ has a unique representation of the form 2.4.1, we write:

$$\mathscr{V}=\mathscr{W}_1\oplus\mathscr{W}_2\oplus\ldots\oplus\mathscr{W}_n$$

There is no reason we would restrict ourselves only to decompositions to subspaces of dimensionality 1, which by the way does not offer any new advantage, since they are equal to bases:

Definition 2.4.1. Let $\mathscr{V}$ be a vector space, and $\mathscr{W}_1,\mathscr{W}_2,\ldots,\mathscr{W}_k$ be arbitrary non-trivial subspaces (having dimensionality at least 1) of $\mathscr{V}$. We say that $\mathscr{V}$ **directly decomposes** to $\{\mathscr{W}_i\}_{i=1}^k$ if any vector $v\in\mathscr{V}$ has a *unique* representation as the sum of exactly k vectors, one from each of $\mathscr{W}_i$:

$$v=w_1+w_2+\ldots+w_k \qquad s.t. \quad w_i\in\mathscr{W}_i$$

and we write:

$$\mathscr{V}=\mathscr{W}_1\oplus\mathscr{W}_2\oplus\ldots\oplus\mathscr{W}_k$$

or in short, $\mathscr{V}=\bigoplus_{i=1}^k\mathscr{W}_i$, and read it as "$\mathscr{V}$ is equal to the **direct sum** of $\mathscr{W}_i$".

Of course for being able to decompose $\mathscr{V}$ completely, $\mathscr{W}_i$ have to satisfy some properties:

Proposition 2.4.1. *Let $\mathscr{V}$ be a vector space, of which $\mathscr{W}_1,\mathscr{W}_2,\ldots,\mathscr{W}_k$ are arbitrary non-trivial subspaces. We have $\mathscr{V}=\bigoplus_{i=1}^k\mathscr{W}_i$ if and only if both the following hold:*

(i) *Any two of the $\mathscr{W}_i$ s are linearly independent of each other (and since their span is themselves again, any two of them should intersect only trivially), i.e for any $i\neq j$ we must have:*

$$\mathscr{W}_i\cap\mathscr{W}_j=\{0\}$$

(ii) *The dimensions of $\mathscr{W}_i$ s add up to that of $\mathscr{V}$:*

$$\sum_{i=1}^k\dim(\mathscr{W}_i)=\dim(\mathscr{V})$$

Proof. If the direct sum decomposition $\mathscr{V}=\bigoplus_{i=1}^k\mathscr{W}_i$ holds and for some $i\neq j$ there exists a vector v in $\mathscr{W}_i\cap\mathscr{W}_j$, the zero vector would immediately have endless distinct representations other than its trivial representation: $0=w_i+w_j$ for any $w_i=\alpha v$ and $w_j=-\alpha v$. Also we notice that if $\mathscr{B}_1,\mathscr{B}_2,\ldots,\mathscr{B}_k$ are bases for $\mathscr{W}_i$, from the fact that any two $\mathscr{W}_i,\mathscr{W}_j$ are linearly independent, we can infer that the (1) the $\mathscr{B}_i$ are disjoint, and (2) $\bigcup_{i=1}^k\mathscr{B}_i$ is a linearly independent set of vectors in $\mathscr{V}$. Now we notice that the direct sum decomposition $\mathscr{V}=\bigoplus_{i=1}^k\mathscr{W}_i$ implies:

$$\mathscr{V}=\text{span}(\mathscr{B})$$

where $\mathscr{B} = \bigcup_{i=1}^{k} \mathscr{B}_i$. Now since $\mathscr{B}$ is a set of linearly independent vectors in $\mathscr{V}$ that spans the whole vector space, we immediately infer that:

$$\dim(\mathscr{V}) = |\mathscr{B}| = \sum_{i=1}^{k} |\mathscr{B}_i| = \sum_{i=1}^{k} \dim(\mathscr{W}_i)$$

We now prove that if $\mathscr{W}_1, \mathscr{W}_2, \ldots, \mathscr{W}_k$ are two by two linearly independent of each other (intersect only trivially) and their dimensions add up to $\dim(\mathscr{V})$, they in fact decompose $\mathscr{V}$. Again we introduce the bases $\mathscr{B}_i$ for $\mathscr{W}_i$, and notice that from linear independence of $\mathscr{W}_i$ of each other, $\mathscr{B}_i$ are disjoint, and hence defining $\mathscr{B} = \bigcup_{i=1}^{k} \mathscr{B}_i$, we have:

$$|\mathscr{B}| = \sum_{i=1}^{k} |\mathscr{B}_i| = \sum_{i=1}^{k} \dim(\mathscr{W}_i) \overset{\text{(ii)}}{=} \sum_{i=1}^{k} \dim(\mathscr{V})$$

Again using the fact that $\mathscr{W}_i$ intersect only trivially we infer that $\mathscr{B}$ as a whole is a linearly independent set of vectors in $\mathscr{V}$. Combining this with the fact that $\mathscr{B}$ consists of exactly $\dim(\mathscr{V})$ vectors, we get that $\mathscr{B}$ is a basis for $\mathscr{V}$:

$$\mathscr{V} = \operatorname{span}(\mathscr{B})$$

As a result any vector $v \in \mathscr{V}$ has a unique representation in terms of $\mathscr{B}$:

$$v = \alpha_1 b_1 + \alpha_2 b_2 + \ldots + \alpha_n b_n$$

We now remember that $\mathscr{B}$ is the result of the union of the disjoint sets $\mathscr{B}_i$. Hence $\dim(\mathscr{W}_1)$ of b_i s belong to $\mathscr{B}_1$, $\dim(\mathscr{W}_2)$ of them belong to $\mathscr{B}_2$, and so on. Grouping corresponding terms of each $\mathscr{W}_i$ in the linear combination $u = \sum_{i=1}^{k} \alpha_i b_i$, we get a representation of v of the form:

$$v = w_1 + w_2 + \ldots + w_k \qquad s.t. \quad w_i \in \mathscr{W}_i$$

the uniqueness of which follows immediately from uniqueness of $\alpha_1, \alpha_2, \ldots, \alpha_n$. □

Corollary 2.4.1. Obviously if $\mathscr{W}_1, \mathscr{W}_2, \ldots, \mathscr{W}_k$ are non-trivial subspaces to which $\mathscr{V}$ directly decomposes, from (ii) in the last proposition it follows that $k \leq \dim(\mathscr{V})$, since all the dimensionalities of $\mathscr{W}_i$ are at least 1.

Proposition 2.4.2. *Let $\mathscr{V}$ be a vector space. Let $\mathscr{B}$ be an arbitrary basis for $\mathscr{V}$. Any splitting of $\mathscr{B}$ to k disjoint subsets $\{\mathscr{B}_i\}_{i=1}^{k}$ induces a direct sum decomposition of $\mathscr{V}$:*

$$\mathscr{V} = \operatorname{span}(\mathscr{B}_1) \oplus \operatorname{span}(\mathscr{B}_2) \oplus \ldots \oplus \operatorname{span}(\mathscr{B}_k)$$

Similarly given a direct sum decomposition of $\mathscr{V}$:

$$\mathscr{V} = \mathscr{W}_1 \oplus \mathscr{W}_2 \oplus \ldots \oplus \mathscr{W}_k$$

for any set of bases $\{\mathscr{B}_i\}_{i=1}^{k}$ for $\mathscr{W}_i$, the disjoints union $\bigcup_{i=1}^{k} \mathscr{B}_i$ is a basis for $\mathscr{V}$.

Corollary 2.4.2. Combining this result with that of proposition 2.2.2, it follows that for any subspace $\mathscr{W}$ of $\mathscr{V}$ there exists another subspace $\mathscr{Y}$ of $\mathscr{V}$ that is linearly independent of $\mathscr{W}$ (i.e. $\mathscr{W} \cap \mathscr{Y} = \{0\}$) and: $\mathscr{V} = \mathscr{W} \oplus \mathscr{Y}$.

Although up to now we have only used the decomposition to sum of subspaces, only for when the decomposition is direct, i.e. the representation of vectors in terms of vectors in the summands is unique, it would be reasonable to define a more relaxed version of decomposition, which would not require the summands to be linearly independent.

Definition 2.4.2. Let $\mathscr{V}$ be a vector space, and $\mathscr{W}_1, \mathscr{W}_2, \ldots, \mathscr{W}_m$ be arbitrary non-trivial subspaces (having dimensionality at least 1) of $\mathscr{V}$. We say that $\mathscr{V}$ **decomposes** (*not necessarily directly*) to $\{\mathscr{W}_i\}_{i=1}^m$ if any vector $v \in \mathscr{V}$ has a representation (*not necessarily unique*) as the sum of exactly m vectors, one from each of $\mathscr{W}_i$:

$$v = w_1 + w_2 + \ldots + w_m \qquad s.t. \quad w_i \in \mathscr{W}_i$$

and we write:

$$V = \mathscr{W}_1 + \mathscr{W}_2 + \ldots + \mathscr{W}_m$$

Obviously direct sums are special cases of the above definition. If $\mathscr{V}$ decomposes to the sum (in the general sense of the above definition) of $\mathscr{W}_i$, the sum of dimensions of $\mathscr{W}_i$ could possibly exceed $\dim(\mathscr{V})$, also m could possibly be greater than $\dim(\mathscr{V})$.

Chapter 3

Linear Transformations

We treat linear transformations in their pure algebraic sense, and differentiate between the identity of a linear transformation and its matrix representation, and prefer to prove properties as independent of matrix representations as possible. In that case the properties of matrices would just flow immediately from respective probabilities of linear transformations.

3.1 Basic properties

For defining linear transformations (or linear maps), some authors define them to be mappings from one vector space $\mathscr{V}$ to another vector space $\mathscr{W}$ (possibly coinciding with $\mathscr{V}$) satisfying the following properties:

Definition 3.1.1. For any two arbitrary vector spaces $\mathscr{V}$ and $\mathscr{W}$ that share the same carrier field $\mathbb{F}$, a mapping $\mathrm{T}\colon \mathscr{V}\to\mathscr{W}$ is a **linear transformation** if it is *linear* in the following sense:

(i) $\mathrm{T}(\alpha u) = \alpha \mathrm{T}u$.

(ii) $\mathrm{T}(u+v) = \mathrm{T}u + \mathrm{T}v$.

We also use the notation $\mathrm{T}C$ to refer to the image of T restricted to a set C of vectors in $\mathscr{V}$.

And according to those authors, using this definition a bunch of properties immediately follow from the above definition, for example:

$$\mathrm{T}0 = 0$$

$$\mathrm{T}(\sum_{i=1}^{k} \alpha_i u_i) = \sum_{i=1}^{k} \alpha_i \mathrm{T}u_i$$

and so on. But it would be worthwhile to note that the above definition is just the definition of a vector space homomorphism (see section 1.3). One can easily see how *respecting* the algebraic structure of vector spaces, reduces to the above two criteria in the special case where the departure and destination vector spaces share their carrier fields, and hence the operation of T is limited to mapping vectors to vectors, leaving scalars untouched. From now on when we say T is a linear transformation from $\mathscr{V}$ to $\mathscr{W}$, we implicitly assume the two share their carrier fields.

Definition 3.1.2. Let $\mathrm{T}\colon \mathscr{V} \to \mathscr{W}$ be a linear transformation. We define the **image** of T as usual to be the following subset of $\mathscr{W}$:

$$\mathrm{Im}[\mathrm{T}] := \mathrm{T}\mathscr{V} = \{w \in \mathscr{W} : \exists u \in \mathscr{V}, w = \mathrm{T}u\}$$

and the **kernel** or the **null space** of T to be the following subset of $\mathscr{V}$:

$$\mathrm{Ker}[\mathrm{T}] := \{u \in \mathscr{V} : \mathrm{T}u = 0\}$$

Proposition 3.1.1. *Let $\mathscr{V}, \mathscr{W}, \mathscr{Y}$ be three subspaces sharing their carrier field. Let $\mathrm{T}\colon \mathscr{V} \to \mathscr{W}$ and $\mathrm{S}\colon \mathscr{W} \to \mathscr{Y}$ be two linear transformations. Obviously* $\mathrm{Ker}[\mathrm{T} \circ \mathrm{S}]$ *is a subspace of $\mathscr{V}$ and* $\mathrm{Im}[\mathrm{T} \circ \mathrm{S}]$ *is a subspace of $\mathscr{Y}$. The following will hold:*

(i) $\mathrm{Ker}[\mathrm{T}] \subseteq \mathrm{Ker}[\mathrm{T} \circ \mathrm{S}]$.

(ii) $\mathrm{Im}[\mathrm{T}] \supseteq \mathrm{Im}[\mathrm{T} \circ \mathrm{S}]$

Corollary 3.1.1. In the special case where $\mathrm{T}\colon \mathscr{V} \to \mathscr{V}$ we will have the following:

(i) $\mathrm{Ker}[\mathrm{T}] \subseteq \mathrm{Ker}[\mathrm{T} \circ \mathrm{T}] \subseteq \mathrm{Ker}[\mathrm{T} \circ \mathrm{T} \circ \mathrm{T}] \subseteq \ldots$.

(ii) $\mathrm{Im}[\mathrm{T}] \supseteq \mathrm{Im}[\mathrm{T} \circ \mathrm{T}] \supseteq \mathrm{Im}[\mathrm{T} \circ \mathrm{T} \circ \mathrm{T}] \supseteq \ldots$

Proposition 3.1.2. *Let $\mathrm{T}\colon \mathscr{V} \to \mathscr{W}$ be a linear transformation. The image of* T *restricted to a subspace $\mathscr{Y}$ of $\mathscr{V}$, is a subspace of $\mathscr{W}$, i.e for any subspace $\mathscr{Y}$ of $\mathscr{V}$, $\mathrm{T}\mathscr{Y}$ is a subspace of $\mathscr{W}$.*

Proof. We just need to prove compactness of $\mathrm{T}\mathscr{Y}$. We first notice that for any $v \in \mathrm{T}\mathscr{Y}$, there is some $u \in \mathscr{Y}$ such that $v = \mathrm{T}u$. It follows that $\alpha v = \alpha \mathrm{T}(u) = \mathrm{T}(\alpha u)$, which is again in $\mathrm{T}\mathscr{Y}$ since $\mathscr{Y}$ is a subspace. Also if $v_1, v_2 \in \mathrm{T}\mathscr{Y}$ we have

$$v_i = \mathrm{T}u_i \Rightarrow v_1 + v_2 = \mathrm{T}(u_1 + u_2)$$

and the proof is complete. □

Proposition 3.1.3. *If $\mathrm{T}\colon \mathscr{V} \to \mathscr{W}$ is a linear transformation, then the image (in $\mathscr{W}$) of the span of a set of vectors in $\mathscr{V}$, is the span of the image of the vectors individually:*

$$\mathrm{span}(\mathrm{T}C) = \mathrm{T}(\mathrm{span}(C))$$

Proof. Let C be the set $\{u_1, u_2, \ldots\}$. We first prove that $\mathrm{span}(\mathrm{T}C) \subseteq \mathrm{T}(\mathrm{span}(C))$, and next the other way around. For any vector $v \in \mathrm{span}(\mathrm{T}C)$ we prove that $v \in \mathrm{T}(\mathrm{span}(C))$. We notice that:

$$\mathrm{T}C = \{\mathrm{T}u_1, \mathrm{T}u_2, \ldots\}$$

and hence

$$v = \sum_{i=1}^{k} \alpha_i \mathrm{T}u_i = \mathrm{T}(\sum_{i=1}^{k} \alpha_i u_i) = \mathrm{T}w$$

for some $w \in \mathrm{span}(C)$ and thus $v \in \mathrm{T}(\mathrm{span}(C))$. Now let's assume $v \in \mathrm{T}(\mathrm{span}(C))$ which translates to:

$$v = \mathrm{T}(\sum_{i=1}^{k} \alpha_i u_i) = \sum_{i=1}^{k} \alpha_i \mathrm{T}u_i$$

and hence $v \in \mathrm{span}(\mathrm{T}C)$. □

Corollary 3.1.2. Let $\mathrm{T}: \mathscr{V} \to \mathscr{W}$ be a linear transformation. Let C be an arbitrary set of vectors in $\mathscr{V}$, and let $\mathscr{B}$ be any basis for $\mathrm{span}(C)$. We notice that $\mathrm{span}(C) = \mathrm{span}(\mathscr{B})$ and that by proposition 3.1.3 we have:

$$\mathrm{span}(\mathrm{T}C) = \mathrm{T}(\mathrm{span}(C))$$

As a result we will have:

$$\mathrm{span}(\mathrm{T}\mathscr{B}) = \mathrm{T}(\mathrm{span}(\mathscr{B})) = \mathrm{T}(\mathrm{span}(C)) = \mathrm{span}(\mathrm{T}C)$$

Proposition 3.1.4. *For any linear transformation* $\mathrm{T} : \mathscr{V} \to \mathscr{W}$*, the sets* Im[T] *and* Ker[T] *are vector spaces, i.e subspaces of* $\mathscr{W}$ *and* $\mathscr{V}$ *respectively.*

Definition 3.1.3. Let $\mathrm{T} : \mathscr{V} \to \mathscr{W}$ be a linear transformation. From 3.1.4 we know that Im[T] and Ker[T] are subspaces of $\mathscr{W}$ and $\mathscr{V}$ respectively. We thus define the **rank** of T to be the dimensionality of its image:

$$\rho(\mathrm{T}) = \dim(\mathrm{Im}[\mathrm{T}]) \leq \dim(\mathscr{W})$$

and the **nullity** of T to be the dimensionality of its kernel:

$$\eta(\mathrm{T}) = \dim(\mathrm{Ker}[\mathrm{T}]) \leq \dim(\mathscr{V})$$

An interesting question would be to see how linear transformations treat linear independence.

Proposition 3.1.5. *Let* $\mathrm{T}: \mathscr{V} \to \mathscr{W}$ *be a linear transformation. For any subspace* $\mathscr{Y}$ *of* $\mathscr{V}$ *we have:*

$$\dim(\mathrm{T}\mathscr{Y}) \leq \dim(\mathscr{Y})$$

Similarly for any arbitrary set C *of vectors in* $\mathscr{V}$ *we have:*

$$\dim^{\mathrm{span}}(\mathrm{T}C) \leq \dim^{\mathrm{span}}(C)$$

Proof. It follows from proposition 3.1.3 that the two versions (set vs subspace) are equivalent. We will thus prove the claim for subspaces. Let $\dim(\mathscr{Y}) = k$. We will prove here that any set of $k+1$ vectors in $\mathrm{T}\mathscr{Y}$ is linearly dependent, which will immediately prove the claim since in that case we would have $\dim(\mathrm{T}\mathscr{Y}) < k+1$. Let $\{b_1, b_2, \ldots, b_{k+1}\}$ be a set of $k+1$ vectors in $\mathrm{T}\mathscr{Y}$. There obviously exists $u_1, \ldots, u_{k+1}$ such that $b_i = \mathrm{T}u_i$. The set $\{u_1, \ldots, u_{k+1}\}$ obviously is linearly dependent, since they are all in a k-dimensional vector space. Thus there exist scalars $\alpha_1, \ldots, \alpha_{k+1}$, not all zero, such that:

$$\sum_{i=1}^{k+1} \alpha_i u_i = 0 \Rightarrow \mathrm{T}(\sum_{i=1}^{k+1} \alpha_i u_i) = \sum_{i=1}^{k+1} \alpha_i \mathrm{T}u_i = 0 \Rightarrow \sum_{i=1}^{k+1} \alpha_i b_i = 0$$

and since this holds for any set $\{b_1, \ldots, b_{k+1}\}$ of vectors in $\mathrm{T}\mathscr{Y}$, the proof is complete. □

Remark. Why does the flow of arguments we just mentioned not hold the other way around, so we could prove $\dim(\mathrm{T}\mathscr{Y}) \geq \dim(\mathscr{Y})$ as well? Let $\dim(\mathrm{T}\mathscr{Y}) = t$ and the same as in the proof, let $u_1, u_2, \ldots, u_{t+1}$ be vectors in $\mathscr{Y}$ and we try to prove they are linearly dependent. The problem in this direction is that from the fact that, by definition of dimensionality for $\mathrm{T}\mathscr{Y}$, there exists scalars $\alpha_1, \ldots, \alpha_{t+1}$, not all zero, such that:

$$\sum_{i=1}^{k+1} \alpha_i \mathrm{T}u_i = \mathrm{T}(\sum_{i=1}^{k+1} \alpha_i u_i) = 0$$

it does not follow that $\sum_{i=1}^{k+1} \alpha_i u_i$ is zero.

From part (iii) of corollary 2.2.3 we remember that as long as Ker[T] has not swallowed all $\mathscr{V}$, the mere fact that $u_1, u_2, \ldots, u_k$ are not individually inside the kernel does not mean no linear combination of them would fall in the kernel as well. However if the set $\{u_1, u_2, \ldots, u_k\}$ as a set, is linearly independent of the kernel, it would by definition mean that all members of $\text{span}\{u_i\}_{i=1}^k$ (all linear combinations of u_i) would as well stir clear of the kernel. In fact, this is the only way a subspace $\mathscr{Y}$ of $\mathscr{V}$ could sustain its dimensionality after being fed to T, where by $\mathscr{Y}$ sustaining its dimensionality under T we mean:

$$\dim(\mathscr{Y}) = \dim(\mathrm{T}\mathscr{Y})$$

We could also think of an arbitrary set C of vectors sustaining their dimensionality under T:

$$\dim^{\text{span}}(C) = \dim^{\text{span}}(\mathrm{T}C)$$

Proposition 3.1.6. *Let* $\mathrm{T}: \mathscr{V} \to \mathscr{W}$ *be a linear transformation. Any subspace* $\mathscr{Y}$ *of (arbitrary set of vectors* C *in)* $\mathscr{V}$ *sustains its dimensionality under* T *if and only if the subspace* $\mathscr{Y}$ *(the set* C*) is linearly independent, in the set-set sense, of* Ker[T]*, i.e:*

$$\mathscr{Y} \cap \text{Ker[T]} = \{0\}$$

$$(\text{span}(C) \cap \text{Ker[T]} = \{0\})$$

Proof. The proof is exactly similar for a subspace and an arbitrary set of vectors, we will thus show it for sets. First, let C be linearly independent of Ker[T]. As a result the only linear combination of members of C that falls into the kernel is the trivial all zero linear combination. Now let $\dim^{\text{span}}(C) = k$ and let $u_1, u_2, \ldots, u_k$ be k members of C that form a linearly independent set. We claim that their respective images $\mathrm{T}u_i$ will be linearly independent as well:

$$\sum_{i=1}^{k} \alpha_i \mathrm{T}u_i = 0 \Rightarrow \mathrm{T}(\sum_{i=1}^{k} \alpha_i u_i) = 0 \Rightarrow \sum_{i=1}^{k} \alpha_i u_i \in \text{Ker[T]} \Rightarrow \forall \alpha_i = 0$$

where the last derivation was possible because C is linearly independent of Ker[T]. As a result $\mathrm{T}C$ has dimensionality at least k. Combining this with proposition 3.1.5 completes the first part of the proof. We will now prove that if C sustains its dimensionality it must be linearly independent of Ker[T]. Let's assume C sustains it dimensionality, but has a non-zero linear combination in its span, that falls in the kernel:

$$v = \sum_{i=1}^{k} \alpha_i u_i \in \text{Ker[T]}$$

From proposition 2.2.2 we know that there exists a basis for $\text{span}(C)$ containing v, namely $\mathscr{B} = \{b_1 = v, b_2, \ldots, b_k\}$ where $k = \dim^{\text{span}}(C)$. We now look at $\mathrm{T}\mathscr{B}$. Since at least one of the members of $\mathrm{T}\mathscr{B}$ is zero ($\mathrm{T}v$=0), by corollary 2.2.3 we have:

$$\dim^{\text{span}}(\mathrm{T}\mathscr{B}) \leq k - 1$$

we also notice that by corollary 3.1.2 we have:

$$\text{span}(\mathrm{T}C) = \text{span}(\mathrm{T}\mathscr{B})$$

and hence:

$$\dim^{\text{span}}(\mathrm{T}C)) = \dim(\text{span}(\mathrm{T}C)) = \dim(\text{span}(\mathrm{T}\mathscr{B})) = \dim^{\text{span}}(\mathrm{T}\mathscr{B})) < k = \dim^{\text{span}}(C)$$

and the proof is complete. □

Now let $\mathscr{Y}$ be a subspace of $\mathscr{V}$ that intersects Ker[T] non-trivially. We know by now that dimensionality of $\mathrm{T}\mathscr{Y}$ is strictly smaller than that of $\mathscr{Y}$, but how much less? Intuitively we might think this has to have something to do with the size of $\mathscr{Y}$'s intersection with Ker[T], a measure of which could be $\dim(\mathscr{Y} \cap \mathrm{Ker}[\mathrm{T}])$. The following proposition proves this intuitive guess, and has propositions 3.1.6 and 3.1.5 as its special case:

Proposition 3.1.7. *Let* $\mathrm{T}\colon \mathscr{V} \to \mathscr{W}$ *be a linear transformation, and* $\mathscr{Y}$ *be an arbitrary subspace of* $\mathscr{V}$*. The the following holds:*

$$\dim(\mathrm{T}\mathscr{Y}) = \dim(\mathscr{Y}) - \dim(\mathscr{Y} \cap \mathrm{Ker}[\mathrm{T}])$$

Proof. Let $\mathscr{B}^\circ$ be a basis for $\mathscr{Y} \cap \mathrm{Ker}[\mathrm{T}]$. By proposition 2.2.2 we know that there exists a basis $\mathscr{B}$ for $\mathscr{Y}$ such that $\mathscr{B}^\circ \subseteq \mathscr{B}$. By letting $k = \dim(\mathscr{Y})$ and $k^\circ = \dim(\mathscr{Y} \cap \mathrm{Ker}[\mathrm{T}])$, we can name members of $\mathscr{B}$ accordingly:

$$\mathscr{B}^\circ = \{b_1, b_2, \ldots, b_{k^\circ}\}$$

$$\mathscr{B} = \{b_1, b_2, \ldots, b_{k^\circ}, b_{k^\circ+1}, \ldots, b_k\}$$

and the claim is that $\dim(\mathrm{T}\mathscr{Y}) = k - k^\circ$. We first notice that by defining $\mathscr{Z}$ to be the following subspace of $\mathscr{Y}$:

$$\mathscr{Z} = \text{span}\{b_i\}_{i=k^\circ+1}^{k}$$

by proposition 3.1.6 we can say that:

$$\dim(\mathrm{T}\mathscr{Z}) = \dim(\mathscr{Z})$$

which results in the following:

$$\dim(\mathrm{T}\mathscr{Y}) \geq \dim(\mathrm{T}\mathscr{Z}) = \dim(\mathscr{Z}) = k - k^\circ$$

and leaves us with just proving $\dim(\mathrm{T}\mathscr{Y}) \leq k - k^\circ$. To prove this we prove that any linearly independent set of vectors in $\mathrm{T}\mathscr{Y}$ has at most $k - k^\circ$ members. Let $\mathrm{T}u_1, \mathrm{T}u_2, \ldots, \mathrm{T}u_t \in \mathrm{T}\mathscr{Y}$ be linearly independent. Of course $u_1, u_2, \ldots, u_t$ are linearly independent as well:

$$\sum_{i=1}^{t} \alpha_i u_i = 0 \Rightarrow \sum_{i=1}^{t} \alpha_i \mathrm{T}u_i = 0$$

We use proposition 3.1.6 again, this time for the subspace $\mathscr{X} = \text{span}\{u_i\}_{i=1}^{t}$. Since both u_i and $\mathrm{T}u_i$ are linearly independent, $\mathscr{X}$ has sustained its dimensionality under T. Thus by proposition 3.1.6, $\mathscr{X}$ has to be linearly independent of Ker[T]. We know look at $\mathscr{X} \oplus \mathrm{Ker}[\mathrm{T}]$. By proposition 2.4.1 we have:

$$\dim(\mathscr{X} \oplus \mathrm{Ker}[\mathrm{T}]) = \dim(\mathscr{X}) + \dim(\mathrm{Ker}[\mathrm{T}]) = t + \eta(\mathrm{T})$$

We now notice that $\mathscr{X} \oplus \text{Ker[T]}$ is a subspace of $\mathscr{Y} \cup \text{Ker[T]}$ and the dimensionality on the left hand side is no larger than $\dim(\mathscr{Y} \cup \text{Ker[T]})$, which in its own turn according to proposition 2.2.3 satisfies:

$$\dim(\mathscr{Y} \cup \text{Ker[T]}) = \dim(\mathscr{Y}) + \dim(\text{Ker[T]}) - \dim(\mathscr{Y} \cap \text{Ker[T]}) = k + \eta(\text{T}) - k^\circ$$

combining the two together we get:

$$\dim(\mathscr{X} \oplus \text{Ker[T]}) = t + \eta(\text{T}) \leq \dim(\mathscr{Y} \cup \text{Ker[T]}) = k + \eta(\text{T}) - k^\circ$$

and hence $t \leq k - k^\circ$ and the proof is complete. □

Corollary 3.1.3 (Rank-Nullity Theorem). In the special case where $\mathscr{Y} = \mathscr{V}$ the result we just proved will be:

$$\dim(\text{T}\mathscr{V}) = \dim(\mathscr{V}) - \dim(\mathscr{V} \cap \text{Ker[T]})$$

But we know that $\mathscr{V} \cap \text{Ker[T]} = \text{Ker[T]}$, and that $\text{T}\mathscr{V} = \text{Im[T]}$. We also have defined the rank and nullity of T to be $\dim(\text{Im[T]})$ and $\dim(\text{Ker[T]})$, respectively. Thus:

$$\rho(\text{T}) + \eta(\text{T}) = \dim(\mathscr{V})$$

which is known as the Rank-Nullity theorem.

We have seen in corollary 3.1.2 that if $\mathscr{Y}$ is a subspace of $\mathscr{V}$ and $\mathscr{B}$ is a basis for $\mathscr{Y}$, we have:

$$\text{T}\mathscr{Y} = \text{span}(\text{T}\mathscr{B})$$

Thus for any such basis, we have access to $\dim(\mathscr{Y})$ vectors in $\mathscr{W}$, that fully span $\text{T}\mathscr{Y}$ and no more. But we have also seen in proposition 3.1.7 that this number of vectors could be redundant, since the dimensionality of $\text{T}\mathscr{Y}$ is possibly smaller than $\dim(\mathscr{Y})$:

$$\dim(\text{T}\mathscr{Y}) = \dim(\mathscr{Y}) - \dim(\mathscr{Y} \cap \text{Ker[T]})$$

Of course if $\mathscr{B}$ was picked such that exactly $\dim(\mathscr{Y} \cap \text{Ker[T]})$ of its members were a spanning set for $\mathscr{Y} \cap \text{Ker[T]}$, the rest of the vectors in $\mathscr{B}$ would be, as a whole, linearly independent of the kernel, and thus would neatly form a basis for $\text{T}\mathscr{Y}$. But we know from corollary 2.2.3 that, unless $\mathscr{Y}$ is completely inside the kernel, $\mathscr{B}$ can be such that *none* of its members actually lie in the kernel.

Corollary 3.1.4. We can see from the Rank-Nullity theorem that the size of the kernel is bounded from below by $\dim(\mathscr{V}) - \dim(\mathscr{W})$. As a result if $\mathscr{V}$ is no larger than $\mathscr{W}$, the kernel has the option of being trivial. But if $\mathscr{V}$ is larger than $\mathscr{W}$, the kernel can *not* be trivial, and there has to be a non-trivial subspace of $\mathscr{V}$ that is completely smashed down to zero by T.

3.2 The vector space of linear maps and their matrix representation

Let us now fix the two vector spaces $\mathscr{V}$ and $\mathscr{W}$, that as always share their carrier field $\mathbb{F}$, and look at all possible linear transformations from $\mathscr{V}$ to $\mathscr{W}$. Defining scalar multiplication and addition in their natural way, one can easily notice that this new space is in fact by itself a new vector space over the same field $\mathbb{F}$:

Definition 3.2.1. Let $\mathscr{V}$ and $\mathscr{W}$ be two (possibly coinciding) vector spaces over the field $\mathbb{F}$. We define $\mathscr{L}(\mathscr{V},\mathscr{W})$ to be the set of all linear transformations $\mathrm{T}\colon \mathscr{V}\to\mathscr{W}$. Defining scalar multiplication for any linear transformation to be:

$$\alpha\mathrm{T} = \mathrm{S} \qquad \text{s.t.} \qquad \mathrm{S}u = \alpha\mathrm{T}u \quad \forall u \in \mathscr{V}$$

and addition of linear transformations by:

$$\mathrm{T}_1 + \mathrm{T}_2 = \mathrm{S} \qquad \text{s.t.} \qquad \mathrm{S}u = \mathrm{T}_1 u + \mathrm{T}_2 u \quad \forall u \in \mathscr{V}$$

one can easily check that $\mathscr{L}(\mathscr{V},\mathscr{W})$ is in fact a vector space over $\mathbb{F}$.

In fact, $\mathscr{L}(\mathscr{V},\mathscr{W})$ is a vector space merely by inheriting the linearity properties of $\mathscr{V}$ and $\mathscr{W}$. But in the special case where $\mathscr{W} = \mathscr{V}$ a multiplication operation can be defined over $\mathscr{L}(\mathscr{V},\mathscr{V})$. Composition of linear transformations can serve as a non-commutative multiplication over $\mathscr{L}(\mathscr{V},\mathscr{V})$ which readily associates with scalar multiplication and distributes over vector addition (where vectors are now linear transformations). Thus $\mathscr{L}(\mathscr{V},\mathscr{V})$ is a *linear algebra* in abstract algebra terminology, and by itself (forgetting about the field of scalars) is now a non-commutative ring, instead of just an abelian group. We will take advantage of this property later on when discussing evaluation of polynomials over linear transformations.

Earlier on we saw that by fixing a basis for $\mathscr{V}$, we have an isomorphism between $\mathscr{V}$ and $\mathbb{F}^{\dim(\mathscr{V})}$. We now prove a similar result for $\mathscr{L}(\mathscr{V},\mathscr{W})$. Of course, one can immediately see how any choice of bases for $\mathscr{V}$ and $\mathscr{W}$, namely $\mathscr{B}$ and $\mathscr{B}'$ respectively, induces an isomorphism between $\mathscr{L}(\mathscr{V},\mathscr{W})$ and $\mathscr{L}(\mathbb{F}^n,\mathbb{F}^m)$. But what exactly are elements in $\mathscr{L}(\mathbb{F}^n,\mathbb{F}^m)$?

Proposition 3.2.1. *For any linear transformation* L *in* $\mathscr{L}(\mathbb{F}^n,\mathbb{F}^m)$*. Recall the natural basis for* $\mathbb{F}^n$*:* $\mathscr{E} = \{e_1, e_2, \ldots, e_n\}$*. Now if we define the* $m \times n$ *matrix* $[\mathrm{L}]$ *of entries in* $\mathbb{F}$ *to be:*

$$[\mathrm{L}] = \begin{bmatrix} \mathrm{L}e_1 & \mathrm{L}e_2 & \ldots & \mathrm{L}e_n \end{bmatrix}$$

then the application of L *is equal to left matrix multiplication by* $[\mathrm{L}]$*, i.e. for all* $u \in \mathbb{F}^n$ *we have:*

$$\mathrm{L}u = [\mathrm{L}] \cdot u$$

where $\cdot$ *denotes matrix multiplication.*

Proof. We know that if $u \in \mathscr{V}$ is $[\alpha_1, \alpha_2, \ldots, \alpha_n]^\intercal$ we have:

$$u = \sum_{i=1}^{n} \alpha_i e_i \Rightarrow \mathrm{L}u = \sum_{i=1}^{n} \alpha_i \mathrm{L}e_i$$

which is obviously equal to the matrix multiplication introduced. □

Corollary 3.2.1. Replacing "application of a function to an n-tuple" with "left multiplication by a matrix", $\mathscr{L}(\mathbb{F}^n,\mathbb{F}^m)$ and $\mathbb{F}^{\dim(\mathscr{W})\times\dim(\mathscr{V})}$ are essentially the same space.

Proposition 3.2.2. *Let* $\mathscr{V}$ *and* $\mathscr{W}$ *be two vector spaces over the field* $\mathbb{F}$ *with dimensionalities* n *and* m*, respectively.* $\mathscr{L}(\mathscr{V},\mathscr{W})$ *is isomorphic to* $\mathbb{F}^{m\times n}$*, the vector space of all* $m \times n$ *matrices of entries from* $\mathbb{F}$*.*

Proof. Let $\mathscr{B}_v = \{b_1, b_2, \ldots, b_n\}$ be a basis for $\mathscr{V}$ and $\mathscr{B}_w = \{b'_1, b'_2, \ldots, b'_m\}$ be a basis for $\mathscr{W}$. By proposition 2.3.1, $\mathscr{B}_v$ induces an isomorphism $\varphi^{(\mathscr{B}_v)}$ between $\mathscr{V}$ and $\mathbb{F}^n$:

$$u \mapsto [u]_{\mathscr{B}}$$

and similarly $\mathscr{B}_w$ induces an isomorphism $\varphi^{(\mathscr{B}_w)}$ between $\mathscr{W}$ and $\mathbb{F}^m$:

$$v \mapsto [v]_{\mathscr{B}'}$$

For any $\mathrm{T} \in \mathscr{L}(\mathscr{V}, \mathscr{W})$, we know that T is uniquely identified by $\mathrm{T}b_1, \mathrm{T}b_2, \ldots, \mathrm{T}b_n$. Obviously all these have unique counterparts in $\mathbb{F}^m$, as $\varphi^{(\mathscr{B}_w)}$ maps them to $[\mathrm{T}b_i]_{\mathscr{B}_w}$. Define the mapping

$$\varphi^{(\mathscr{B}_v\mathscr{B}_w)} \colon \mathscr{L}(\mathscr{V}, \mathscr{W}) \to \mathscr{L}(\mathbb{F}^n, \mathbb{F}^m)$$

to be one that maps any linear transformation T in $\mathscr{L}(\mathscr{V}, \mathscr{W})$ to a linear transformation L in $\mathscr{L}(\mathbb{F}^n, \mathbb{F}^m)$ such that:

$$[\mathrm{T}u]_{\mathscr{B}_w} = \mathrm{L}[u]_{\mathscr{B}_v} = [\mathrm{L}] \cdot [u]_{\mathscr{B}_v}$$

where the right hand side is a matrix multiplication. In other words:

$$[\varphi^{(\mathscr{B}_v\mathscr{B}_w)}(\mathrm{T})] \cdot [u]_{\mathscr{B}_v} = [\mathrm{T}u]_{\mathscr{B}_w}$$

but we already know how to find the $\mathbb{F}^{m\times n}$ matrix, left multiplication by which, is equivalent to a application of a linear transformation in $\mathscr{L}(\mathbb{F}^n, \mathbb{F}^m)$:

$$[\mathrm{L}] = \begin{bmatrix} \mathrm{L}e_1 & \mathrm{L}e_2 & \ldots & \mathrm{L}e_n \end{bmatrix}$$

but we know that b_i is the vector in $\mathscr{V}$ that is mapped to e_i by $[\cdot]_{\mathscr{B}}$, as a result $\varphi^{(\mathscr{B}\mathscr{B}')}$ operates in the following way:

$$\mathrm{T} \xrightarrow{\varphi^{(\mathscr{B}_v\mathscr{B}_w)}} \begin{bmatrix} [\mathrm{T}b_1]_{\mathscr{B}_w} & [\mathrm{T}b_2]_{\mathscr{B}_w} & \ldots & [\mathrm{T}b_n]_{\mathscr{B}_w} \end{bmatrix}$$

Since $\mathscr{B}_v$ is a basis and thus spans all $\mathscr{V}$, this mapping is injective, i.e. to any possible T, $\varphi^{(\mathscr{B}_v\mathscr{B}')_w}$ assigns an $m \times n$ matrix of scalars. One can easily observe that it is surjective as well, since the columns of any matrix A of scalars induces n vectors in $\mathscr{W}$ to be $\mathrm{T}b_i$s and hence uniquely identifies a linear transformation in $\mathscr{L}(\mathscr{V}, \mathscr{W})$. Also by uniqueness of representation, $\varphi^{(\mathscr{B}_v\mathscr{B}_w)}$ is one-to-one. As a result if we can prove that $\varphi^{(\mathscr{B}\mathscr{B}')}$ is a homomorphism, it, in fact, is an isomorphism. To prove the latter we need to show that $\varphi^{(\mathscr{B}_v\mathscr{B}_w)}$ respects the algebraic structure of vector spaces. Specifically, since all the vector spaces $\mathscr{V}$, $\mathscr{W}$, $\mathbb{F}^n$, and $\mathbb{F}^m$ share their underlying field, we just have to show that the following hold for all $\mathrm{T}_1, \mathrm{T}_2 \in \mathscr{L}(\mathscr{V}, \mathscr{W})$ and all $\alpha \in \mathbb{F}$:

$$\varphi^{(\mathscr{B}_v\mathscr{B}_w)}(\mathrm{T}_1 + \mathrm{T}_2) = \varphi^{(\mathscr{B}_v\mathscr{B}_w)}(\mathrm{T}_1) + \varphi^{(\mathscr{B}_v\mathscr{B}_w)}(\mathrm{T}_2)$$

$$\alpha\varphi^{(\mathscr{B}_v\mathscr{B}_w)}(\mathrm{T}_1) = \varphi^{(\mathscr{B}_v\mathscr{B}_w)}(\alpha\mathrm{T}_1)$$

both of which quite simply follow from uniqueness of representation in terms of bases. □

Corollary 3.2.2. Let $\mathscr{V}$ and $\mathscr{W}$ be two (possibly coinciding) vector spaces over the field $\mathbb{F}$. We have:

$$\dim(\mathscr{L}(\mathscr{V}, \mathscr{W})) = \dim(\mathscr{V}) \times \dim(\mathscr{W}).$$

Definition 3.2.2. As of now, based on proposition 3.2.2 we adopt the following notation to refer the matrix counterpart of a linear transformation. For any linear transformation $\mathrm{T}\colon \mathscr{V} \to \mathscr{W}$ and any basis $\mathscr{B} = \{b_1, b_2, \ldots, b_n\}$ for $\mathscr{V}$ and any basis $\mathscr{B}'$ for $\mathscr{W}$, we use $[\mathrm{T}]_{\mathscr{B}\mathscr{B}'}$ to denote the following $\dim(\mathscr{W}) \times \dim(\mathscr{V})$ matrix of entries in $\mathbb{F}$:

$$\begin{bmatrix}[\mathrm{T}b_1]_{\mathscr{B}'} & [\mathrm{T}b_2]_{\mathscr{B}'} & \ldots & [\mathrm{T}b_n]_{\mathscr{B}'}\end{bmatrix}$$

Obviously we have:

$$[\mathrm{T}u]_{\mathscr{B}'} = [\mathrm{T}]_{\mathscr{B}\mathscr{B}'}[u]_{\mathscr{B}}$$

since if the representation of u in terms of $\mathscr{B}$ is $[\alpha_1, \alpha_2, \ldots, \alpha_n]^\intercal$ we would have:

$$\mathrm{T}u = \sum_{i=1}^{n} \alpha_i \mathrm{T}b_i \Rightarrow [\mathrm{T}u]_{\mathscr{B}'} = \begin{bmatrix}[\mathrm{T}b_1]_{\mathscr{B}'} & [\mathrm{T}b_2]_{\mathscr{B}'} & \ldots & [\mathrm{T}b_n]_{\mathscr{B}'}\end{bmatrix} \begin{bmatrix}\alpha_1 \\ \alpha_2 \\ \vdots \\ \alpha_n\end{bmatrix}$$

where we have dropped the $\cdot$ notation for matrix multiplication.

Proposition 3.2.3 (Linear map composition equals Matrix multiplication)**.** *Let $\mathrm{T}, \mathrm{S}\colon \mathscr{V} \to \mathscr{W}$ be any two arbitrary linear transformations, and let $\mathscr{B}$ and $\mathscr{B}'$ be bases, respectively for $\mathscr{V}$ and $\mathscr{W}$. The following holds:*

$$[\mathrm{T} \circ \mathrm{S}]_{\mathscr{B}\mathscr{B}'} = [\mathrm{T}]_{\mathscr{B}\mathscr{B}'}[\mathrm{S}]_{\mathscr{B}\mathscr{B}'}$$

Proof. The claim can be proved by a few applications of the defintion:

$$\begin{aligned}
[\mathrm{T} \circ \mathrm{S}]_{\mathscr{B}\mathscr{B}'} &= \begin{bmatrix}[\mathrm{TS}b_1]_{\mathscr{B}'} & [\mathrm{TS}b_2]_{\mathscr{B}'} & \ldots & [\mathrm{TS}b_n]_{\mathscr{B}'}\end{bmatrix} \\
&= \begin{bmatrix}[\mathrm{T}]_{\mathscr{B}\mathscr{B}'}[\mathrm{S}b_1]_{\mathscr{B}'} & [\mathrm{T}]_{\mathscr{B}\mathscr{B}'}[\mathrm{S}b_2]_{\mathscr{B}'} & \ldots & [\mathrm{T}]_{\mathscr{B}\mathscr{B}'}[\mathrm{S}b_n]_{\mathscr{B}'}\end{bmatrix} \\
&= [\mathrm{T}]_{\mathscr{B}\mathscr{B}'}\begin{bmatrix}[\mathrm{S}]_{\mathscr{B}\mathscr{B}'}[b_1]_{\mathscr{B}} & [\mathrm{S}]_{\mathscr{B}\mathscr{B}'}[b_2]_{\mathscr{B}} & \ldots & [\mathrm{S}]_{\mathscr{B}\mathscr{B}'}[b_n]_{\mathscr{B}}\end{bmatrix} \\
&= [\mathrm{T}]_{\mathscr{B}\mathscr{B}'}[\mathrm{S}]_{\mathscr{B}\mathscr{B}'}\underbrace{\begin{bmatrix}[b_1]_{\mathscr{B}} & [b_2]_{\mathscr{B}} & \ldots & [b_n]_{\mathscr{B}}\end{bmatrix}}_{\mathrm{I}_{n\times n}} \\
&= [\mathrm{T}]_{\mathscr{B}\mathscr{B}'}[\mathrm{S}]_{\mathscr{B}\mathscr{B}'}
\end{aligned}$$

and the proof is complete. □

The last proposition of the section is the change of basis formula. We derive the change of basis for the general case of $\mathrm{T}\colon \mathscr{V} \to \mathscr{W}$.

Proposition 3.2.4 (Change of basis)**.** *Let $\mathrm{T}\colon \mathscr{V} \to \mathscr{W}$ be a linear transformation, $\mathscr{B} = \{b_1, \ldots, b_n\}$ and $\mathscr{C}$ be two bases for $\mathscr{V}$, and $\mathscr{B}'$ and $\mathscr{C}'$ be two bases for $\mathscr{W}$. The representations of T in terms of $\mathscr{B}\mathscr{B}'$ and $\mathscr{C}\mathscr{C}'$ are related by the following formula:*

$$[\mathrm{T}]_{\mathscr{B}\mathscr{B}'} = \mathrm{P}[\mathrm{T}]_{\mathscr{C}\mathscr{C}'}\mathrm{Q}$$

where P and Q both have the same form as in in proposition 2.3.2 of the general form of basis change matrices we introduced earlier, $\mathrm{P} = \mathrm{P}_{(\mathscr{C}' \to \mathscr{B}')}$ and $\mathrm{Q} = \mathrm{P}_{(\mathscr{B} \to \mathscr{C})}$. We could more succinctly write:

$$[\mathrm{T}]_{\mathscr{B}\mathscr{B}'} = [\mathscr{C}']_{\mathscr{B}'}[\mathrm{T}]_{\mathscr{C}\mathscr{C}'}[\mathscr{B}]_{\mathscr{C}}$$

Proof. We first notice that:

$$[\mathrm{T}]_{\mathscr{B}\mathscr{B}'} = \begin{bmatrix}[\mathrm{T}b_1]_{\mathscr{B}'} & [\mathrm{T}b_2]_{\mathscr{B}'} & \dots & [\mathrm{T}b_n]_{\mathscr{B}'}\end{bmatrix}$$

Now by proposition 2.3.2 we know that for all $[\mathrm{T}b_i]_{\mathscr{B}'}$ columns we can write:

$$[\mathrm{T}b_i]_{\mathscr{B}'} = \mathrm{P}[\mathrm{T}b_i]_{\mathscr{C}'}$$

and hence:

$$[\mathrm{T}]_{\mathscr{B}\mathscr{B}'} = \mathrm{P}\begin{bmatrix}[\mathrm{T}b_1]_{\mathscr{C}'} & [\mathrm{T}b_2]_{\mathscr{C}'} & \dots & [\mathrm{T}b_n]_{\mathscr{C}'}\end{bmatrix}$$

we know that:

$$[\mathrm{T}b_i]_{\mathscr{C}'} = [\mathrm{T}]_{\mathscr{C}\mathscr{C}'}[b_i]_{\mathscr{C}}$$

therefore we get:

$$[\mathrm{T}]_{\mathscr{B}\mathscr{B}'} = \mathrm{P}[\mathrm{T}]_{\mathscr{C}\mathscr{C}'}\begin{bmatrix}[b_1]_{\mathscr{C}} & [b_2]_{\mathscr{C}} & \dots & [b_n]_{\mathscr{C}}\end{bmatrix}$$

using again proposition 2.3.2 we notice that:

$$\begin{bmatrix}[b_1]_{\mathscr{C}} & [b_2]_{\mathscr{C}} & \dots & [b_n]_{\mathscr{C}}\end{bmatrix} = \mathrm{P}_{(\mathscr{B}\to\mathscr{C})}$$

and thus we finally get:

$$[\mathrm{T}]_{\mathscr{B}\mathscr{B}'} = \mathrm{P}[\mathrm{T}]_{\mathscr{C}\mathscr{C}'}\mathrm{Q}$$

which completes the proof. □

Corollary 3.2.3. In the special case where $\mathscr{W} = \mathscr{V}$, and we want to get $[\mathrm{T}]_{\mathscr{B}\mathscr{B}}$ from $[\mathrm{T}]_{\mathscr{B}'\mathscr{B}'}$ we get the following:

$$[\mathrm{T}]_{\mathscr{B}\mathscr{B}} = \mathrm{P}[\mathrm{T}]_{\mathscr{B}'\mathscr{B}'}\mathrm{P}^{-1}$$

where P is the matrix $\mathrm{P} = \mathrm{P}_{(\mathscr{B}'\to\mathscr{B})}$. We could also write:

$$[\mathrm{T}]_{\mathscr{B}\mathscr{B}} = [\mathscr{B}']_{\mathscr{B}}[\mathrm{T}]_{\mathscr{B}'\mathscr{B}'}[\mathscr{B}]_{\mathscr{B}'}$$

As seen in the corollary 3.2.3, a single linear map $\mathrm{T}\colon \mathscr{V}\to\mathscr{V}$ has countless matrix representations A depending on the choice of basis, all of which are related by a transformation of the form:

$$\mathrm{A}' = \mathrm{PAP}^{-1}$$

But is there any constraint on the kind of matrices that could be $\mathrm{P}_{(\mathscr{B}'\to\mathscr{B})}$ for some $\mathscr{B}$ and $\mathscr{B}'$. On part of $\mathscr{B}$ and $\mathscr{B}'$, all we require is that the two are linearly independent sets of $\dim(\mathscr{V})$ vectors. From a corollary to proposition 2.3.1 we remember that given $\mathscr{B}$ is a basis, $\mathscr{B}'$ is a basis if and only if the matrix $\mathrm{P}_{(\mathscr{B}'\to\mathscr{B})}$ has linearly independent columns (each column treated as a vector in $\mathbb{F}^n$), or in matrix terminology has full column rank. As we shall see later, this is again equivalent to $\mathrm{P}_{(\mathscr{B}'\to\mathscr{B})}$ being a non-singular matrix. Also we notice that this matrix, has no trace of the actual identity of $\mathscr{B}$, but only the representation of $\mathscr{B}'$ in terms of $\mathscr{B}$. Hence for *any* choice of basis $\mathscr{B}$ for $\mathscr{V}$, the vectors identified by those linear combinations of them identified by columns of $\mathrm{P}_{(\mathscr{B}'\to\mathscr{B})}$, again is a basis for $\mathscr{V}$. It is completely in tune with what we mentioned earlier about the different isomorphisms that different choices of basis induce between an n-dimensional vector space $\mathscr{V}$ and $\mathbb{F}^n$. And we can safely conclude that the operation described above, transforming A to PAP^{-1} for some P, induces an equivalence relation:

Definition 3.2.3 (Similar Matrices)**.** Let A and B be two $n \times n$ matrices of scalars in $\mathbb{F}$. We say that A is **similar** to B and write $A \sim B$ if there exists a non-singular (and hence[1] $n \times n$) matrix P such that:

$$A = PBP^{-1}$$

Proposition 3.2.5. *The matrix similarity relationship $\sim$ defined above is an equivalence relationship: reflexive, symmetric, and transitive. Furthermore, for any vector space $\mathscr{V}$ of dimensionality n over $\mathbb{F}$, any basis $\mathscr{B}$ for $\mathscr{V}$ induces a one-to-one mapping between the equivalence classes of $n \times n$ matrices and linear transformations $T : \mathscr{V} \to \mathscr{V}$. To be more exact, by fixing a basis $\mathscr{B}$ for $\mathscr{V}$, all the members of an equivalence class are different representations of the very same linear transformation over $\mathscr{V}$. Of course, by changing the basis $\mathscr{B}$, the mapping between equivalence classes and linear transformations over $\mathscr{V}$ changes, but the equivalence classes themselves are untouched, and still all correspond to one single (now different) linear transformation on $\mathscr{V}$.*

Proposition 3.2.6. *The only member of the equivalence class of $I_{n\times n}$ is itself. The only member of the equivalence class of the zero matrix is itself.*

Remark. The similarity relation is defined over scalar matrices, and not linear transformations. Linear transformations do not get to be "similar". Because the essence of similarity is that "two things are similar, if they are essentially the same unique object, and an appropriate change of *basis* will show this unity", and linear transformations do not get their identity from bases, so that a change of basis would "disguise" their true identity. However, when talking about $\mathscr{V} = \mathbb{F}^n$ since linear transformations *are* matrices (not just *represented* by matrices), one could talk about two linear transformations being similar, meaning that they have been the same matrix, but a change of basis made them look different; and this is the only sense in which one could talk about similarity of linear transformations.

3.3 Inverses and In(Sur)jective transformations

Linear transformations before anything are functions, and thus it would be reasonable to think of their inverses. It would be worthwhile here to mention the two quite different functionalities we refer to when we talk about inverses. The first one is what we are going to call a *left* inverse:

Definition 3.3.1. Let $T : \mathscr{V} \to \mathscr{W}$ be any arbitrary mapping between two vector spaces. The map $S : \mathscr{W} \to \mathscr{V}$ is the **left inverse** of T if for any $u \in \mathscr{V}$, S maps Tu back to u.

Obviously such an inverse will be well-defined only if T is injective. Furthermore, unless T is surjective, S is not defined all over $\mathscr{W}$, and could have any arbitrary behavior on vectors lying in $\mathscr{W} - \mathrm{Im}[T]$. Figure 3.1 depicts the behavior of a left inverse. The behavior of a left inverse actually is defined with respect to the elements in the *departure* space of T, in the fashion that S tries to cancel out the effect of T when *left*-composed with T:

$$ST u = u$$

But there is another functionality we at times refer to when we talk about inverses. A function

[1] The reason for this will be clear in chapter 3.3

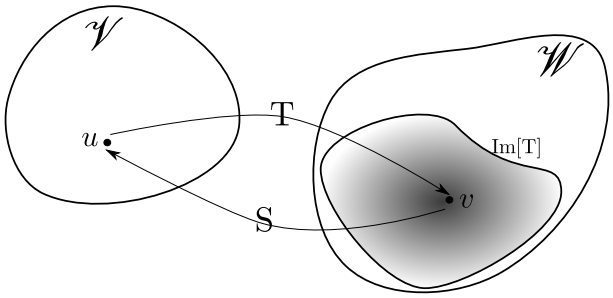

Figure 3.1: Left Inverse

that looks at a vector v in the destination space of T, and looks back into the departure space, and tries to find a vector u that would have been the reason v is in the image of T. We call this second function the right inverse:

Definition 3.3.2. Let $\mathrm{T}\colon \mathscr{V} \to \mathscr{W}$ be any arbitrary mapping between two vector spaces. The map $\mathrm{S}\colon \mathscr{W} \to \mathscr{V}$ is the **right inverse** of T if for any $v \in \mathscr{W}$, S maps v back to a u that satisfies $\mathrm{T}u = v$.

Obviously such an inverse will be well-defined only if T is surjective. Furthermore, unless T is injective, S is not well-defined over $\mathscr{W}$, since there could exist multiple vectors u satisfying $\mathrm{T} = v$. Figure 3.2 depicts the behavior of a right inverse. The behavior of a right inverse actually is defined with respect to the elements in the *destination* space of T, in the fashion that S tries to cancel out the effect of T when *right*-composed with T:

$$\mathrm{TS}u = u$$

Of course, if both inverses exist, they are bound to coincide; however, one could exist while the

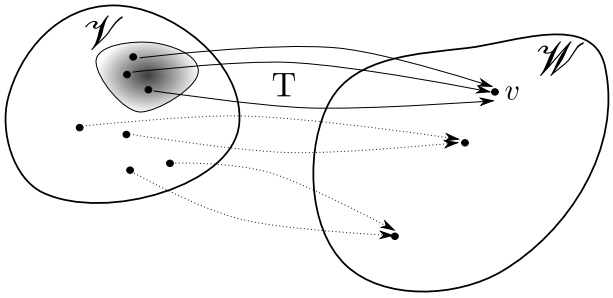

Figure 3.2: Right Inverse

other does not! Now that the connection of these two types of inverses are established with injectivity and surjectivity properties of T, we observe some results in that regard, for the case where T is a linear mapping.

It can easily be seen that injectivity is closely related to the nullity of a linear transformation, and surjectivity to its rank.

Proposition 3.3.1. *Let* $\mathrm{T}\colon \mathscr{V} \to \mathscr{W}$ *be a linear transformation.* T *is injective if and only if it has zero nullity, i.e it has a trivial kernel.*

Proof. First assume $\eta(\mathrm{T}) = 0$ and there exists distinct $u, v \in \mathscr{V}$ such that $\mathrm{T}u = \mathrm{T}v$. It would immediately follow that $\mathrm{T}(u - v) = 0$ which is a contradiction since $u - v \neq 0$. Now assume T is injective, and let $\mathrm{T}u = 0$ for some nonzero vector $u \in \mathscr{V}$. This is again a contradiction since we already have two vectors, namely u and 0, both mapped to the zero vector in $\mathscr{W}$. □

Corollary 3.3.1. By using the Rank-Nullity theorem it follows that the following statements are equivalent:

(i) $\mathrm{T}: \mathscr{V} \to \mathscr{W}$ is injective.

(ii) $\rho(\mathrm{T}) = \dim(\mathscr{V})$.

(iii) $\eta(\mathrm{T}) = 0$.

We saw in corollary 3.1.4 that if $\mathscr{V}$ is larger than $\mathscr{W}$, the kernel does not have the option of being trivial, and hence no linear transformation T can be injective in this case. We now turn our attention to the problem of surjectivity.

Proposition 3.3.2. *Let* $\mathrm{T} : \mathscr{V} \to \mathscr{W}$ *be a linear transformation.* T *is surjective if and only if* $\rho(\mathrm{T}) = \dim(\mathscr{W})$.

Proof. If we have $\rho(\mathrm{T}) = \dim(\mathscr{W})$, since Im[T] already is a subspace of $\mathscr{W}$, it follows from it having equal dimensionality to $\mathscr{W}$, that Im[T] swallows the whole destination space, and hence T is surjective. The other direction is trivial. □

Corollary 3.3.2. By using the Rank-Nullity theorem it follows that the following statements are equivalent:

(i) $\mathrm{T}: \mathscr{V} \to \mathscr{W}$ is surjective.

(ii) $\rho(\mathrm{T}) = \dim(\mathscr{W})$.

(iii) $\eta(\mathrm{T}) = \dim(\mathscr{V}) - \dim(\mathscr{W})$.

Again referring to corollary 3.1.4, we notice that if $\mathscr{V}$ is smaller than $\mathscr{W}$, the right hand side of the last equation in the last corollary would be negative, and no linear transformation T would not have the option of being surjective, to sum up:

Corollary 3.3.3. Let $\mathscr{V}$ and $\mathscr{W}$ be two vector spaces. If $\mathscr{V}$ is larger than $\mathscr{W}$ no linear transformation $\mathrm{T}: \mathscr{V} \to \mathscr{W}$ could possibly be injective, and if $\mathscr{W}$ is larger than $\mathscr{V}$ no linear transformation $\mathrm{T}: \mathscr{V} \to \mathscr{W}$ could possibly be surjective. The rest depends on the behavior of different linear transformations.

Notice that we would not know for granted that the right and left inverses *function*s (in case of existence) would also satisfy nice linearity properties, and hence be linear transformations.

Proposition 3.3.3. *Let* $\mathrm{T}: \mathscr{V} \to \mathscr{W}$ *be an injective linear transformation, then* T *has a unique left inverse* $\mathrm{S}: \mathrm{Im}[\mathrm{T}] \to \mathscr{V}$*, and* S *is a bijective linear transformation, and therefore* Im[T] *and* $\mathscr{V}$ *would be isomorphic.*

Proof. Obviously S is well defined for all members of Im[T]. By assumption, for any $v \in \text{Im}[\text{T}]$ we have a unique $u \in \mathscr{V}$ such that $\text{T}u = v$ and hence by definition of S we have $\text{S}v = \text{ST}u = u$, and hence the unique behavior of S all over its domain (which is Im[T]) is inherited by injectivity of T. To see why S is a linear transformation notice that for all $v_1, v_2 \in \text{Im}[\text{T}]$ where $v_i = \text{T}u_i$ we have:

$$\text{S}(\alpha v_1) = \text{ST}(\alpha u_1) = \alpha u_1 = \alpha \text{S}v_1$$

and:

$$\text{S}(v_1 + v_2) = \text{S}(\text{T}u_1 + \text{T}u_2) = \text{ST}(u_1 + u_2) = u_1 + u_2 = \text{S}v_1 + \text{S}v_2$$

The fact that S is surjective and injective is trivial. □

Unfortunately the world of right inverses is not at all this neat. Since a right inverse might exist (as we mentioned iff T is surjective), but would not be unique unless T is injective as well. Let's assume for now that T is not injective, and investigate the problem of tracing back a vector $v \in \mathscr{W}$ to the set of vectors in $\mathscr{V}$ satisfying $\text{T}u = v$.

Definition 3.3.3. Let $\text{T}: \mathscr{V} \to \mathscr{W}$ be an arbitrary linear transformation. For any $v \in \mathscr{W}$, we call the equation $\text{T}u = v$ a **linear system**, the solution to which is the **set** of vectors $u \in \mathscr{V}$ that satisfy $\text{T}u = v$, and use the following notation:

$$\text{T}_v^{-1} = \{u \in \mathscr{V} : \text{T}u = v\}$$

We insist on not making the impression that T^{-1} is actually a mapping, since unless T is *bijective* there does *not* exist any well-defined linear mapping with the properties of a right inverse.

Remark. The set T_v^{-1} could be regarded as a generalization of the notion of *level sets*. This is the set of all members in the departure space whose result after applying T is the same vector v in the destination space.

Obviously for $v \notin \text{Im}[\text{T}]$ the set T_v^{-1} would be empty. Notice that by fixing any two bases for $\mathscr{V}$ and $\mathscr{W}$ the above linear system will translate into the standard matrix formulation of linear systems of equations. But notice that since the matrix formulation is completely dependent on the arbitrary choice of bases, it would be more reasonable to consider the problem abstractly.

Proposition 3.3.4. *Let $\mathscr{V}$ and $\mathscr{W}$ be vector spaces over the field $\mathbb{F}$, and let $\mathscr{V}$ have dimensionality n. Let $\text{T}: \mathscr{V} \to \mathscr{W}$ be an arbitrary linear transformation and $\mathscr{B}$ be a basis for $\mathscr{V}$. For any $v \in \mathscr{W}$, the set $[\text{T}_v^{-1}]_{\mathscr{B}} \subseteq \mathbb{F}^n$ is a compact convex set.*

Proof. Let $u_1, u_2 \in \text{T}_v^{-1}$. For any $\lambda \in (0, 1)$ we would have:

$$\text{T}(\lambda u_1 + (1-\lambda)u_2) = \lambda \text{T}u_1 + (1-\lambda)\text{T}u_2 = \lambda v + (1-\lambda)v = v$$

and it follows that $\lambda u_1 + (1-\lambda)u_2 \in \text{T}_v^{-1}$. Since T_v^{-1} and $[\text{T}_v^{-1}]_{\mathscr{B}}$ are isomorphic, for any two vectors $[u_1]_{\mathscr{B}}, [u_2]_{\mathscr{B}} \in [\text{T}_v^{-1}]_{\mathscr{B}}$, $\lambda[u_1]_{\mathscr{B}} + (1-\lambda)[u_2]_{\mathscr{B}}$ also falls inside $[\text{T}_v^{-1}]_{\mathscr{B}}$ which completes the proof for convexity. To see why $[\text{T}_v^{-1}]_{\mathscr{B}}$ is closed as well, we assume the Euclidean metric over $\mathbb{F}^n$, fix a basis $\mathscr{B}'$ over $\mathscr{W}$ and notice that for any sequence of elements $\{u_i\}_{i=1}^{\infty}$ in $[\text{T}_v^{-1}]_{\mathscr{B}}$ converging to some $[u_\circ]_{\mathscr{B}}$ we have:

$$[\text{T}u_\circ]_{\mathscr{B}'} = [\text{T}]_{\mathscr{B}\mathscr{B}'}[u_\circ]_{\mathscr{B}} = [\text{T}]_{\mathscr{B}\mathscr{B}'} \lim_{i\to\infty} [u_i]_{\mathscr{B}} \overset{?}{=} \lim_{i\to\infty} [\text{T}]_{\mathscr{B}\mathscr{B}'}[u_i]_{\mathscr{B}} = \lim_{i\to\infty} [\text{T}u_i]_{\mathscr{B}'}$$

where in the last derivation on the right hand side, all the elements are in fact, by definition of u_i, equal to $[v]_{\mathscr{B}'}$, and the proof would be complete if the derivation with the question mark on top would be correct. This is *not* trivial, but a simple investigation of the nature of matrix multiplication of scalars, shows its truth. □

An interesting inquiry, now, would be the dimensionality of T_v^{-1}, or to be more exact $\dim^{\mathrm{span}}(\mathrm{T}_v^{-1})$. As we saw for $v \notin \mathrm{Im}[\mathrm{T}]$ the set T_v^{-1} is empty, while for other vectors it is not, but one might ask how much variation can there be in the size of T_v^{-1} by different choices of v, and how much intrinsic properties of T limit T_v^{-1}. More importantly other than $\mathrm{T}_0^{-1} = \mathrm{Ker}[\mathrm{T}]$ what other sets T_v^{-1} could possibly swallow up a whole subspace of $\mathscr{V}$?
The answer to the second question is quite easy. No v but 0 could make T_v^{-1} be a subspace in $\mathscr{V}$. To see this, let $\mathrm{T}u = v$ and for $v \neq 0$. If T_v^{-1} is a subspace, it should as well contain αu for any scalar α, which leads to $\mathrm{T}\alpha u = \alpha v = v$, which can hold for non-zero scalars only if $v = 0$. Now let's investigate the first question. With the following observation and in light of the derivations of proposition 2.2.7 we can see completely what the structure of T_v^{-1} is:

Proposition 3.3.5. *For any two arbitrary vectors u_1, u_2 in T_v^{-1}, we will have $u_1 - u_2 \in \mathrm{Ker}[\mathrm{T}]$.*

This coincides exactly with the definition of *shifted subspaces* (which are not subspaces themselves) earlier in section 2.2. Obviously for any $v \in \mathrm{Im}[\mathrm{T}]$, there exists some $u_\circ$ such that $\mathrm{T}u_\circ = v$ and all the other members of T_v^{-1} are uniquely identified by the shifted subspace $u_\circ + \mathrm{Ker}[\mathrm{T}]$ (which is different from $\mathrm{span}\{u_\circ\} + \mathrm{Ker}[\mathrm{T}]$).

Proposition 3.3.6. *Let $\mathrm{T} : \mathscr{V} \to \mathscr{W}$ be an arbitrary linear transformation. For any $v \in \mathrm{Im}[\mathrm{T}]$, the level set T_v^{-1} is a shifted subspace of the kernel. If $v = 0$ this shifted subspace will itself be a subspace coinciding with* Ker[T]*, and otherwise it will not be a subspace and we would have:*

$$\dim^{\mathrm{span}}(\mathrm{T}_v^{-1}) = \eta(\mathrm{T}) + 1$$

Proof. The equation about the dimensionality follows from proposition 2.2.7. □

Corollary 3.3.4. We have seen that for the case where $v = 0$ the dimensionality of the level set T_v^{-1} is exactly $\eta(\mathrm{T})$ (and the set in fact coincides with the kernel) and in other cases it will be exactly $\eta(\mathrm{T}) + 1$. We could write the following inequality in general, for all v:

$$\forall v \in \mathrm{Im}[\mathrm{T}] : \ 0 \leq \dim^{\mathrm{span}}(\mathrm{T}_v^{-1}) - \eta(\mathrm{T}) \leq 1$$

Now that we have understood the structure of the set T_v^{-1} for the case where T is not injective, we wrap up the study of inverses. We saw earlier that if T is injective it has a unique left inverse $\mathrm{S} : \mathrm{Im}[\mathrm{T}] \to \mathscr{V}$. The result of our study of the shifted subspaces T_v^{-1} would be the following:

Proposition 3.3.7. *Let $\mathrm{T} : \mathscr{V} \to \mathscr{W}$ be an injective linear transformation. Then* T *has a unique right inverse $\mathrm{S} : \mathrm{Im}[\mathrm{T}] \to \mathscr{V}$, and* S *is a bijective linear transformation (an isomorphism) between* Im[T] *and $\mathscr{V}$.*

Proof. We know by proposition 3.3.1 that since T is injective it has a trivial kernel. As a result for any $v = \mathrm{T}u \in \mathrm{Im}[\mathrm{T}]$ by proposition 3.3.6 we know that the set T_v^{-1} is the shifted subspace $u + \mathrm{Ker}[\mathrm{T}]$, which since kernel is trivial, contains the unique vector u. The fact that S is bijective, and hence an isomorphism is trivial. □

Proposition 3.3.8. *The following statements are equivalent:*

(i) $\mathrm{T}\colon \mathscr{V} \to \mathscr{W}$ *is bijective.*

(ii) $\dim(\mathscr{V}) = \dim(\mathscr{W})$ *and* $\eta(\mathrm{T}) = 0$.

(iii) $\rho(\mathrm{T}) = \dim(\mathscr{W})$ *and* $\eta(\mathrm{T}) = 0$.

Furthermore if $\dim(\mathscr{V}) = \dim(\mathscr{W})$, *any linear transformation* $\mathrm{T}\colon \mathscr{V} \to \mathscr{W}$ *is injective if and only if it is surjective.*

Corollary 3.3.5. Combining the last proposition with 3.3.1 we get the following matrix property: Let A be a square matrix of scalars. A is non-singular if and only if there exists no non-trivial solution to the system $\mathrm{A}x = 0$.

Chapter 4

Duality and Transposition

Definition 4.0.4. Let $\mathscr{V}$ be a vector space over the field $\mathbb{F}$, we call $f\colon \mathscr{V}\to\mathbb{F}$ a **linear functional** over $\mathscr{V}$ if it belongs to the space $\mathscr{L}(\mathscr{V},\mathbb{F})$ where $\mathbb{F}$ is regarded as the trivial 1-dimensional vector space over itself. Obviously f assigns a *scalar value* to all vectors in $\mathscr{V}$, and satisfies the following linearity properties:

$$f(u+v) = f(u) + f(v)$$
$$f(\alpha u) = \alpha f(u)$$

Definition 4.0.5. We have already defined the notion of transformation spaces $\mathscr{L}(\mathscr{V},\mathscr{W})$ for arbitrary vector spaces sharing their carrier field, and have seen they are vector spaces if equipped with natural addition and scalar multiplication operators inherited from $\mathscr{V}$ and $\mathscr{W}$. For the special case of $\mathscr{W}$ being the trivial 1-dimensional vector space $\mathbb{F}$ constitutes over itself [1], we call the vector space $\mathscr{L}(\mathscr{V},\mathbb{F})$ the **dual space** of $\mathscr{V}$ and refer to it by $\mathscr{V}^*$.

Proposition 4.0.9. *It follows from proposition 3.2.2 that $\mathscr{V}^*$ shares the dimensionality of $\mathscr{V}$:*

$$\dim(\mathscr{V}^*) = \dim(\mathscr{V})$$

As an example, recall that we introduced in section 1.4 the polynomial ring $\mathbb{F}[x]$ of a field[2]. Now taking the subset of $\mathbb{F}[x]$ containing all degree n polynomials over $\mathbb{F}$ together with natural scalar multiplication, would be an n-dimensional vector space over $\mathbb{F}$: $(\mathbb{F}_n[x],\cdot,\mathbb{F},\cdot,+)$. An example of a linear functional $f_t\colon \mathscr{V}\to\mathbb{F}$ would be one that assigns to any $p(x)\in\mathbb{F}_n[x]$ its evaluation at a fixed point $t\in\mathbb{F}$, all over the vector space:

$$p(x) = \alpha_0 + \alpha_1 x + \alpha_2 x^2 + \ldots + \alpha_n x^n \xrightarrow{\ f\ } \alpha_0 + \alpha_1 t + \alpha_2 t^2 + \ldots + \alpha_n t^n$$

This can easily be checked to be a linear functional. Again from proposition 3.2.2 we know that for any basis $\mathscr{B}$ for $\mathscr{V}$, assuming the trivial basis for the trivial 1-dimensional vector space $\mathbb{F}$, the matrix representation of any linear functional would be a row vector (a $1\times n$ matrix):

$$[f]_{\mathscr{B}} = \begin{bmatrix} f(b_1) & f(b_2) & \ldots & f(b_n) \end{bmatrix}$$

[1] To be more exact the vector space defined as $(\mathbb{F},\cdot,\mathbb{F},+,\cdot)$ where $\cdot$ and $+$ are the operators of $\mathbb{F}$.

[2] As we mentioned, there the polynomial ring is completely well defined over rings, but the special case of much interest is polynomials over fields.

and for any vector $u \in \mathscr{V}$ with its representation in $\mathscr{B}$ being $[\alpha_1, \alpha_2, \ldots, \alpha_n]^\intercal$ we would have:

$$f(u) = [f]_{\mathscr{B}}[u]_{\mathscr{B}} = \begin{bmatrix} f(b_1) & f(b_2) & \ldots & f(b_n) \end{bmatrix} \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \vdots \\ \alpha_n \end{bmatrix}$$

Proposition 4.0.10. *Let $\mathscr{V}$ be an n-dimensional vector space over the field $\mathbb{F}$. It follows immediately from the Rank-Nullity theorem that any linear function f over $\mathscr{V}$ is either the zero functional with $\eta(f) = n$:*

$$f(\cdot) = 0$$

or would have $\eta(f) = n - 1$, in which case it would be surjective:

$$\forall t \in \mathbb{F},\ \exists u \in \mathscr{V} :\ f(u) = t$$

The case where f is not the zero functional is worth more investigation. Notice that $\mathrm{Ker}[f]$ has dimensionality $n - 1$, which is a hyperplane in $\mathscr{V}$. Of course the set of all the vectors lying outside $\mathrm{Ker}[f]$ has $\dim^{\mathrm{span}}$ equal to n (refer to corollary 2.2.3), but one could fix a single vector u outside $\mathrm{Ker}[f]$ and it would follow that $\mathrm{span}\{u\} \oplus \mathrm{Ker}[f] = \mathscr{V}$, and hence building a basis $\mathscr{B}$ whose first member is $u_\circ = \frac{1}{f(u)}u$ and the rest lie inside the kernel, it would follow that for any vector v with representation $[v]_{\mathscr{B}} = [\alpha, *, *, \ldots]^\intercal$ we would have:

$$f(v) = \alpha$$

Thus we have proved the following:

Proposition 4.0.11. *Let $\mathscr{V}$ be a vector space. Any choice of (i) a hyperplane $\mathscr{Y}$ in $\mathscr{V}$ and (ii) a vector $u_\circ$ lying outside this hyperplane, would uniquely identify a linear functional, with the understanding that $f(u_\circ) = 1$ and $\mathscr{Y} = \mathrm{Ker}[f]$.*

Of course, as we saw earlier, the choice of $u_\circ$ is not a unique choice. We could pick *any* vector u outside of the kernel and pick $u_\circ = \frac{1}{f(u)}u$, and everything work the same, but with all these different choices $u_\circ$ would not fall just *anywhere*. We know from proposition 3.3.6 that level sets of f (the sets corresponding to the vectors satisfying $f(\cdot) = c$ for some $c \in \mathbb{F}$) are *shifted subspaces* resulting from shifting the hyperplane $\mathrm{Ker}[f]$. Any level set $f_c^{\text{-}1}$ is a shifted subspace of the kernel. Since it is obviously not the kernel itself as long as c is not zero, it follows from proposition 3.3.6 that the level set $f_c^{\text{-}1}$ has $\dim^{\mathrm{span}}$ equal to $\eta(f) + 1$, which from the fact that f has nullity equal to $n - 1$ it follows:

$$\dim^{\mathrm{span}}(f_c^{\text{-}1}) = n$$

As a result the unique identification process we described in proposition 4.0.11, assigns to a lot of different choices of $\mathscr{Y}$ and $u_\circ$ the same linear functional f.

Proposition 4.0.12. *Let $\mathscr{V}$ be a vector space. Consider all the different choices of $\mathscr{Y}$ and $u_\circ$ that according to the unique identification process we described in proposition 4.0.11, all result in the same non-zero linear functional f. Among all these choices $\mathscr{Y}$ should be constantly the same hyperplane $\mathrm{Ker}[f]$, but $u_\circ$ could vary all over the shifted hyperplane $f_c^{\text{-}1}$ for $c = 1$.*

Proposition 4.0.13. *Let $\mathscr{V}$ be an n-dimensional vector space over the field $\mathbb{F}$, and let $f_1, f_2 \in \mathscr{V}^*$ be two linear functionals with the same kernel hyperplane. There would exist a scalar $\gamma \in \mathbb{F}$ such that $f_1(\cdot) = \gamma f_2(\cdot)$.*

Proof. Let's call the shared kernel hyperplane $\mathscr{Y} = \text{Ker}[f_i]$. As we mentioned before, one could fix a single vector u outside $\mathscr{Y}$ and it would follow that by building two bases $\mathscr{B}_1$ and $\mathscr{B}_2$ whose first members are respectively $\frac{1}{f_1(u)}u$ and $\frac{1}{f_2(u)}u$ and the rest being the same $n-1$ vectors for $\mathscr{B}_1$ and $\mathscr{B}_2$ lying inside $\mathscr{Y}$, for any vector v with respective representations $[v]_{\mathscr{B}_i} = [\alpha_u^{(i)}, *, *, \ldots]^\intercal$ we would have:

$$f_i(v) = \alpha_u^{(i)}$$

It can easily be seen that, since $\mathscr{B}_1$ and $\mathscr{B}_2$ are the same bases except for a rescaling their first member, regardless of the choice of v we would have:

$$\frac{\alpha_u^{(1)}}{f_1(u)} = \frac{\alpha_u^{(2)}}{f_2(u)}$$

And hence picking $\gamma = \frac{f_1(u)}{f_2(u)}$ we would have $f_1(\cdot) = \gamma f_2(\cdot)$. □

4.1 Linear functionals as Superposition of coordinate-pickers

As for any other linear transformation, uniquely defining $f(b_1), \ldots, f(b_n)$ would uniquely identify the behavior of the linear functional, and there is no constraint (of consistency or of any other matter) to the choice of scalars f would assign to any of b_i. Since, from proposition 3.2.2, the space of linear functionals $\mathscr{V}^* = \mathscr{L}(\mathscr{V}, \mathbb{F})$ is isomorphic to the space of row vectors of the size of $\mathscr{V}$'s dimensionality, for any choice of basis $\mathscr{B}$ for $\mathscr{V}$, an immediate natural basis for $\mathscr{V}^*$ would be the following:

$$\begin{aligned} [f_1]_{\mathscr{B}} &= [1, 0, 0, \ldots, 0] \\ [f_2]_{\mathscr{B}} &= [0, 1, 0, \ldots, 0] \\ &\vdots \\ [f_n]_{\mathscr{B}} &= [0, 0, 0, \ldots, 1] \end{aligned}$$

We will call these functionals coordinate-pickers:

Definition 4.1.1. Let $\mathscr{V}$ be an n-dimensional vector space over the field $\mathbb{F}$ and $\mathscr{B}$ be a basis for $\mathscr{V}$. We call the linear functional $f_i : \mathscr{V} \to \mathbb{F}$ that behaves in the following fashion the ***i*-th coordinate picker** with respect to the basis $\mathscr{B}$:

$$[u]_{\mathscr{B}} = \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \vdots \\ \alpha_n \end{bmatrix} \Rightarrow u \xrightarrow{f_i} \alpha_i$$

Definition 4.1.2. Let $\mathscr{V}$ be an n-dimensional vector space over the field $\mathbb{F}$ and $\mathscr{V}^*$ be its dual space. For any basis $\mathscr{B}$ for $\mathscr{V}$ we define the **dual basis** $\mathscr{B}^*$ to be the set of all coordinate pickers with respect to the basis $\mathscr{B}$, i.e. the n linear functionals: $\{f_1, f_2, \ldots, f_n\}$ where f_i is the linear functional assigning 1 to b_i and 0 to all the other b_j:

$$f_i(b_j) = \delta_{i,j}$$

where the latter is the Kronecker delta function.

Remark. We know from proposition 2.3.1 that, by regarding $\mathscr{V}$ and $\mathscr{V}^*$ as two arbitrary vector spaces sharing their dimensionality, any basis $\mathscr{B}$ for $\mathscr{V}$ together with any basis $\mathscr{C}$ for $\mathscr{V}^*$ would induce an isomorphism between $\mathscr{V}$ and $\mathscr{V}^*$. But the two have a more intimate relationship: Any basis $\mathscr{B}$ for $\mathscr{V}$ induces a natural basis for $\mathscr{V}^*$, i.e. the dual basis $\mathscr{B}^*$ containing all the coordinate pickers with respect to $\mathscr{B}$, and hence any basis $\mathscr{B}$ for $\mathscr{V}$ by itself induces an isomorphism between $\mathscr{V}$ and $\mathscr{V}^*$.

Of course as in the case of any other vector space the representation of linear functionals in terms of $\mathscr{B}^*$ would be a column vector, while their representation as linear transformations on $\mathscr{V}$ in terms of the basis $\mathscr{B}$ would be a row vector:

$$[f]_{\mathscr{B}}^{\mathsf{T}} = [f]_{\mathscr{B}^*}$$

This transition between row and column vectors is exactly what we meant by the isomorphism that the basis $\mathscr{B}$ induces between $\mathscr{V}$ and $\mathscr{V}^*$: For any vector u with representation in $\mathscr{B}$ being $[\alpha_1, \alpha_2, \ldots, \alpha_n]^{\mathsf{T}}$, take the matrix transpose of this column vector, and *look* at this row vector as the representation of a linear functional f. Of course we would have:

$$[u]_{\mathscr{B}} = [f]_{\mathscr{B}^*} \in \mathbb{F}^n$$

4.2 Nature of Duality

We now prove a series of results that are in the core of the nature of duality, and demonstrate clearly why we named $\mathscr{L}(\mathscr{V}, \mathbb{F})$ the *dual* of $\mathscr{V}$. Let $f_1, f_2, \ldots, f_n$ be a basis for $\mathscr{V}^*$.

Proposition 4.2.1 (Any basis for $\mathscr{V}^*$ is the dual of a unique basis for $\mathscr{V}$)**.** *Let $\mathscr{V}$ be a vector space, and $\mathscr{V}^*$ be its dual space. Let $\mathscr{C} = \{f_1, \ldots, f_n\}$ be any arbitrary basis for $\mathscr{V}^*$. In the manner of proposition 4.0.12 we could identify any of the f_i by a unique hyperplane $\mathscr{Y}_i$ and a variety of choices for $u_{\circ_i}$. We claim that one could find $u_{\circ_i}$s such that $\mathscr{B} = \{u_{\circ_1}, u_{\circ_2}, \ldots, u_{\circ_n}\}$ forms a basis for $\mathscr{V}$, and furthermore $\mathscr{C}$ would be equal to $\mathscr{B}^*$, the dual of the basis $\mathscr{B}$. Furthermore, this choice of $u_{\circ_i}$s is unique.*

Proof. We first prove that we could choose $u_{\circ_i}$s in a fashion that for any i, $u_{\circ_i}$ would be in the kernel of all f_j for $j \neq i$. In this case, obviously, f_is would satisfy:

$$f_i(u_{\circ_j}) = \delta_{i,j}$$

and as a result if we could show that $u_{\circ_j}$s form a basis for $\mathscr{V}$, $\mathscr{C}$ would immediately be its dual. Thus, we then will show that our choice of $u_{\circ_i}$s is unique and will lead to linearly independent vectors in $\mathscr{V}$. First notice that Ker$[f_i]$s are hyperplanes in $\mathscr{V}$. By proposition 4.0.13 we know that f_is can not share their respective kernels. We also notice that since $\mathscr{C}$ is a basis, none of the f_is could be the zero functional, and by proposition 4.0.10 all satisfy:

$$\dim(\mathscr{Y}_i) = \eta(f_i) = n - 1$$

We claim now, the proof of which is provided in proposition 4.2.2, that for any i the intersection of the kernels of the rest of f_js: $\mathscr{Z}_i = \bigcap_{j\neq i} \mathscr{Y}_j$ has dimensionality 1. Notice also that, again by proposition 4.2.2, the line $\mathscr{Z}_i$ can not fall inside $\mathscr{Y}_i$, and hence for every i, there exists a unique choice for $u_{\circ_i}$ to satisfy the following for all j:

$$f_j(u_{\circ_i}) = \delta_{i,j}$$

that leaves us to just prove that with this choice of $u_{\circ_i}$s, they would form a linearly independent set of vectors. Let there be a vanishing linear combination of $u_{\circ_i}$s:

$$\sum_{i=1}^{n} \alpha_i u_{\circ_i} = 0$$

then for every j we would have:

$$f_j(\sum_{i=1}^{n} \alpha_i u_{\circ_i}) = \sum_{i=1}^{n} \alpha_i f_j(u_{\circ_i}) = \sum_{i=1}^{n} \alpha_i \delta_{i,j} = \alpha_j = 0$$

which shows $\mathscr{B} = \{u_{\circ_1}, u_{\circ_2}, \ldots, u_{\circ_n}\}$ will be a basis for $\mathscr{V}$ and from the behavior of f_i over the members of $\mathscr{B}$ it results that $\mathscr{C}$ is the dual of the basis $\mathscr{B}$, and the proof is complete. □

Corollary 4.2.1. In matrix terminology, this result yields to a weak statement about square matrices. If the rows of a matrix A are linearly independent (form a basis for the dual space of $\mathbb{F}^n$) there exists a unique ordered set S of vectors in $\mathbb{F}^n$ with respect to which, those rows are coordinate pickers. Putting the vectors in S as the columns of a matrix B, this means that B is the inverse of the matrix A.

Proposition 4.2.2. *Let $\mathscr{V}$ be an n-dimensional vector space, and $\mathscr{W}_1, \mathscr{W}_2, \ldots, \mathscr{W}_k$ be k distinct hyperplanes, i.e $(n-1)$-dimensional subspaces, in $\mathscr{V}$. The following holds:*

$$\dim\left(\bigcap_{i=1}^{k} \mathscr{W}_i\right) = n - k$$

Proof. We recall the inclusion exclusion theorem for dimensionalities from proposition 2.2.4:

$$\dim(\bigcup_{i=1}^{k} \mathscr{W}_i) = \sum_{i=1}^{k} \dim(\mathscr{W}_i) - \sum_{i<j} \dim(\mathscr{W}_i \cap \mathscr{W}_j) + \sum_{i<j<t} \dim(\mathscr{W}_i \cap \mathscr{W}_j \cap \mathscr{W}_t) - \ldots + (-1)^{k+1} \dim(\bigcap_{i=1}^{k} \mathscr{W}_i)$$

or more succinctly:

$$\dim(\bigcup_{i=1}^{k} \mathscr{W}_i) = \sum_{l=1}^{k} (-1)^l \left[\sum_{1 \le t_1 < t_2 < \ldots < t_l \le k} \dim(\bigcap_{i=1}^{l} \mathscr{W}_{t_i}) \right]$$

and notice that since $\mathscr{W}_i$ are distinct hyperplanes we have:

$$\mathscr{W}_i \cup \mathscr{W}_j = \mathscr{V}$$

and hence:

$$\dim(\bigcup_{i=1}^{k} \mathscr{W}_i) = n$$

We now perform induction on k. For $k = 1$ the result is obvious. For $k = 2$ we have:

$$\dim(\mathscr{W}_1 \cap \mathscr{W}_2) = \dim(\mathscr{W}_1) + \dim(\mathscr{W}_2) - \dim(\mathscr{W}_1 \cup \mathscr{W}_2) = 2(n-1) - n = n - 2$$

Now let's assume that the result holds for any number of less than k hyperplanes. The left hand side of the inclusion exclusion formula would be just n, and all the summands in the right hand side are instances of the induction hypothesis, except for the one with $l = k$ which is the quantity under investigation:

$$n = \sum_{l=1}^{k-1} (-1)^{l+1} \left[\sum_{1 \le t_1 < t_2 < \ldots < t_l \le k} \dim(\bigcap_{i=1}^{l} \mathscr{W}_{t_i}) \right] + (-1)^{k+1} \dim\left(\bigcap_{i=1}^{k+1} \mathscr{W}_i \right)$$

Using the induction hypothesis we would get:

$$\begin{aligned}
(-1)^{k+1} \dim\left(\bigcap_{i=1}^{k} \mathscr{W}_i \right) &= n + \sum_{l=1}^{k-1} (-1)^l \left[\sum_{1 \le t_1 < t_2 < \ldots < t_l \le k} \dim(\bigcap_{i=1}^{l} \mathscr{W}_{t_i}) \right] \\
&= n + \sum_{l=1}^{k-1} (-1)^l \left[\sum_{1 \le t_1 < t_2 < \ldots < t_l \le k} (n-l) \right] \\
&= n + \sum_{l=1}^{k-1} (-1)^l \binom{k}{l} (n-l) \\
&= n + n \left[\sum_{l=1}^{k-1} (-1)^l \binom{k}{l} \right] - \sum_{l=1}^{k} (-1)^l \binom{k}{l} l
\end{aligned}$$

But we know from the binomial expansion that:

$$(1+x)^k = \sum_{l=0}^{k} \binom{k}{l} x^l$$

and hence:

$$0 = (1-1)^k = \sum_{l=0}^{k} \binom{k}{l} (-1)^l$$

which leads to:

$$\sum_{l=1}^{k-1}\binom{k}{l}(-1)^l = 0 - \binom{k}{k}(-1)^k - \binom{k}{0}(-1)^0 = (-1)^{k+1} - 1$$

We also know for $k > 1$ that:

$$k(1+x)^{k-1} = \sum_{l=0}^{k} l\binom{k}{l}x^{l-1}$$

multiplying both sides by x we get:

$$kx(1+x)^{k-1} = \sum_{l=0}^{k} l\binom{k}{l}x^{l}$$

and hence:

$$\sum_{l=1}^{k}(-1)^l\binom{k}{l}l = 0 - k\binom{k}{k}(-1)^k = k(-1)^{k+1}$$

Getting back to our manipulations of the inclusion exclusion formula, and plugging in these results we get:

$$\begin{aligned}(-1)^{k+1}\dim\left(\bigcap_{i=1}^{k}\mathscr{W}_i\right) &= n + n\left[\sum_{l=1}^{k-1}(-1)^l\binom{k}{l}\right] - \sum_{l=1}^{k}(-1)^l\binom{k}{l}l \\ &= n + n[(-1)^{k+1} - 1] - k(-1)^{k+1} \\ &= n + n(-1)^{k+1} - n - k(-1)^{k+1} \\ &= (n-k)(-1)^{k+1}\end{aligned}$$

and by canceling $(-1)^{k+1}$ from both sides, the proof is complete. □

Now let's think about $\mathscr{L}(\mathscr{V}^*, \mathbb{F})$, the dual space of $\mathscr{V}^*$. What would be a linear functional over $\mathscr{V}^*$? Obviously fixing a vector $u \in \mathscr{V}$ and assigning to any functional f in $\mathscr{V}^*$ the scalar $f(u)$, would result into a linear functional over $\mathscr{V}^*$. But what other kinds of linear functionals are there over $\mathscr{V}^*$? It turns out, nothing! In other words all the linear functionals over $\mathscr{V}^*$ are equivalent to the *evaluation* map for some fixed vector in $\mathscr{V}$.

Proposition 4.2.3. *Let $\ell \in \mathscr{L}(\mathscr{V}^*, \mathbb{F})$ be an arbitrary linear functional over the dual space of the arbitrary vector space $\mathscr{V}$ over the field $\mathbb{F}$. There exists a unique vector $u \in \mathscr{V}$ such that for all $f \in \mathscr{V}^*$ we have:*

$$\ell(f) = f(u)$$

Proof. Now let $\mathscr{C} = \{\ell_1, \ldots, \ell_n\}$ be a basis for the double dual space $\mathscr{V}^{**} = \mathscr{L}(\mathscr{V}^*, \mathbb{F})$. By proposition 4.2.1, $\mathscr{C}$ is the dual of some basis for $\mathscr{V}^*$ which in its own turn is a dual of some basis $\mathscr{B}$ for $\mathscr{V}$. But what exactly is the double dual of a basis? Let $\mathscr{B} = \{b_1, \ldots, b_n\}$, and let $\mathscr{B}^* = \{f_1, \ldots, f_n\}$. We know that f_is are coordinate pickers with respect to $\mathscr{B}$, and ℓ_is are coordinate pickers with respect to $\mathscr{B}^*$. But there is a subtle twist here! We have seen that by fixing a basis $\mathscr{B}$ for $\mathscr{V}$, a

functional f over $\mathscr{V}$ whose representation in terms of the dual basis $\mathscr{B}^*$ is $[\xi_1, \xi_2, \ldots, \xi_n]^\intercal$, is one that is the linear combination of coordinate pickers with weights being the ξ_i:

$$f = \xi_1 f_1 + \xi_2 f_2 + \ldots + \xi_n f_n$$

As a result f is a linear functional that assigns to b_i the scalar ξ_i. Or equivalently b_i is a vector that assigns to the functionals f its i-th coordinate in terms of the dual basis, i.e. ξ_i, and thus has the exact same behavior as that of the i-th *coordinate picker* in $\mathscr{V}^*$ with respect to $\mathscr{B}^*$: a function that assigns to any such f as above the scalar ξ_i! We could have gone through the argument from another direction. Let's fix a basis $\mathscr{T}$ for $\mathscr{V}^*$. For any f, b_i is the *unique* i-th member of that unique basis of $\mathscr{V}$ whose dual is $\mathscr{T}$, and has the property that $f(b_i)$ is the i-th coordinate of f in terms of $\mathscr{T}$. As a result any possible functional over $\mathscr{V}^*$ that, talking in terms of $\mathscr{T}$ (for any possible choice of $\mathscr{T}$), is a linear combination of coordinate pickers, behaves in fact exactly the same as uniquely finding the corresponding linear combination of b_i (which is a unique vector in $\mathscr{V}$) and applying f to that vector. In less verbal and more mathematical terms, ℓ_i is such that we have:

$$\ell_i(f_j) = \delta_{i,j}$$

and f_j is such that for all i we have:

$$f_j(b_i) = \delta_{i,j}$$

As a result for all i, j we have:

$$\ell_i(f_j) = f_j(b_i)$$

and hence if the representation of ℓ in terms of the double dual basis $\mathscr{B}^{**}$ is $[\zeta_1, \zeta_2, \ldots, \zeta_n]^\intercal$, we would have:

$$\ell(f) = \sum_{i=1}^{n} \zeta_i \ell_i(f_j) = \sum_{i=1}^{n}\sum_{j=1}^{n} \zeta_i \xi_j \ell_i(f_j) = \sum_{i=1}^{n}\sum_{j=1}^{n} \zeta_i \xi_j f_j(b_i) = f\left(\sum_{i=1}^{n} \zeta_i b_i\right) = f(u)$$

where u is such that $[u]_{\mathscr{B}} = [\zeta_1, \zeta_2, \ldots, \zeta_n]^\intercal$, the exact same representation as that of ℓ in terms of $\mathscr{B}^{**}$, and the proof is complete. □

Any two vector spaces of the same dimensionality are isomorphic (through fixing some basis), but the isomorphism between $\mathscr{V}$ and $\mathscr{V}^{**}$ is way more intimate, since it is independent of any choice of a basis. Any vector $u \in \mathscr{V}$ is uniquely mapped, regardless any choice of basis, to a unique member ℓ of $\mathscr{V}^{**}$, since the two have constantly the very same representation in terms of any choice of basis in their respective vector spaces: $[u]_{\mathscr{B}}$ and $[\ell]_{\mathscr{B}^{**}}$ are the very same members of $\mathbb{F}^n$ for *any* choice of $\mathscr{B}$. As a result one could think of $\mathscr{V}^{**}$ as indistinguishable from the original vector space $\mathscr{V}$:

Proposition 4.2.4. *For any vector space $\mathscr{V}$, the double dual space $\mathscr{V}^{**}$ is **canonically** isomorphic to $\mathscr{V}$.*

It can easily be seen that this kind of relationship does not hold between $\mathscr{V}$ and the first dual $\mathscr{V}^*$. The dual space does have the nice property that for any linear functional f, its representations in terms of $\mathscr{B}$ (as a linear transformation) and in terms of $\mathscr{B}^*$ (as a vector in $\mathscr{V}^*$) are matrix transposes of each other. But these two vectors (the latter a column vector and the former a row

vector) are completely dependent on the choice of basis.
As it can be seen from the fact that we have made extensive use of the property that the relationship between a basis for $\mathscr{V}$ and its dual basis in $\mathscr{V}^*$ is a bijective map, the duality of $\mathscr{V}$ and $\mathscr{V}^*$ can be best be understood through the duality of their bases:

$$\mathscr{B} = \{b_1, \ldots, b_n\}$$

$$\mathscr{B}^* = \{f_1, \ldots, f_n\}$$

f_is are coordinate pickers in $\mathscr{V}$ with respect to $\mathscr{B}$, and b_i are coordinate pickers in $\mathscr{V}^*$ with respect to $\mathscr{B}^*$. For any functional $f \in \mathscr{V}^*$, f is a creature that is identified by a bunch of n scalar weights ξ_i (which happen to be the coordinates of f in terms of $\mathscr{B}^*$), and while marching all over the vector space $\mathscr{V}$, f assigns to any vector u a linear combination of u's coordinates (in terms of $\mathscr{B}$) with weights being ξ_i. And this happens to cover all the possible linear transformations from $\mathscr{V}$ to $\mathbb{F}$. Symmetrically, any vector u is a creature that is identified by a bunch of n scalar weights α_i (which happen to be the coordinates of u in terms of $\mathscr{B}$), and while marching all over the vector space $\mathscr{V}^*$ assigns to any vector f a linear combination of f's coordinates (in terms of $\mathscr{B}^*$) with weights being α_i. And this happens to cover all the possible linear transformations from $\mathscr{V}^*$ to $\mathbb{F}$.

4.3 Annihilator spaces and Transformation$^{\mathrm{t}}$ transposes

Fixing a linear functional f over a vector space $\mathscr{V}$, we know that the set of all vectors that vanish to zero under application to f, is the kernel of a linear transformation (f) and hence is a subspace of $\mathscr{V}$. But in the light of canonical isomorphism of $\mathscr{V}^{**}$ and $\mathscr{V}$, we made some observations at the end of the last section, which rises the question that what about fixing a vector u, and looking at the set of linear functionals f that vanish to zero under application by u. This is obviously the kernel of a linear transformation (that unique member of $\mathscr{V}^{**}$ corresponding to u) and hence is a subspace of $\mathscr{V}^*$, which in a similar fashion as in proposition 4.0.10 has a dimensionality of either n or $n-1$. This, regarding u as a linear functional over $\mathscr{V}^*$, is its kernel, and we will, just for the sake of conformity with the literature, call it the *annihilator* of u in $\mathscr{V}^*$.

Definition 4.3.1. Let $\mathscr{V}$ be a vector space of which $\mathscr{W}$ is a subspace. We define the **annihilator** of $\mathscr{W}$, denoted by $\mathscr{W}^\circ$, to be that subspace of the dual space $\mathscr{V}^*$ containing all the linear functionals f vanishing to zero when applied to any of the vectors in $\mathscr{W}$. The fact that $\mathscr{W}^\circ$ is a subspace follows from the arguments preceding these lines.

Obviously picking any basis $\mathscr{B}_w$ for $\mathscr{W}$, the annihilator is the intersection of the corresponding hyperplanes in $\mathscr{V}^*$ each being the single-vector annihilator of a member of $\mathscr{B}_w$. Using proposition 4.2.2, if $\mathscr{W}$ has dimensionality k, its annihilator, being the intersection of k distinct hyperplanes, will have dimensionality $n-k$, where n is the dimensionality of $\mathscr{V}^*$:

Proposition 4.3.1. *Let $\mathscr{V}$ be a vector space of which $\mathscr{W}$ is a subspace. The following holds regarding the dimensionality of the annihilator of $\mathscr{W}$:*

$$\dim(\mathscr{W}^\circ) + \dim(\mathscr{W}) = \dim(\mathscr{V})$$

It obviously follows from duality, that the annihilator of $\mathscr{W}^\circ$ would be $\mathscr{W}$ itself:

Proposition 4.3.2. *Let $\mathscr{V}$ be a vector space. For any subspace $\mathscr{W}$ of $\mathscr{V}$ we have:*

$$\mathscr{W}^{\circ^\circ} = \mathscr{W}$$

We now turn our attention to the problem of transformation transposes. We are already familiar with the well defined notion of matrix transposes. We, using the notion of dual spaces, define in this section the notion of transposes of linear transformations, and observe that it is intuitively consistent with matrix transposes. Throughout the rest of this section we fix two vector spaces $\mathscr{V}$ and $\mathscr{W}$, of dimensionality n and m respectively, over the same field $\mathbb{F}$, and two respective bases $\mathscr{B}_v$ and $\mathscr{B}_w$ for $\mathscr{V}$ and $\mathscr{W}$. We already know that any linear transformation $\mathrm{T}: \mathscr{V} \to \mathscr{W}$ is uniquely identified by its behavior on the members of $\mathscr{B}_v$. This corresponds to decomposing the matrix representation $[\mathrm{T}]_{\mathscr{B}_v\mathscr{B}_w}$ into column vectors, each being the representation of $\mathrm{T}b_i$ in terms of $\mathscr{B}_w$. But now that we are comfortable with the notion of linear functionals, we could decompose T in another fashion. If we let $f_1, f_2, \ldots, f_m$ be coordinate pickers with respect to $\mathscr{B}_w$ (and as a result form a basis for $\mathscr{W}^*$), obviously for any $u \in \mathscr{V}$ we could write:

$$[\mathrm{T}u]_{\mathscr{B}_w} = \begin{bmatrix} f_1(\mathrm{T}u) \\ f_2(\mathrm{T}u) \\ \vdots \\ f_m(\mathrm{T}u) \end{bmatrix} \in \mathbb{F}^m$$

As a result any linear transformation T induces m functionals over its departure space $\mathscr{V}$, namely $t_1, t_2, \ldots, t_m$ in $\mathscr{V}^*$, such that for any $u \in \mathscr{V}$ we have:

$$[\mathrm{T}u]_{\mathscr{B}_w} = \begin{bmatrix} t_1(u) \\ t_2(u) \\ \vdots \\ t_m(u) \end{bmatrix}$$

The t_is would be defined as:

$$t_i: \mathscr{V} \to \mathbb{F}$$
$$t_i = f_i \circ \mathrm{T}$$

As we know by now, when represented in terms of $\mathscr{B}_v$, all linear functionals over $\mathscr{V}$ correspond to row vectors of scalars of size n; and hence any linear transformation T, having fixed the bases $\mathscr{B}_w$ and $\mathscr{B}_v$ for the destination and departure spaces, induces m row vectors of scalars of size n, which correspond exactly to the row decomposition of $[\mathrm{T}]_{\mathscr{B}_v\mathscr{B}_w}$.
The important thing that happened when we defined the t_is was that we took the members of a basis for $\mathscr{W}^*$ (to be exact, $\mathscr{B}_w^*$), and assigned to them elements in $\mathscr{V}^*$. We know that any such choice corresponds to a linear transformation from $\mathscr{W}^*$ to $\mathscr{V}^*$. It also easy to see that since we mapped $f_i \in \mathscr{B}_w^*$ to $f_i \circ \mathrm{T}$, *any* other functional $f \in \mathscr{W}^*$ would be mapped by this uniquely defined transformation to $f \circ \mathrm{T} \in \mathscr{V}^*$. We call this linear transformation the *transpose* of T. In fact the behavior of the transpose transformation, is completely independent of the choice of bases $\mathscr{B}_w, \mathscr{B}_v$.

Definition 4.3.2. Let $\mathrm{T}\colon \mathscr{V} \to \mathscr{W}$ be a linear transformation. We define the **transpose** of T, and denote it by T^{t}, to by the linear transformation $\mathrm{T}^{\mathrm{t}}\colon \mathscr{W}^* \to \mathscr{V}^*$, that assigns to any functional $f \in \mathscr{W}^*$ the functional $f \circ \mathrm{T} \in \mathscr{V}^*$. The verification of T^{t} being a linear transformation is straightforward.

In more graphical terms, we have a linear transformation $\mathscr{V} \xrightarrow{\mathrm{T}} \mathscr{W}$, the transpose of T is the mapping that gets a functional $\mathscr{W} \xrightarrow{f} \mathbb{F}$ and sends it to another functional $\mathscr{V} \xrightarrow{\mathrm{T}^{\mathrm{t}} f} \mathbb{F}$ in the following manner:

$$\begin{array}{ccc} \square & \xrightarrow{\quad\mathrm{T}\quad} & \lozenge \\ \left[\square \xrightarrow{\mathrm{T}} \lozenge \xrightarrow{f} \mathbb{F}\right] : t & \xleftarrow{\quad\mathrm{T}^{\mathrm{t}}\quad} & f : \left[\lozenge \xrightarrow{f} \mathbb{F}\right] \end{array}$$

Let's now take a second transpose, and use the last observations of the previous section about the nature of duality:

$$\begin{array}{ccc} \square^* & \xleftarrow{\quad\mathrm{T}^{\mathrm{t}}\quad} & \lozenge^* \\ \left[\square^* \xrightarrow{u} \mathbb{F}\right] : u & \xrightarrow{\quad\mathrm{T}^{\mathrm{t}^{\mathrm{t}}}\quad} & v : \left[\lozenge^* \xrightarrow{\mathrm{T}^{\mathrm{t}}} \square^* \xrightarrow{u} \mathbb{F}\right] \end{array}$$

But what does $\mathrm{T}^{\mathrm{t}^{\mathrm{t}}}$ assign to $u \in \mathscr{V}$? We saw that T^{t} assigns to a functional f over $\mathscr{W}$ the functional t over $\mathscr{V}$ such that for every $u \in \mathscr{V}$ we have $t(u) = f(\mathrm{T}u)$. In the same fashion $\mathrm{T}^{\mathrm{t}^{\mathrm{t}}}$ assigns to a functional u over $\mathscr{V}^*$ (where u by duality is a vector in $\mathscr{V}$) the functional v over $\mathscr{W}^*$ (a vector in $\mathscr{W}$) such that for every $f \in \mathscr{W}$ we have $v(f) = u(\mathrm{T}^{\mathrm{t}} f)$, where $v(f)$ is the result of looking at v as a functional over $\mathscr{W}^*$ and f as a vector in $\mathscr{W}^*$, and its scalar value if $f(v)$:

$$f(v) = (\mathrm{T}^{\mathrm{t}} f)(u) = f(\mathrm{T}u)$$

which means that for every $u \in \mathscr{V}$, $v = \mathrm{T}^{\mathrm{t}^{\mathrm{t}}} u$ is such that for every functional $f \in \mathscr{W}^*$ we have $f(v) = f(\mathrm{T}u)$. This means that the vector $v - \mathrm{T}u$ is in the kernel of *any single* linear functional f, which is impossible unless $v - \mathrm{T}u$ is zero, which completes the proof of the following claim:

Proposition 4.3.3. *For any linear transformation* $\mathrm{T}\colon \mathscr{V} \to \mathscr{W}$*, with transposition defined as above, we have:*

$$\mathrm{T}^{\mathrm{t}^{\mathrm{t}}} = \mathrm{T}$$

Proposition 4.3.4. *Let* $\mathrm{T}\colon \mathscr{V} \to \mathscr{W}$ *be a linear transformation, and* T^{t} *be its transpose defined as above.* $\mathrm{Ker}[\mathrm{T}^{\mathrm{t}}]$ *is a subspace of* $\mathscr{W}^*$ *by definition. We claim that* $\mathrm{Ker}[\mathrm{T}^{\mathrm{t}}]$ *is in fact the annihilator of* $\mathrm{Im}[\mathrm{T}]$ *which is a subspace of* $\mathscr{W}$*:*

$$\mathrm{Ker}[\mathrm{T}^{\mathrm{t}}] = \mathrm{Im}[\mathrm{T}]^{\circ}$$

and also:

$$\mathrm{Im}[\mathrm{T}^{\mathrm{t}}] = \mathrm{Ker}[\mathrm{T}]^{\circ}$$

Proof. The kernel of the transpose of T is, as we defined the transpose, the space of those functionals $f\colon \mathscr{W} \to \mathbb{F}$, such that the functional over $\mathscr{V}$ that results from $f \circ \mathrm{T}$, i.e the functional $\mathscr{V} \xrightarrow{\mathrm{T}} \mathscr{W} \xrightarrow{f} \mathbb{F}$, is the zero functional. The fact that all such f completely vanish to zero all of $\mathrm{Im}[\mathrm{T}]$ is obvious, and as a result $\mathrm{Ker}[\mathrm{T}^{\mathrm{t}}] \subseteq \mathrm{Im}[\mathrm{T}]^{\circ}$. But let's take an arbitrary functional $g \in \mathscr{W}^*$ from $\mathrm{Im}[\mathrm{T}]^{\circ}$. What

would $\mathrm{T}^{\mathrm{t}}g\colon \mathscr{V}\to\mathbb{F}$ be? For any $u\in\mathscr{V}$ we have: $(\mathrm{T}^{\mathrm{t}}g)u = g(\mathrm{T}u)$ which by the choice of g being from $\mathrm{Im}[\mathrm{T}]^\circ$ would be zero, and hence $\mathrm{Im}[\mathrm{T}]^\circ \subseteq \mathrm{Ker}[\mathrm{T}^{\mathrm{t}}]$ and the proof of the first equation is complete. We now write the first equality for T^{t}:

$$\mathrm{Ker}[\mathrm{T}^{\mathrm{t}^{\mathrm{t}}}] = \mathrm{Im}[\mathrm{T}^{\mathrm{t}}]^\circ$$

where from proposition 4.3.3 the left hand side is $\mathrm{Ker}[\mathrm{T}]$, and also, by proposition 4.3.2, taking annihilators of both sides yields the second equality:

$$\mathrm{Ker}[\mathrm{T}]^\circ = \mathrm{Im}[\mathrm{T}^{\mathrm{t}}]^{\circ^\circ} = \mathrm{Im}[\mathrm{T}^{\mathrm{t}}]$$

□

Proposition 4.3.5. *Let $\mathrm{T}\colon \mathscr{V}\to\mathscr{W}$ be a linear transformation, and T^{t} be its transpose. We have:*

$$\rho(\mathrm{T}) = \rho(\mathrm{T}^{\mathrm{t}})$$

Proof. The wanted equality immediately follows from Rank-Nullity theorem and the results of propositions 4.3.4 and 4.3.1:

$$\rho(\mathrm{T}^{\mathrm{t}}) = \dim(\mathrm{Im}[\mathrm{T}^{\mathrm{t}}]) = \dim(\mathrm{Ker}[\mathrm{T}]^\circ) = \dim(\mathscr{V}) - \dim(\mathrm{Ker}[\mathrm{T}]) = \dim(\mathrm{Im}[\mathrm{T}]) = \rho(\mathrm{T})$$

□

And finally here we prove the liberating relationship between the transformation transpose and the matrix transpose:

Proposition 4.3.6. *: Let $\mathrm{T}\colon \mathscr{V}\to\mathscr{W}$ be a linear transformation, and T^{t} be its transpose defined as above. For any choice of bases $\mathscr{B}_v$ and $\mathscr{B}_w$, respectively for $\mathscr{V}$ and $\mathscr{W}$, we have:*

$$[\mathrm{T}^{\mathrm{t}}]_{\mathscr{B}_w^*\mathscr{B}_v^*} = [\mathrm{T}]^{\intercal}_{\mathscr{B}_v\mathscr{B}_w}$$

where $\cdot^{\mathrm{t}}$ denotes transformation transpose, $\cdot^{\intercal}$ denotes matrix transpose, and $\mathscr{B}_\cdot^*$ denotes dual bases.

Proof. We just need proposition 3.2.3 to write the following for any functional f:

$$[\mathrm{T}^{\mathrm{t}}f]_{\mathscr{B}_v} = [f\circ\mathrm{T}]_{\mathscr{B}_v^*} = [f]_{\mathscr{B}_w}[\mathrm{T}]_{\mathscr{B}_v\mathscr{B}_w}$$

and hence:

$$[\mathrm{T}^{\mathrm{t}}f]_{\mathscr{B}_v^*} = [\mathrm{T}^{\mathrm{t}}f]^{\intercal}_{\mathscr{B}_v} = [\mathrm{T}]^{\intercal}_{\mathscr{B}_v\mathscr{B}_w}[f]^{\intercal}_{\mathscr{B}_w}$$

We know denote the members of $\mathscr{B}_w^*$ by $\{f_1, f_2, \ldots, f_m\}$, and derive the wanted equality as follows:

$$\begin{aligned}
[\mathrm{T}^{\mathrm{t}}]_{\mathscr{B}_w^*\mathscr{B}_v^*} &= \begin{bmatrix}[\mathrm{T}^{\mathrm{t}}f_1]_{\mathscr{B}_v^*} & [\mathrm{T}^{\mathrm{t}}f_2]_{\mathscr{B}_v^*} & \cdots & [\mathrm{T}^{\mathrm{t}}f_m]_{\mathscr{B}_v^*}\end{bmatrix}\\
&= \begin{bmatrix}[\mathrm{T}]^{\intercal}_{\mathscr{B}_v\mathscr{B}_w}[f_1]^{\intercal}_{\mathscr{B}_w} & [\mathrm{T}]^{\intercal}_{\mathscr{B}_v\mathscr{B}_w}[f_2]^{\intercal}_{\mathscr{B}_w} & \cdots & [\mathrm{T}]^{\intercal}_{\mathscr{B}_v\mathscr{B}_w}[f_m]^{\intercal}_{\mathscr{B}_w}\end{bmatrix}\\
&= [\mathrm{T}]^{\intercal}_{\mathscr{B}_v\mathscr{B}_w}\begin{bmatrix}[f_1]^{\intercal}_{\mathscr{B}_w} & [f_2]^{\intercal}_{\mathscr{B}_w} & \cdots & [f_m]^{\intercal}_{\mathscr{B}_w}\end{bmatrix}\\
&= [\mathrm{T}]^{\intercal}_{\mathscr{B}_v\mathscr{B}_w}\mathrm{I}_{m\times m}\\
&= [\mathrm{T}]^{\intercal}_{\mathscr{B}_v\mathscr{B}_w}
\end{aligned}$$

where in the last simplification we have used the fact that f_is, being members of the dual basis $\mathscr{B}_w^*$, are coordinate pickers in $\mathscr{W}$ with respect to the basis $\mathscr{B}_w$. □

Corollary 4.3.1. For any $m\times n$ matrix A with scalar entities, the row rank and column rank of A are equal.

Chapter 5

The zero fever and the basic matrix decompositions

We just love it when we face organized chunks of zero in matrix representations. That this provides us with computational efficiency is fair obvious: We know that multiplying two $n \times n$ matrices takes $2n^3$ flops. Now assume I have the, seemingly benign, knowledge that one specific entry of one of the matrices is zero. It would reduce the computation cost to $2n^3 - n$ by not performing the n flops of $0 \times \cdot$ we would otherwise perform[1]. If a whole column is zero we get to save $O(n^2)$ flops, and if the bulk of zeros we have "prior knowledge" about gets to $O(n^2)$ we actually can reduce the cost by a factor (say for example of both matrices are upper (lower) triangular we get to reduce $2n^3$ flops to $\frac{n^3}{2}$, a reduction by a factor of 4. If this knowledge of existence of zeros goes to the extent of knowing that a specific one of the matrices is diagonal we could finish the multiplication by n^2 flops and by n flops if we know that both are diagonal!

The study of diagonalization and triangularization of matrices as well as the study of eigenvalues and canonical forms could be viewed from this viewpoint. However, the existence of "well-organized" "bulk"s of zero in the matrix representation of a linear transformation, is not just favored due to massive computational savings, but also for the fact that they usually tell us a fair lot about the behavior of a linear transformation. This can go to the extent of enabling us to describe the behavior of linear transformations in one single verbal sentence. In this section we build some elementary results that we are going to use extensively in the treatise of canonical forms.

5.1 Some basic properties of envelopes

Definition 5.1.1. Let $\mathscr{B} = \{b_1, \ldots, b_n\}$ be an arbitrary basis for the vector space $\mathscr{V}$. We define the function

$$\zeta_{\mathscr{B}} : \mathscr{B} \to \mathbb{Z}^{\geq 0}$$

[1] Of course for the sake of finding one single zero, or for that matter any larger number of zeros we might find, "attain"ing such knowledge, if it is not "given" somehow, would not be worth its computational cost, since it takes n^2 flops to go through all the entries. But this is, anyway, an exaggerated example to see how much we could save from the knowledge of existence of zeros.

such that for any $u \in \mathscr{V}$, $\zeta_{\mathscr{B}}(u)$ denotes the number of leading zeros in $[u]_{\mathscr{B}}$. To be more exactly if $[u]_{\mathscr{B}} = [\alpha_1, \alpha_2, \ldots, \alpha_n]^{\intercal}$, then $\zeta_{\mathscr{B}}(u) = k$ where k is such that:

$$\forall i \leq k \ : \ \alpha_i = 0$$

and $\alpha_{k+1} = 0$.

Obviously from this definition one can infer that for any vector $u \in \mathscr{V}$ we have:

$$u \in \text{span}\{b_i\}_{i=\zeta_{\mathscr{B}}(u)+1}^{n}$$

Proposition 5.1.1. *Let $\mathscr{B}$ be an arbitrary basis for the n-dimensional vector space $\mathscr{V}$. The non-zero vectors $u_1, u_2, \ldots, u_k$ form a linearly independent set if (and not only if):*

$$\forall i \neq j : \ \zeta_{\mathscr{B}}(u_i) \neq \zeta_{\mathscr{B}}(u_j)$$

Proof. Without loss of generality we can assume: $\zeta_{\mathscr{B}}(u_1) < \zeta_{\mathscr{B}}(u_2) < \ldots < \zeta_{\mathscr{B}}(u_k) < n$. If there exists some vanishing linear combination of u_is:

$$\sum_{i=1}^{k} \alpha_i u_i = 0$$

Looking at the $(\zeta_{\mathscr{B}}(u_1)+1)$-th coordinate (component of both sides in the direction of $b_{\zeta_{\mathscr{B}}(u_1)+1}$) we could easily see that: $\alpha_1 = 0$. Continuing in the same fashion by looking at the $(\zeta_{\mathscr{B}}(u_i)+1)$-th coordinate at both sides we could get $\alpha_i = 0$. □

Proposition 5.1.2. *Let $\mathscr{B}$ be an arbitrary basis for the n-dimensional vector space $\mathscr{V}$, and $u_1, u_2, \ldots, u_k$ be an arbitrary set of non-zero vectors in $\mathscr{V}$. The dimensionality of this set of vectors is bounded tightly by the following quantities:*

$$s \leq \dim^{\text{span}}\{u_i\}_{i=1}^{k} \leq n - \min_i\{\zeta_{\mathscr{B}}(u_i)\}$$

where s is the largest number of u_is one could select with distinct $\zeta_{\mathscr{B}}(u_i)$.

Proof. Let $\mathrm{A} = \begin{bmatrix}[u_1]_{\mathscr{B}}[u_2]_{\mathscr{B}} \ldots [u_k]_{\mathscr{B}}\end{bmatrix}$ and let $\min_i\{\zeta_{\mathscr{B}}(u_i) \geq t$ for all i. A is a $n \times k$ matrix. We know that $\dim^{\text{span}}$ is equal to the column rank of A, which in its own turn is equal to the row rank of A. Since the first t rows of A are completely zero, it will obviously follow that it can not have row rank larger than $n - t$. This bound is obviously tight, since if $t = 0$, $\dim^{\text{span}}\{u_i\}_{i=1}^{k}$ can be as large as n. We now prove the lower bound. Let $u_{\pi_1}, u_{\pi_2}, \ldots, u_{\pi_s}$ be s of the u_is with distinct $\zeta_{\mathscr{B}}(u_{\pi_i})$. The inequality follows from proposition 5.1.1:

$$s = \dim^{\text{span}}\{u_{\pi_i}\}_{i=1}^{s} \leq \dim^{\text{span}}\{u_i\}_{i=1}^{k}$$

This bound is also tight, since if all u_i have distinct $\zeta_{\mathscr{B}}(u_i)$ then by the same proposition $\dim^{\text{span}}\{u_i\}_{i=1}^{k}$ will be exactly equal to $s = k$. □

Corollary 5.1.1. Let A be an $n \times n$ triangular matrix with k non-zero diagonal entries. We will have:

$$k \leq \rho(\mathrm{A}) \leq n$$

As a result a square triangular matrix is non-singular if and only if it is non-zero all over its diagonal.

Proposition 5.1.3. *Let* A *be an* $n \times n$ *diagonal matrix with* k *non-zero diagonal entries. We will have:*

$$\rho(\mathrm{A}) = k$$

Proposition 5.1.4. *The inverse of a lower (upper) triangular matrix is again a lower (upper) triangular. The multiplication of two lower (upper) triangular matrices is again a lower (upper) triangular matrix. Furthermore the diagonal of the multiplication is equal to the entry-wise multiplication of the respective diagonals.*

5.2 Invariant subspaces and Invariant direct sum decompositions

Definition 5.2.1. Let $\mathrm{T}: \mathscr{V} \to \mathscr{V}$ be a linear transformation and $\mathscr{W}$ be a subspace of $\mathscr{V}$. Obviously T imposes a linear transformation from $\mathscr{W}$ to $\mathscr{V}$ which we will call its **confinement** to the subspace $\mathscr{W}$, and denote it by $\mathrm{T}_{\mathscr{W}}$.

Proposition 5.2.1. *Let* $\mathrm{T}: \mathscr{V} \to \mathscr{V}$ *be a linear transformation and* $\mathscr{B}_w = \{b_1, b_2, \ldots, b_k\}$ *be any arbitrary basis for the subspace* $\mathscr{W}$ *of* $\mathscr{V}$*. For any basis* $\mathscr{B}_v$ *for* $\mathscr{V}$ *we would have:*

$$[\mathrm{T}_{\mathscr{W}}]_{\mathscr{B}_w \mathscr{B}_v} = [\mathrm{T}]_{\mathscr{B}_v} [\mathscr{B}_w]_{\mathscr{B}_v}$$

where the left most and right most matrices are, of course, $n \times k$ *matrices of scalars. And* $[\mathscr{B}_w]_{\mathscr{B}_v}$ *is the matrix with the representation of members of* $\mathscr{B}_w$ *in terms of* $\mathscr{B}_v$ *in its columns.*

Proof. Let $\mathscr{B}_w = \{b'_1, b'_2, \ldots, b'_k\}$. By definition, we would have:

$$[\mathrm{T}_{\mathscr{W}}]_{\mathscr{B}_w \mathscr{B}_v} = \begin{bmatrix} [\mathrm{T}_{\mathscr{W}} b'_1]_{\mathscr{B}_v} & [\mathrm{T}_{\mathscr{W}} b'_2]_{\mathscr{B}_v} & \ldots & [\mathrm{T}_{\mathscr{W}} b'_k]_{\mathscr{B}_v} \end{bmatrix}$$

By definition of $\mathrm{T}_{\mathscr{W}}$ we know that $\mathrm{T}_{\mathscr{W}} u = \mathrm{T}u$ as long as $u \in \mathscr{W}$. As a result:

$$[\mathrm{T}_{\mathscr{W}}]_{\mathscr{B}_w \mathscr{B}_v} = \begin{bmatrix} [\mathrm{T} b'_1]_{\mathscr{B}_v} & [\mathrm{T} b'_2]_{\mathscr{B}_v} & \ldots & [\mathrm{T} b'_k]_{\mathscr{B}_v} \end{bmatrix} = \begin{bmatrix} [\mathrm{T}]_{\mathscr{B}_v} [b'_1]_{\mathscr{B}_v} & [\mathrm{T}]_{\mathscr{B}_v} [b'_2]_{\mathscr{B}_v} & \ldots & [\mathrm{T}]_{\mathscr{B}_v} [b'_k]_{\mathscr{B}_v} \end{bmatrix}$$

which obviously simplifies to:

$$[\mathrm{T}_{\mathscr{W}}]_{\mathscr{B}_w \mathscr{B}_v} = [\mathrm{T}]_{\mathscr{B}_v} \begin{bmatrix} [b'_1]_{\mathscr{B}_v} & [b'_2]_{\mathscr{B}_v} & \ldots & [b'_k]_{\mathscr{B}_v} \end{bmatrix} = [\mathrm{T}]_{\mathscr{B}_v} [\mathscr{B}_w]_{\mathscr{B}_v}$$

□

Corollary 5.2.1. If we extend $\mathscr{B}_w$ to a basis $\mathscr{B}_v$ for $\mathscr{V}$ (refer to proposition 2.2.2) then for any choice of such extension we would have:

$$[\mathrm{T}_{\mathscr{W}}]_{\mathscr{B}_w \mathscr{B}_v} = [\mathrm{T}]_{\mathscr{B}_v} \begin{bmatrix} \mathrm{I}_{k \times k} \\ [0]_{(n-k) \times k} \end{bmatrix}$$

Proof. Since $\mathscr{B}_v$ is an extension of $\mathscr{B}_w$ for all $b'_1, \ldots, b'_k$ we would have:

$$[b'_i]_{\mathscr{B}_v} = \begin{bmatrix} 0 \\ \vdots \\ 1 \\ \vdots \\ 0 \end{bmatrix} (i)$$

and hence:

$$[\mathrm{T}_{\mathscr{W}}]_{\mathscr{B}_w\mathscr{B}_v} = [\mathrm{T}]_{\mathscr{B}_v}[\mathscr{B}_w]_{\mathscr{B}_v} = [\mathrm{T}]_{\mathscr{B}_v} \left[\begin{array}{cccc} 1 & 0 & \ldots & 0 \\ 0 & 1 & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \ldots & 1 \\ \hline 0 & 0 & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \ldots & 0 \end{array}\right] = [\mathrm{T}]_{\mathscr{B}_v} \begin{bmatrix} \mathrm{I}_{k\times k} \\ [0]_{(n-k)\times k} \end{bmatrix}$$

which completes the proof. □

Remark. Since $\mathscr{B}_w \subseteq \mathscr{B}_v$ we, at times, would simply drop $\mathscr{B}_w$ in the subscript of the matrix representation of $\mathrm{T}_{\mathscr{W}}$ and write:

$$[\mathrm{T}_{\mathscr{W}}]_{\mathscr{B}_v} = [\mathrm{T}]_{\mathscr{B}_v} \begin{bmatrix} \mathrm{I}_{k\times k} \\ [0]_{(n-k)\times k} \end{bmatrix}$$

Proposition 5.2.2. *Let $\mathrm{T} : \mathscr{V} \to \mathscr{V}$ be a linear transformation and $\mathscr{W}$ be a subspace of $\mathscr{V}$. The nullity of $\mathrm{T}_{\mathscr{W}}$ would be equal to:*

$$\eta(\mathrm{T}_{\mathscr{W}}) = \dim(\mathscr{V}) - \dim(\mathscr{W}) + \dim(\mathscr{W} \cap \mathrm{Ker}[\mathrm{T}])$$

and:

$$\rho(\mathrm{T}_{\mathscr{W}}) = \dim(\mathscr{W}) - \dim(\mathscr{W} \cap \mathrm{Ker}[\mathrm{T}])$$

Proof. We first notice that by Rank-Nullity theorem for $\mathrm{T}_{\mathscr{W}}$ we can write:

$$\eta(\mathrm{T}_{\mathscr{W}}) = \dim(\mathscr{V}) - \rho(\mathrm{T}_{\mathscr{W}})$$

By definition we can write:

$$\rho(\mathrm{T}_{\mathscr{W}}) = \dim(\mathrm{Im}[\mathrm{T}_{\mathscr{W}}]) = \dim(\mathrm{T}\mathscr{W})$$

and now using the generalized Rank-Nullity theorem (proposition 3.1.7) we know that:

$$\dim(\mathrm{T}\mathscr{W}) = \dim(\mathscr{W}) - \dim(\mathscr{W} \cap \mathrm{Ker}[\mathrm{T}])$$

which results in:

$$\eta(\mathrm{T}_{\mathscr{W}}) = \dim(\mathscr{V}) - \dim(\mathscr{W}) + \dim(\mathscr{W} \cap \mathrm{Ker}[\mathrm{T}])$$

□

Remark. Obviously for some two distinct linear transformations T, S, their respective confinements to some subspace $\mathscr{W}$ might coincide:

$$\mathrm{T}_{\mathscr{W}} = \mathrm{S}_{\mathscr{W}}$$

Definition 5.2.2. Let $\mathrm{T}: \mathscr{V} \to \mathscr{V}$ be a linear transformation and $\mathscr{W}$ be a subspace of $\mathscr{V}$. We say that $\mathscr{W}$ is T**-invariant** if:

$$\mathrm{T}\mathscr{W} \subseteq \mathscr{W}$$

Proposition 5.2.3. *Let $\mathrm{T}: \mathscr{V} \to \mathscr{V}$ be a linear transformation and $\mathscr{W}$ be a subspace of $\mathscr{V}$. Let $\mathscr{B}_w = \{b'_1, \ldots, b'_k\}$ be a basis for $\mathscr{W}$ and $\mathscr{B}_v = \{b_1, \ldots, b_n\}$ be a basis for $\mathscr{V}$ such that $\mathscr{B}_w \subseteq \mathscr{B}_v$ and $b'_i = b_{\pi_i}$ for $i = 1, \ldots, k$. Then by corollary 5.2.1, $\mathscr{W}$ is T-invariant if and only if all the entries in the π_i-the column of $[\mathrm{T}]_{\mathscr{B}_v}$ are zero (for all i) except for possibly the ones on the intersection with the π_j-th row for some j. If we rearrange $\mathscr{B}_v$ such that the first k of them are members of $\mathscr{B}_w$, then $\mathscr{W}$ is T-invariant if and only if $[\mathrm{T}]_{\mathscr{B}_v}$ has gets the following form :*

$$[\mathrm{T}]_{\mathscr{B}_v} = \left[\begin{array}{c|c} \times & \times \\ \hline & \\ 0 & \times \\ & \end{array}\right]$$

Corollary 5.2.2. Let $\mathrm{T}: \mathscr{V} \to \mathscr{V}$ be a linear transformation, and $\mathscr{V}$ have a direct sum decomposition $\mathscr{V} = \bigoplus_{i=1}^{k} \mathscr{W}_i$, such that all $\mathscr{W}_i$s are T-invariant. Then for *any* choice of bases $\mathscr{B}_1, \ldots, \mathscr{B}_k$ for $\mathscr{W}_1, \ldots, \mathscr{W}_k$ the matrix $\mathrm{A} = [\mathrm{T}]_{\mathscr{B}}$, where $\mathscr{B} = \bigcup_{i=1}^{k} \mathscr{B}_i$, is a block diagonal matrix:

$$\mathrm{A} = \begin{bmatrix} \mathrm{A}_{1,1} & 0 & \ldots & 0 \\ 0 & \mathrm{A}_{2,2} & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \ldots & \mathrm{A}_{k,k} \end{bmatrix} = \begin{bmatrix} [T_{\mathscr{W}_1}]_{\mathscr{B}_1} & 0 & \ldots & 0 \\ 0 & [T_{\mathscr{W}_2}]_{\mathscr{B}_2} & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \ldots & [T_{\mathscr{W}_k}]_{\mathscr{B}_k} \end{bmatrix}$$

where $\mathrm{A}_{i,i}$ is a square matrix of size $\dim(\mathscr{W}_i) \times \dim(\mathscr{W}_i)$.

Proposition 5.2.4. *Let $\mathrm{T}: \mathscr{V} \to \mathscr{V}$ be a linear transformation, and $\mathscr{V}$ have a T-invariant direct sum decomposition $\mathscr{V} = \bigoplus_{i=1}^{k} \mathscr{W}_i$. The following would hold:*

$$\rho(\mathrm{T}) = \sum_{i=1}^{k} \rho(\mathrm{T}_{\mathscr{W}_i})$$

$$\eta(\mathrm{T}) = \sum_{i=1}^{k} \eta(\mathrm{T}_{\mathscr{W}_i})$$

as a result, T is an isomorphism (bijective map) if and only if so are all $\mathrm{T}_{\mathscr{W}_i}$s.

Corollary 5.2.3. Let A be a square matrix of scalars having a block diagonalization of the following form:

$$\mathrm{A} = \begin{bmatrix} \mathrm{A}_{1,1} & 0 & \dots & 0 \\ 0 & \mathrm{A}_{2,2} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \mathrm{A}_{k,k} \end{bmatrix}$$

the following would hold:

$$\rho(\mathrm{A}) = \sum_{i=1}^{k} \rho(\mathrm{A_{i,i}})$$

$$\eta(\mathrm{A}) = \sum_{i=1}^{k} \eta(\mathrm{A_{i,i}})$$

as a result, A is nonsingular if and only if so are all $\mathrm{A}_{i,i}$s.

Proposition 5.2.5. *Let* $\mathrm{T}: \mathscr{V} \to \mathscr{V}$ *be a linear transformation, and* $\mathscr{V}$ *have a* T*-invariant direct sum decomposition* $\mathscr{V} = \bigoplus_{i=1}^{k} \mathscr{W}_i$*. There exist* k *unique linear transformations* $\mathrm{T}^{(1)}, \dots, \mathrm{T}^{(k)}$ *over* $\mathscr{V}$ *such that:*

$$\mathrm{T}^{(i)}_{\mathscr{W}_i} = \mathrm{T}_{\mathscr{W}_i}$$

and:

$$\mathrm{T}^{(i)}_{\mathscr{W}_j} = \mathrm{I}$$

As a result we would have:

$$\mathrm{T} = \mathrm{T}^{(\pi_1)} \circ \mathrm{T}^{(\pi_2)} \circ \dots \circ \mathrm{T}^{(\pi_k)}$$

for any permutation $\pi = \langle \pi_1, \dots, \pi_k \rangle$ *of* $1, \dots, k$*.*

Proof. We fix some k bases $\mathscr{B}_1, \dots, \mathscr{B}_k$ for $\mathscr{W}_1, \dots, \mathscr{W}_k$, and for any $b \in \mathscr{B} = \cup_{i=1}^{k} \mathscr{B}_i$ we define $\mathrm{T}^{(i)}b$ in the following fashion:

$$\mathrm{T}^{(i)}b = \begin{cases} \mathrm{T}b & \text{if } b \in \mathscr{B}_i \\ b & \text{if } b \notin \mathscr{B}_i \end{cases}$$

The fact that such definition of the $\mathrm{T}^{(i)}$ satisfies the required behavior regardless of the choice of $\mathscr{B}_i$ follows from the fact that $\mathscr{W}_i$ are T-invariant subspaces of $\mathscr{V}$. □

Corollary 5.2.4. Let A be a square matrix of scalars having a block diagonalization of the following form:

$$\mathrm{A} = \begin{bmatrix} \mathrm{A}_{1,1} & 0 & \dots & 0 \\ 0 & \mathrm{A}_{2,2} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \mathrm{A}_{k,k} \end{bmatrix}$$

Then defining $n \times n$ matrices $\mathrm{A}^{(i)}$ in the following fashion:

$$\mathrm{A}^{(i)} = \left[\begin{array}{c|c|c} 1 & & \\ \ddots & & \\ & 1 & \\ \hline & \mathrm{A}_{i,i} & \\ \hline & & \begin{array}{ccccc} 1 & & & & \\ & 1 & & & \\ & & \ddots & & \\ & & & 1 & \\ & & & & 1 \end{array} \end{array}\right]$$

Then we would have:

$$\mathrm{A} = \mathrm{A}^{(\pi_1)} \times \mathrm{A}^{(\pi_2)} \times \ldots \times \mathrm{A}^{(\pi_k)}$$

for any permutation $\pi = \langle \pi_1, \ldots, \pi_k \rangle$ of $1, \ldots, k$.

I will pick up this issue later in chapter 8 which is devoted to the problem of finding good *descriptions* that maximally (and hopefully) uniquely *simplify* the behavior of a linear transformation; where by *description* we technically mean invariant direct sum decompositions, and by *simplicity* we refer to our ever lasting love of zeros and block diagonal (or ideally just diagonal) matrices.

5.3 LU and Cholesky decompositions

Chapter 6

Bilinar Forms, Orthogonality and Inner Product Spaces

Bilinear forms can be regarded as generalization of the notion of inner products. Obviously inner products are complex concepts, and finding nice ones over any arbitrary vector space is not always possible.

Definition 6.0.1. Let $\mathscr{V}$ be a vector space defined over the field $\mathbb{F}$. We call a function $f : \mathscr{V} \times \mathscr{V} \to \mathbb{F}$ a **bilinear form** if it is a linear functional in both its arguments, *individually*. In other words fixing any vector $u \in \mathscr{V}$, both $f(u, \cdot)$ and $f(\cdot, u)$ are linear functionals over $\mathscr{V}$.

We know that for any $u \in \mathscr{V}$ the linear functional $f(u, \cdot)$ belongs to the dual space $\mathscr{V}^*$ and hence would have the transpose of its row vector representation in terms of a basis $\mathscr{B}$, in terms of $\mathscr{B}^*$:

$$[f(u, \cdot)]_{\mathscr{B}^*} = [f(u, \cdot)]^\intercal_{\mathscr{B}}$$

We also know that for any fixed u we would have:

$$f(u, v) = [f(u, \cdot)]_{\mathscr{B}}[v]_{\mathscr{B}}$$

and hence:

$$f(u, v) = [f(u, \cdot)]^\intercal_{\mathscr{B}^*}[v]_{\mathscr{B}}$$

Now let's look at the mapping that takes u and produces the linear functional $f(u, \cdot)$. Let's name this mapping $\mathrm{F} : \mathscr{V} \to \mathscr{V}^*$:

$$\mathrm{F}u = f(u, \cdot) \in \mathscr{V}^*$$

From the bilinearity of f it follows that F is a linear transformation, and hence has a matrix representation for any basis $\mathscr{B}$:

$$[f(u, \cdot)]_{\mathscr{B}^*} = [\mathrm{F}u]_{\mathscr{B}^*} =$$

and thus for any u and v we would have:

$$f(u, v) = ([\mathrm{F}]_{\mathscr{B}\mathscr{B}^*}[u]_{\mathscr{B}})^\intercal[v]_{\mathscr{B}} = [u]^\intercal_{\mathscr{B}}[\mathrm{F}]^\intercal_{\mathscr{B}\mathscr{B}^*}[v]_{\mathscr{B}}$$

Thus by fixing a basis $\mathscr{B}$, the effect of our bilinear function f would be equal to:

$$[u]^\intercal_{\mathscr{B}}\mathrm{A}[v]_{\mathscr{B}}$$

for some square scalar matrix A. This relationship as all the other *matrix representation*s we have seen until now is a bijective homomorphism, and hence we have proved that the space of all bilinear forms $\mathscr{L}(\mathscr{V} \times \mathscr{V}, \mathbb{F})$ is isomorphic to the space of $n \times n$ scalar matrices. Let's investigate more the content of the matrix A. Let $b_1, b_2, \ldots, b_n$ be the members of the basis $\mathscr{B}$:

$$\begin{aligned}
\mathrm{A} &= [\mathrm{F}]^{\intercal}_{\mathscr{B}\mathscr{B}^*} \\
&= \begin{bmatrix}[\mathrm{F}b_1]_{\mathscr{B}^*} & [\mathrm{F}b_2]_{\mathscr{B}^*} & \ldots & [\mathrm{F}b_n]_{\mathscr{B}^*}\end{bmatrix}^{\intercal} = \begin{bmatrix}[\mathrm{F}b_1]^{\intercal}_{\mathscr{B}^*} \\ [\mathrm{F}b_2]^{\intercal}_{\mathscr{B}^*} \\ \vdots \\ [\mathrm{F}b_n]^{\intercal}_{\mathscr{B}^*}\end{bmatrix} = \begin{bmatrix}[\mathrm{F}b_1]_{\mathscr{B}} \\ [\mathrm{F}b_2]_{\mathscr{B}} \\ \vdots \\ [\mathrm{F}b_n]_{\mathscr{B}}\end{bmatrix} \\
&= \begin{bmatrix}[f(b_1,\cdot)]_{\mathscr{B}} \\ [f(b_2,\cdot)]_{\mathscr{B}} \\ \vdots \\ [f(b_n,\cdot)]_{\mathscr{B}}\end{bmatrix} = \begin{bmatrix}f(b_1,b_1) & f(b_1,b_2) & \ldots & f(b_1,b_n) \\ f(b_2,b_1) & f(b_2,b_2) & \ldots & f(b_2,b_n) \\ \vdots & \vdots & \ddots & \vdots \\ f(b_n,b_1) & f(b_n,b_2) & \ldots & f(b_n,b_n)\end{bmatrix}
\end{aligned}$$

Definition 6.0.2. Let $f : \mathscr{V} \times \mathscr{V} \to \mathbb{F}$ be a bilinear form over the n-dimensional vector space $\mathscr{V}$. For any basis $\mathscr{B}$ for $\mathscr{V}$ we define the following $n \times n$ matrix of scalars to be the matrix representation of f in terms of $\mathscr{B}$:

$$[f]_{\mathscr{B}} = [f(b_i, b_j)]_{n\times n}$$

Proposition 6.0.1. *Let $f : \mathscr{V} \times \mathscr{V} \to \mathbb{F}$ be a bilinear form. For any basis $\mathscr{B}$ for $\mathscr{V}$, for every $u, v \in \mathscr{V}$ we have:*

$$f(u, v) = [u]^{\intercal}_{\mathscr{B}}[f]_{\mathscr{B}}[v]_{\mathscr{B}}$$

Also the two functionals induced by f and a vector u satisfy the following:

$$[f(u,\cdot)]_{\mathscr{B}} = [u]^{\intercal}_{\mathscr{B}}[f]_{\mathscr{B}} = \Big[[f]^{\intercal}_{\mathscr{B}}[u]_{\mathscr{B}}\Big]^{\intercal}$$

$$[f(\cdot,u)]_{\mathscr{B}} = [u]^{\intercal}_{\mathscr{B}}[f]^{\intercal}_{\mathscr{B}} = \Big[[f]_{\mathscr{B}}[u]_{\mathscr{B}}\Big]^{\intercal}$$

Proposition 6.0.2. *Let $f : \mathscr{V} \times \mathscr{V} \to \mathbb{F}$ be a bilinear form and $\mathscr{B}$ be any arbitrary basis for $\mathscr{V}$. f is a symmetric bilinear form (i.e. for all u, v, $f(u, v) = f(v, u)$ holds), if and only if the matrix $[f]_{\mathscr{B}}$ is symmetric.*

Proof. If $\mathrm{A} = [f]_{\mathscr{B}}$ is symmetric, the symmetricity of f immediately follows. Now assume f is symmetric:

$$\forall u, v \in \mathscr{V} : \ [u]^{\intercal}\mathrm{A}[v] = [v]^{\intercal}\mathrm{A}[u]$$

If we pick u and v to be the vectors corresponding to the representations of the form $[0, \ldots, 0, 1, 0, \ldots, 0]^{\intercal}$ with the 1 in u's case be in its i-th position and in v's case in its j-th position, in that case $[u]^{\intercal}\mathrm{A}[v]$ would select the i, j entry of A and $[v]^{\intercal}\mathrm{A}[u]$ would select the j, i entry of A, and the equality holds for all i, j. □

Proposition 6.0.3. *Let $f : \mathscr{V} \times \mathscr{V} \to \mathbb{F}$ be a bilinear form over the vector space $\mathscr{V}$ defined over the totally ordered field $\mathbb{F}$, and $\mathscr{B}$ be any arbitrary basis for $\mathscr{V}$. The following are equivalent:*

(i) *For any non-zero vector $u \in \mathscr{V}$ we have $f(u, u) > 0$*

(ii) *The matrix* $[f]_{\mathscr{B}}$ *is a positive definite matrix.*

and hence we consistently thus define the property of positive definiteness for bilinear forms.

Proposition 6.0.4 (Change of Basis)**.** *Let* $\mathscr{V}$ *be a vector space defined over* $\mathbb{F}$ *and* $f : \mathscr{V} \times \mathscr{V} \to \mathbb{F}$ *be a bilinear form over* $\mathscr{V}$*. Let* $\mathscr{B}$ *and* $\mathscr{B}'$ *be two bases for* $\mathscr{V}$*. Defining the basis change matrix* $\mathrm{P}_{\mathscr{B}' \to \mathscr{B}}$ *the same as in 2.3.2 to be:*

$$\mathrm{P}_{(\mathscr{B}' \to \mathscr{B})} = [\mathscr{B}']_{\mathscr{B}}$$

we will have:

$$[f]_{\mathscr{B}'} = \mathrm{P}^{\intercal}[f]_{\mathscr{B}}\mathrm{P}$$

Proof.

$$[f]_{\mathscr{B}'} = \Big[f(b'_i, b'_j)\Big]_{n\times n} = \Big[[b'_i]^{\intercal}_{\mathscr{B}}[f]_{\mathscr{B}}[b'_j]_{\mathscr{B}}\Big]_{n\times n} = [\mathscr{B}']^{\intercal}_{\mathscr{B}}[f]_{\mathscr{B}}[\mathscr{B}']_{\mathscr{B}} = \mathrm{P}^{\intercal}[f]_{\mathscr{B}}\mathrm{P}$$

□

Remark. Notice the difference between this formula and that derived in 3.2.4. In the case of linear transformations by defining the same basis change matrix $\mathrm{P} = [\mathscr{B}']_{\mathscr{B}}$ we had:

$$[\mathrm{T}]_{\mathscr{B}'} = \mathrm{P}[\mathrm{T}]_{\mathscr{B}}\mathrm{P}^{-1}$$

but for the case of bilinear forms we have:

$$[f]_{\mathscr{B}'} = \mathrm{P}^{\intercal}[f]_{\mathscr{B}}\mathrm{P}$$

We *can not* infer that all matrix representations $[f]_{\mathscr{B}}$ of f are similar matrices.

Remark. As we have seen by now, the vector space of bilinear forms $\mathscr{L}(\mathscr{V} \times \mathscr{V}, \mathbb{F})$ is isomorphic to $\mathbb{F}^{n\times n}$ by fixing a basis for $\mathscr{V}$, whereas the space of all linear transformations on $\mathscr{V}$, namely $\mathscr{L}(\mathscr{V}, \mathscr{V})$ is as well, by a choice of basis, isomorphic to the same space $\mathbb{F}^{n\times n}$. The two isomorphisms are consistent, as can be seen through the change of bases (similarity) relationships each linear transformation and each bilinear form induces on $\mathbb{F}^{n\times n}$, and hence each similarity class in $\mathbb{F}^{n\times n}$, by fixing a basis, corresponds to one and only one linear transformation and one and only one bilinear form.

The relationship we just mentioned can be best seen in light of duality. As we mentioned earlier any bilinear form f induces two linear transformations one mapping a vector u to $f(\cdot, u)$ and the other to $f(u, \cdot)$. We have seen that the behavior of a bilinear form can be imitated by a matrix operation of the following form:

$$f(u, v) = [u]^{\intercal}_{\mathscr{B}}[f]_{\mathscr{B}}[v]_{\mathscr{B}}$$

This would be an exotic magical form fallen merely out of nowhere, unless we try to interpret this operation in terms of linear transformations operating on row and column vectors (i.e. members of $\mathscr{V}$ and $\mathscr{V}^*$):

$$f(u, v) = \begin{bmatrix}\alpha_1 & \alpha_2 & \dots \alpha_n\end{bmatrix} \begin{bmatrix} f(b_1, b_1) & f(b_1, b_2) & \dots & f(b_1, b_n) \\ f(b_2, b_1) & f(b_2, b_2) & \dots & f(b_2, b_n) \\ \vdots & \vdots & \ddots & \vdots \\ f(b_n, b_1) & f(b_n, b_2) & \dots & f(b_n, b_n) \end{bmatrix} \begin{bmatrix}\beta_1 \\ \beta_2 \\ \vdots \\ \beta_n\end{bmatrix}$$

where $[\alpha_1, \alpha_2, \ldots, \alpha_n]^\intercal$ and $[\beta_1, \beta_2, \ldots, \beta_n]^\intercal$ are representations of u, v in terms of $\mathscr{B}$. We decompose the complex behavior of $f(\cdot,\cdot)$ to the two simple behaviors of $f(u,\cdot)$ and $f(\cdot,u)$. Let's first look at $f(\cdot,u)$, we will call this functional $f_{\cdot,u}(\cdot)$ from now on:

$$f_{\cdot,u}(v) = \begin{bmatrix}\beta_1 & \beta_2 & \ldots \beta_n\end{bmatrix} \underbrace{\begin{bmatrix} f(b_1,b_1) & f(b_1,b_2) & \ldots & f(b_1,b_n) \\ f(b_2,b_1) & f(b_2,b_2) & \ldots & f(b_2,b_n) \\ \vdots & \vdots & \ddots & \vdots \\ f(b_n,b_1) & f(b_n,b_2) & \ldots & f(b_n,b_n) \end{bmatrix}}_{[f_{\cdot,u}]_{\mathscr{B}}} \begin{bmatrix}\alpha_1 \\ \alpha_2 \\ \vdots \\ \alpha_n\end{bmatrix}$$

As we can see $f_{\cdot,u}$ can is a linear functional over $\mathscr{V}^*$, i.e. a vector of $\mathscr{V}$. We know use the term $\mathrm{F}_{\cdot,u}$ to denote the linear map:

$$u \xrightarrow{\mathrm{F}_{\cdot,u}} f_{\cdot,u}(\cdot)$$

and hence:

$$\mathrm{F}_{\cdot,u} : \mathscr{V} \to \mathscr{V}$$

Now let's look at the other side of the story:

$$f_{u,\cdot}(v) = \underbrace{\begin{bmatrix}\alpha_1 & \alpha_2 & \ldots \alpha_n\end{bmatrix} \begin{bmatrix} f(b_1,b_1) & f(b_1,b_2) & \ldots & f(b_1,b_n) \\ f(b_2,b_1) & f(b_2,b_2) & \ldots & f(b_2,b_n) \\ \vdots & \vdots & \ddots & \vdots \\ f(b_n,b_1) & f(b_n,b_2) & \ldots & f(b_n,b_n) \end{bmatrix}}_{[f_{u,\cdot}]_{\mathscr{B}}} \begin{bmatrix}\beta_1 \\ \beta_2 \\ \vdots \\ \beta_n\end{bmatrix}$$

As we can see both u and $f_{u,\cdot}$ are linear functionals over $\mathscr{V}$, i.e. vectors of the dual space $\mathscr{V}^*$. We know use the term $\mathrm{F}_{\cdot,u}$ to denote the linear map:

$$u \xrightarrow{\mathrm{F}_{u,\cdot}} f_{u,\cdot}(\cdot)$$

and we could see that:

$$\mathrm{F}_{u,\cdot} : \mathscr{V}^* \to \mathscr{V}^*$$

It is necessary to notice that the element u in $f(u,\cdot)$ is of a different nature than the one in $f(\cdot,u)$. The latter is a member of $\mathscr{V}$ (and hence a linear functional over $\mathscr{V}^*$), and the former is a member of $\mathscr{V}^*$ (a linear functional over $\mathscr{V}$).
What is the relationship between $\mathrm{F}_{u,\cdot}$ and $\mathrm{F}_{\cdot,u}$? We claim that they are the transpose of each other. It can easily be seen from their corresponding matrix representations. But let's examine this claim in abstract algebraic terms.

Proposition 6.0.5. *Let $\mathscr{V}$ be a vector space defined over $\mathbb{F}$ and $f : \mathscr{V} \times \mathscr{V} \to \mathbb{F}$ be a bilinear form over $\mathscr{V}$. Define the linear transformations $\mathrm{F}_{\cdot,u} : \mathscr{V} \to \mathscr{V}$ and $\mathrm{F}_{u,\cdot} : \mathscr{V}^* \to \mathscr{V}^*$ to be the following:*

$$\mathrm{F}_{\cdot,u} u = f(\cdot, u)$$

$$\mathrm{F}_{u,\cdot} u = f(u, \cdot)$$

The following would hold:

(i) *For any basis $\mathscr{B}$ for $\mathscr{V}$ we have:*

$$[\mathrm{F}_{\cdot,u}]_{\mathscr{B}} = [f]_{\mathscr{B}}$$

(ii) $\mathrm{F}_{u,\cdot}$ *is the transformation transpose of* $\mathrm{F}_{\cdot,u}$*:*

$$\mathrm{F}_{u,\cdot} = \mathrm{F}^{\mathrm{t}}_{\cdot,u}$$

and hence:

$$[\mathrm{F}_{u,\cdot}]_{\mathscr{B}^*} = [\mathrm{F}_{\cdot,u}]^{\intercal}_{\mathscr{B}} = [f]^{\intercal}_{\mathscr{B}}$$

Proof. (i) is trivial from the arguments leading to the proposition. We here prove the second part. (ii) We recall that the transpose of a transformation $\mathrm{T}: \mathscr{V} \to \mathscr{V}$ is a linear transformation $\mathrm{T}^{\mathrm{t}}: \mathscr{V}^* \to \mathscr{V}^*$ such that for every linear functional $g \in \mathscr{V}^*$, $\mathrm{T}^{\mathrm{t}} g$ is the linear functional (again on $\mathscr{V}$) that is defined by $g \circ \mathrm{T}$. In our case let $g \in \mathscr{V}^*$ be some linear functional (a row vector). $\mathrm{F}_{u,\cdot} g$ is the functional defined by $f_{g,\cdot}$, which by definition we know satisfies:

$$[\mathrm{F}_{u,\cdot} g]_{\mathscr{B}} = [g]_{\mathscr{B}} [f]_{\mathscr{B}}$$

and as a result:

$$\mathrm{F}_{u,\cdot} g = g \circ \mathrm{F}_{\cdot,u}$$

and the proof is complete. □

Now that we have established the canonical homomorphism between bilinear forms and linear transformations over $\mathscr{V}$:

$$\mathscr{L}(\mathscr{V} \times \mathscr{V}, \mathbb{F}) \xrightarrow{\varphi} \mathscr{L}(\mathscr{V}, \mathscr{V})$$

that behaves in the following fashion:

$$f \mapsto \mathrm{F}_{\cdot,u}$$

we could safely define the notions of rank and nullity of a bilinear form:

Definition 6.0.3. Let $\mathscr{V}$ be a vector space defined over $\mathbb{F}$ and $f: \mathscr{V} \times \mathscr{V} \to \mathbb{F}$ be a bilinear form over $\mathscr{V}$. We define the rank and nullity of f to be the following:

$$\rho(f) = \rho(\mathrm{F}_{\cdot,\mathrm{u}}) = \rho(\mathrm{A})$$

$$\eta(f) = \eta(\mathrm{F}_{\cdot,\mathrm{u}}) = \eta(\mathrm{A})$$

where A could be any of the possible matrix representations of f, i.e. $\mathrm{A} = [f]_{\mathscr{B}}$ for some bass $\mathscr{B}$.

We could see that the neat behavior we would expect from a bilinear form to be consistent with our intuitions of orthogonality requires that the matrix representation of f be an SPD matrix (although this is *not* a sufficient condition for a bilinear form to be an inner product). This would be no surprise, since we know that SPD is conserved through the quasi-similarity relationship of matrices induced by $\mathrm{B} = \mathrm{P}^{\intercal}\mathrm{AP}$. Other than the property of symmetricity, however, we would expect from our inner product intuitions that the kernel of the linear functional $f(\cdot, u)$ be a hyperplane in $\mathscr{V}$. But we know from proposition 4.0.10 that a linear functional does have the option of being the zero functional, and hence its kernel swallowing the whole of $\mathscr{V}$. However, $f(u, \cdot)$ could be the zero functional while $f(\cdot, u)$ is not! For the sake of this study I have focused on symmetric bilinear forms, an assumption which sets us free of such problems. We will pick up the treatise of bilinear forms just as we leave them here in the next part.

6.1 Symmetric (degenerate) bilinear forms

Confining ourselves to the case of symmetric bilinear forms, we can define the notions of perpendicularity and projection, although they do not satisfy all the neat properties we would expect:

Definition 6.1.1. For any two vectors $u, v \in \mathscr{V}$, and any two sets of vectors (possibly subspaces) $C_1, C_2 \subseteq \mathscr{V}$ we define the following:

(i) $u \perp v$ *with respect to* f, if $f(u, v) = 0$.

(ii) $C_1 \perp C_2$ *with respect to* f, if for any $u \in C_1$ and any $v \in C_2$ we have $u \perp_f v$

Remark. The perpendicularity relationship in this general form that we have defined is *not* transitive! First notice that u can be a vector such that the kernel of the linear functional $f(u, \cdot)$ is the whole space $\mathscr{V}$ (zero functional). In this case *any* two vectors v, w would satisfy $w \perp u$ and $u \perp v$ w.r.t f, from which it does not follow that $v \perp w$ w.r.t f. Even if $f(u, \cdot)$ is not the zero functional, and hence has a kernel which is a hyperplane in $\mathscr{V}$, again *any* two vectors v, w in that hyperplane would satisfy $w \perp u$ and $u \perp v$, from which it does not again follow that $v \perp w$.

Remark. According to the derivation of the matrix representation of a bilinear form f which as we saw has the form $[f(b_i, b_j)]_{n\times n}$, we could notice that if $\mathscr{B}$ is an f-orthogonal basis for $\mathscr{V}$ (a basis all the members of which are mutually orthogonal w.r.t f) then $[f]_{\mathscr{B}}$ would be diagonal.

Definition 6.1.2. Let $\mathscr{V}$ be a vector space defined over a subfield $\mathbb{F}$ of $\mathbb{C}$ and $f : \mathscr{V} \times \mathscr{V} \to \mathbb{F}$ be a symmetric bilinear form over $\mathscr{V}$. We define the **quadratic form** corresponding to f to be the linear functional $q \in \mathscr{V}^*$ defined as:

$$q(u) = f(u, u)$$

Proposition 6.1.1. *Let $\mathscr{V}$ be a vector space defined over a subfield $\mathbb{F}$ of $\mathbb{C}$ and $f : \mathscr{V} \times \mathscr{V} \to \mathbb{F}$ be a symmetric bilinear form over $\mathscr{V}$. The behavior of f is uniquely identified by the behavior of its quadratic form.*

Proof. For any two arbitrary vectors u, v we have:

$$\begin{aligned} f(u,v) &= f(\frac{u+v}{2} + \frac{u-v}{2}, \frac{u+v}{2} - \frac{u-v}{2}) \\ &= f(\frac{u+v}{2}, \frac{u+v}{2}) - f(\frac{u-v}{2}, \frac{u-v}{2}) - f(\frac{u+v}{2}, \frac{u-v}{2}) + f(\frac{u-v}{2}, \frac{u+v}{2}) \\ &= \frac{1}{4}\big[q(u+v) - q(u-v)\big] \end{aligned}$$

□

To avoid all the confusions I faced while reading bilinear forms, I here introduce all the needed notions and then prove their properties altogether; instead of spending time on proving properties that seem trivial as long as a new notion is not introduced. Throughout the rest of this section since we are going to make extensive use of the assumption of *existence of some total ordering* over $\mathbb{F}$, we confine ourselves to the case where $\mathbb{F}$ is either of $\mathbb{R}$ or $\mathbb{C}$, or more generally, and in less verbose and more relaxed terms, we assume $\mathbb{F}$ is a subfield of $\mathbb{C}$.

Definition 6.1.3. Let $\mathscr{W}$ be a subspace of $\mathscr{V}$ over which a symmetric bilinear form f is defined. We define the **perpendicular subspace** of $\mathscr{W}$ **w.r.t** f to be the set of all vectors u such that $u \perp \mathscr{W}$, and denote it by $\mathscr{W}^{\perp}$.

Definition 6.1.4. Let $\mathscr{V}$ be a vector space defined over $\mathbb{F}$ and $f : \mathscr{V} \times \mathscr{V} \to \mathbb{F}$ be a symmetric bilinear form over $\mathscr{V}$. We call the vector $u \in \mathscr{V}$ a **degeneracy** of f if the linear functional $f(\cdot, u)$ (or equivalently $f(u, \cdot)$) is the zero functional. In other words a vector $u \in \mathscr{V}$ is a degeneracy if $u \perp \mathscr{V}$. The set of all degeneracies of f can easily be checked to be a special case of $\mathscr{W}^{\perp}$ and hence a subspace of $\mathscr{V}$. We refer to this subspace by the **radical** of f on $\mathscr{V}$ and denote it by $\mathscr{V}^{\perp}$. We call a symmetric bilinear form f **degenerate** if its radical is non-trivial, and non-degenerate otherwise.

Corollary 6.1.1. If f is a degenerate bilinear form over the vector space $\mathscr{V}$, there exists some vector $u \in \mathscr{V}$ such that:

$$\forall v \in \mathscr{V} : \ f(u, v) = f(v, u) = 0$$

Definition 6.1.5. Let $\mathscr{V}$ be a vector space defined over $\mathbb{F}$ and $f : \mathscr{V} \times \mathscr{V} \to \mathbb{F}$ be a symmetric bilinear form over $\mathscr{V}$. We call the vector $u \in \mathscr{V}$ **isotropic w.r.t** f if $u \perp u$ (i.e. $f(u, u) = 0$). And we call the *set* of all f-isotropic vectors $\mathscr{V}_{(0)}$. We will see later that $\mathscr{V}_{(0)}$ is not a subspace.

Remark. It might seem quite counter-intuitive that $\mathscr{V}_{(0)}$ and $\mathscr{V}^{\perp}$ do not coincide. We will investigate this more, but as a quick example in $\mathscr{V} = \mathbb{R}^3$. Take for instance the matrix:

$$\mathrm{A} = \begin{bmatrix} 1 & 2 & 3 \\ 2 & 1 & 0 \\ 3 & 0 & 5 \end{bmatrix}$$

and the following vector in $\mathbb{R}^3$:

$$u_{\circ} = \begin{bmatrix} -1 \\ 0 \\ +1 \end{bmatrix}$$

It can easily be checked that although u is not a degeneracy of $f(u, v) = u^{\intercal}\mathrm{A}v$, but it is isotropic w.r.t f: $f(u, u) = 0$.

Definition 6.1.6. Special care needs to be given to f-isotropic non-degeneracies, and hence we introduce a name for these kinds of vectors. We call a vector $u \in \mathscr{V}$ an f-**weird** vector if $u \in \mathscr{V}_{(0)} - \mathscr{V}^{\perp}$, i.e. it is isotropic w.r.t f but not a degeneracy.

Proposition 6.1.2. *Let $\mathscr{V}$ be a vector space and f be a bilinear form over $\mathscr{V}$. If $\mathscr{V}_{(0)} = \mathscr{V}$ then for any u, v we will have:*

$$f(u, v) = -f(v, v)$$

In other words, the only case where $\mathscr{V}_{(0)}$ swallows up $\mathscr{V}$ is when f is skew symmetric.

Proof. Fix a basis $\mathscr{B}$ for $\mathscr{V}$ and denote the matrix $[f]_{\mathscr{B}}$ by $\mathrm{A} = [a_{i,j}]_{n \times n}$. We know that if the representation of u in terms of $\mathscr{B}$ is $[\alpha_1, \alpha_2, \ldots, \alpha_n]^{\intercal}$ we will have:

$$f(u, u) = [u]^{\intercal}\mathrm{A}[u]_{\mathscr{B}} = \sum_{i=1}^{n} a_{i,i}\alpha_i + \sum_{i,j} a_{i,j}\alpha_i\alpha_j$$

Since the sum above is zero for *any* $u \in \mathscr{V}$ it would immediately follow that:

$$a_{i,i} = 0$$

and that:

$$a_{i,j} = -a_{j,i}$$

and hence:

$$\mathrm{A}^\intercal = \mathrm{A}$$

which completes the proof. □

Corollary 6.1.2. If the bilinear form f is symmetric and non-zero, it follows from what we proved that $\mathscr{V}_{(0)} \subset \mathscr{V}$ and hence there exist vectors such as u such that $f(u,u) \neq 0$.

Definition 6.1.7. For any f-non-weird vector v one can define a well-defined "**projection on** v " operation over $\mathscr{V}$ (i.e. even if v is a degeneracy of f): Let v be a vector, not weird with respect to the symmetric bilinear form f. For any $u \in \mathscr{V}$ we define the projection of u on v *with respect to* f in the following fashion:

$$\mathrm{proj}_v(u) = \frac{f(u,v)}{f(v,v)}v$$

If v is a degeneracy both the numerator and the denominator of the fraction are zero, and thus we safely define $\mathrm{proj}_v(\cdot) = 0$ when v is a degeneracy.

Proposition 6.1.3. *Let f be a symmetric bilinear form defined over the vector space $\mathscr{V}$, and $v \in \mathscr{V}$ be not weird w.r.t f ($f(v,v) = 0$ only if v is a degeneracy). Then for any $u \in \mathscr{V}$ the projection of u on v w.r.t f satisfies:*

$$v \perp u - \mathrm{proj}_v(u)$$

As a result the following operator is the "projector" operator on the subspace $v^\perp$:

$$\mathrm{T} = \mathrm{I} - \frac{f(\cdot, v)}{f(v,v)}v$$

Proof. If v is a degeneracy the property obviously holds since $\mathrm{proj}_v(\cdot) = 0$. Otherwise, we notice that we have defined the projection w.r.t f to be:

$$\mathrm{proj}_v(u) = \frac{f(u,v)}{f(v,v)}v$$

Plugging this in the perpendicularity definition we would get:

$$f(v, u - \mathrm{proj}_v(u)) = f(v,u) - f\left(v, \frac{f(u,v)}{f(v,v)}v\right) = f(v,u) - \frac{f(u,v)}{f(v,v)}f(v,v) = 0$$

The fact that $\mathrm{T}u$ for any u lies in $v^\perp$ is obvious. The only other thing to prove is that $\mathrm{T}^2 = \mathrm{T}$, which is again obvious, since the projection of a vector in $v^\perp$ on v is zero and thus T leaves such vectors untouched. □

Proposition 6.1.4. *Let $\mathscr{V}$ be a vector space defined over $\mathbb{F}$ and $f : \mathscr{V} \times \mathscr{V} \to \mathbb{F}$ be a symmetric bilinear form over $\mathscr{V}$. Then $\mathscr{V}^\perp$, the radical of f on $\mathscr{V}$, is a subspace of $\mathscr{V}$ and* $\dim(\mathscr{V}^\perp) = \eta(f)$.

Proof. We defined the nullity of f to be $\eta(\mathrm{F}_{\cdot,\mathrm{u}})$. It can easily be seen that $\mathscr{V}^\perp$ is in fact $\mathrm{Ker}[\mathrm{F}_{\cdot,\mathrm{u}}]$. Take any arbitrary vector $u_\circ \in \mathscr{V}$. u is in the radical if and only if $f(\cdot, u_\circ)$ is the zero functional, which is if and only if $\mathrm{F}_{\cdot,u}u_\circ$ is zero, which is if and only if $u_\circ \in \mathrm{Ker}[\mathrm{F}_{\cdot,\mathrm{u}}]$. □

We know that for any vector $u \in \mathscr{V}$ the set of all vectors in $u^\perp$ w.r.t f are in fact the Kernel of a linear functional and by proposition 4.0.10 has nullity equal to either $\dim(\mathscr{V})$ or $\dim(\mathscr{V}) - 1$. In the former case u would be a degeneracy of f, and in the latter case $u^\perp$ is a hyperplane in $\mathscr{V}$ (that might contain u or not, we will discuss this situation more later). Now for the case of a subspace $\mathscr{W}$, the way we have defined $\mathscr{W}^\perp$ w.r.t f is the following:

$$\mathscr{W}^\perp = \bigcap_{v \in \mathscr{W}} v^\perp$$

From this we could easily see that $\mathscr{W}^\perp$ is a subspace of $\mathscr{V}$. However, this definition is not quite functional. It would be reasonable to find a way to perform the intersection on a finite number of $v^\perp$s and build $\mathscr{W}^\perp$ from there. The first simplification is that from all the scalar multiples of a vector v only one of them needs to be counted in the intersection above. We here make the observation that if we fix any arbitrary basis for $\mathscr{W}$, we could find $\mathscr{W}^\perp$ by taking the intersection of their respective $b^\perp$s. Let $\mathscr{B}_w = \{b'_1, b'_2, \ldots, b'_k\}$ be an arbitrary basis for $\mathscr{W}$. It would be easy to see that $u \perp \mathscr{W}$ if and only if $u \perp b'_i$ for all i: If $u \in \mathscr{W}^\perp$ it is obviously perpendicular w.r.t f to all b'_is. Also if $u \perp b'_i$ for all i, for any $v \in \mathscr{W}$ with representation in terms of $\mathscr{B}_w$ being $[\alpha_1, \alpha_2, \ldots, \alpha_k]^\intercal$ we would have:

$$f(u, v) = \sum_{i=1}^{k} \alpha_i f(u, b'_i) = 0 \Rightarrow u \perp v$$

which immediately results in:

$$\mathscr{W}^\perp = \bigcap_{i=1}^{k} {b'_i}^\perp$$

Proposition 6.1.5. *Let $\mathscr{W}$ be a subspace of $\mathscr{V}$ over which a symmetric bilinear form f is defined. The dimensionality of $\mathscr{W}^\perp$ defined w.r.t f would be:*

$$\dim(\mathscr{W}^\perp) = \dim(\mathscr{V}) - \dim(\mathscr{W}) + \dim(\mathscr{W} \cap \mathscr{V}^\perp)$$

where $\mathscr{V}^\perp$ is the radical of f.

Proof. For any basis $\mathscr{B}_v$ for $\mathscr{V}$ we could translate the problem of finding the dimensionality of $\mathscr{W}^\perp$ to the problem of finding the dimensionality of $[\mathscr{W}^\perp]_{\mathscr{B}_v}$. This would be, according to (ii), equivalent to finding the intersection of the following level sets for all i:

$$[v]^\intercal_{\mathscr{B}_v} [f]_{\mathscr{B}_v} [b'_i]_{\mathscr{B}_v} = 0$$

which is equivalent to solving the dimensionality of the following level set:

$$[v]^\intercal_{\mathscr{B}_v} \begin{bmatrix} [f]_{\mathscr{B}_v}[b'_1]_{\mathscr{B}_v} \\ [f]_{\mathscr{B}_v}[b'_2]_{\mathscr{B}_v} \\ \vdots \\ [f]_{\mathscr{B}_v}[b'_k]_{\mathscr{B}_v} \end{bmatrix} = [v]^\intercal_{\mathscr{B}_v} [f]_{\mathscr{B}_v} [\mathscr{B}_w]_{\mathscr{B}_v} = 0$$

and hence we could write:

$$[\mathscr{W}^{\perp}]_{\mathscr{B}_v} = \mathrm{Ker}\Big[[\mathscr{B}_w]^{\mathsf{T}}_{\mathscr{B}_v}[f]^{\mathsf{T}}_{\mathscr{B}_v}\Big]$$

Using proposition 4.3.5 we rewrite the above as:

$$[\mathscr{W}^{\perp}]_{\mathscr{B}_v} = \mathrm{Ker}\Big[[f]_{\mathscr{B}_v}[\mathscr{B}_w]_{\mathscr{B}_v}\Big]$$

In order to get rid of $[f]_{\mathscr{B}_v}$ and be able to use the Rank-Nullity theorem, we turn again to our definition of the linear transformation $\mathrm{F}_{\cdot,u} : \mathscr{V} \to \mathscr{V}$ and remember that:

$$[f]_{\mathscr{B}_v} = [\mathrm{F}_{\cdot,u}]_{\mathscr{B}_v}$$

Now what we remember that by corollary 5.2.1:

$$[\mathrm{F}_{\cdot,u}]_{\mathscr{B}_v}[\mathscr{B}_w]_{\mathscr{B}_v} = [\mathrm{F}_{(\cdot,u)_{\mathscr{W}}}]_{\mathscr{B}_v}$$

where $\mathrm{F}_{(\cdot,u)_{\mathscr{W}}} : \mathscr{W} \to \mathscr{V}$ is the confinement of the linear transformation $\mathrm{F}_{\cdot,u}$ to the subspace $\mathscr{W}$. Thus we have got to the following:

$$\dim(\mathscr{W}^{\perp}) = \dim([\mathscr{W}^{\perp}]_{\mathscr{B}_v}) = \dim\left(\mathrm{Ker}\Big[[\mathrm{F}_{(\cdot,u)_{\mathscr{W}}}]_{\mathscr{B}_v}\Big]\right) = \eta\left(\mathrm{F}_{(\cdot,u)_{\mathscr{W}}}\right)$$

We now write proposition 5.2.2 for $\mathrm{F}_{\cdot,u}$:

$$\eta\left(\mathrm{F}_{(\cdot,u)_{\mathscr{W}}}\right) = \dim(\mathscr{V}) - \dim(\mathscr{W}) + \dim(\mathscr{W} \cap \mathrm{Ker}[\mathrm{F}_{\cdot,\mathrm{u}}])$$

This would rewrite as:

$$\dim(\mathscr{W}^{\perp}) = \dim(\mathrm{Ker}[f_{\mathscr{W}}]) = \dim(\mathrm{Ker}[\mathrm{F}_{(\cdot,u)_{\mathscr{W}}}]) = \dim(\mathscr{V}) - \dim(\mathscr{W}) + \dim(\mathscr{W} \cap \mathrm{Ker}[\mathrm{F}_{\cdot,\mathrm{u}}])$$

and we have seen earlier that $\mathrm{Ker}[\mathrm{F}_{\cdot,\mathrm{u}}]$ is in fact $\mathscr{V}^{\perp}$, and hence the proof is complete. □

Corollary 6.1.3. Two special cases of what we just proved are the following

(i) If $\mathscr{W} = \mathrm{span}\{u\}$ we get the following:

$$\dim(u^{\perp}) = \dim(\mathscr{V}) - 1 + \mathbb{I}\{u \in \mathscr{V}^{\perp}\}$$

(ii) If we know that f is non-degenerate, again as a special case of (iii) we could write:

$$\dim(\mathscr{W}^{\perp}) = \dim(\mathscr{V}) - \dim(\mathscr{W})$$

6.2 The signature of symmetric bilinear forms and their diagonalization

Proposition 6.2.1. *Let $\mathscr{V}$ be a vector space defined over $\mathbb{F}$ and $f : \mathscr{V} \times \mathscr{V} \to \mathbb{F}$ be a bilinear form over $\mathscr{V}$. There exists a basis $\mathscr{B} = \{b_1, \ldots, b_n\}$ for $\mathscr{V}$ such that all b_i are mutually f-orthogonal. In this case as we mentioned earlier $[f]_{\mathscr{B}}$ would be diagonal. Furthermore if the zero entries of the diagonal of $[f]_{\mathscr{B}}$ are the entries at positions $\pi_1, \pi_2, \ldots, \pi_k$ the following would hold:*

$$\eta(f) = k$$

$$\mathscr{V}^{\perp} = \mathrm{span}\{b_{\pi_i}\}_{i=1}^{k}$$

Proof. We build $\mathscr{B}$ constructively by induction, and the idea would be to first take as much non-degeneracies, and then when all that is left is degeneracies (we reach a subspace on which the radical of f swallows the whole of the subspace) any choice of vectors would work. We perform induction on $\dim(\mathscr{V})$. The basis is obvious. Now assume for all vector spaces of dimensionality smaller than $\dim(\mathscr{V})$ we can find an f-orthogonal basis. At this point either, $\mathscr{V}^\perp$, the radical of f, coincides with $\mathscr{V}$, in which case we are already done. Otherwise there would exist a vector $u \in \mathscr{V}$ that is not a degeneracy. By corollary 6.1.2 we know that we can choose u to not be weird, and hence projection on u would be well-defined. By corollary 6.1.3 we know that $u^\perp$ has dimensionality $\dim(\mathscr{V}) - 1$ and furthermore since $f(u,u) \neq 0$ it would follow that $u \notin u^\perp$. By the hypothesis of induction there exists an f-orthogonal basis $\mathscr{B}'$ for $u^\perp$. Obviously since u is not weird, $u \cup \mathscr{B}'$ is a basis for $\mathscr{V}$. We would then easily construct in a Gram-Schmidt fashion an f-orthogonal basis from $u \cup \mathscr{B}'$, by striking out one at a time, the non-orthogonal component of each of b'_i.
Now since $\eta(f) = \dim(\mathscr{V}^\perp) = k$ all we need to show is that b_{π_i} all belong to the radical. In other words we have to show that b_{π_i}s are all degeneracies. To see this we notice that:

$$[f(\cdot, b_{\pi_i})]_{\mathscr{B}} = \mathrm{A}[b_{\pi_i}]_{\mathscr{B}}$$

but we notice that A is diagonal with a zero at its π_i-th diagonal, and $[b_{\pi_i}]_{\mathscr{B}}$ is an all zero column vector, except for a 1 at its π_i-th position. The result of the above multiplication, therefore, is obviously the zero linear functional, and the proof is complete. □

Remark. Notice that what we have proved, in matrix terminology, is *not* the following pretentious statement: "Any symmetric matrix is similar to a diagonal matrix". Remember that the change of basis formula for bilinear forms (proposition 6.0.4) is:

$$[f]_{\mathscr{B}'} = \mathrm{P}^\intercal [f]_{\mathscr{B}} \mathrm{P}$$

where, although P is a basis change matrix, and hence non-singular, but the above is definitely different from the similarity relationship:

$$\mathrm{A} \sim \mathrm{P}^{-1} \mathrm{A} \mathrm{P}$$

and we have not proved the existence of a P in the latter sense (well, it does *not* exist!).

Remark. Since the rank of f happens to be equal to the number of non-zero diagonal entries of $[f]_{\mathscr{B}}$ if $\mathscr{B}$ is such that the representation matrix becomes diagonal, we conclude that for all the bases $\mathscr{B}$ that this happens, the number of zero and non-zero diagonals is the same, $\eta(f)$ and $\rho(f)$ respectively.

Corollary 6.2.1. If $\mathbb{F} = \mathbb{C}$ we could modify the f-orthogonal basis $\mathscr{B}$ for $\mathscr{V}$, the existence of which is guaranteed by proposition 6.2.1, in the following fashion:

$$b_i \leftarrow \frac{1}{\sqrt{f(b_i, b_i)}} b_i$$

for those of b_i that are not degeneracies, and after a trivial rearrangement the modified basis $\mathscr{B}$ would satisfy:

$$[f]_{\mathscr{B}} = \begin{bmatrix} \mathrm{I}_{r \times r} & 0 \\ 0 & 0 \end{bmatrix}$$

where $r = \rho(f)$, the rank of f. In the same fashion, in the case of $\mathbb{F} = \mathbb{R}$, with the following updates:

$$b_i \leftarrow \frac{1}{\sqrt{|f(b_i, b_i)|}} b_i$$

we will get:

$$[f]_{\mathscr{B}} = \begin{bmatrix} \mathrm{I}^{\pm}_{r\times r} & 0 \\ 0 & 0 \end{bmatrix}$$

where $\mathrm{I}^{\pm}_{r\times r}$ is the following matrix:

$$\mathrm{I}^{\pm}_{r\times r} = \begin{bmatrix} \pm 1 & 0 & \dots & 0 \\ 0 & \pm 1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \pm 1 \end{bmatrix}$$

Definition 6.2.1. As we just proved for any symmetric bilinear form, one can find bases for $\mathscr{V}$ such that none of the members are *weird* (isotropic non-degeneracies). For such bases projection on the members of the basis is completely well-defined. We will investigate this issue more later. We call such a basis, which is f-orthogonal and weird-free an f-**good basis**!

Proposition 6.2.2. *Let $\mathscr{V}$ be a vector space defined over $\mathbb{F}$ and $f : \mathscr{V}\times\mathscr{V}\to\mathbb{F}$ be a symmetric bilinear form over $\mathscr{V}$. From proposition 6.0.4 we can not infer that all matrix representations $[f]_{\mathscr{B}}$ of f are similar matrices. However it can easily be seen that (non-)singularity is preserved through different matrix representations of f. We claim that f is degenerate if and only if its matrix representations are singular.*

Proof. We recall from proposition 4.0.10 that any linear functional is either the zero functional or has nullity equal to $n-1$. The bilinear form f is, by definition, non-degenerate if and only if there is no vector $u \in \mathscr{V}$, such that $f(\cdot, u)$ is zero. We fix some basis $\mathscr{B}$ and refer to $[f]_{\mathscr{B}}$ by A. We know that:

$$[f(\cdot, u)]_{\mathscr{B}} = \mathrm{A}[u]_{\mathscr{B}}$$

As a result, $f(\cdot, u)$ is zero if and only if $\mathrm{A}[u]_{\mathscr{B}}$ is the zero vector. We know from corollary 3.3.5 that the square matrix A is non-singular if and only if $\mathrm{A}x = 0$ has no non-trivial solution, and hence f is non-degenerate if and only if A is non-singular. □

Definition 6.2.2. We here extend the notion of $\mathscr{V}_{(0)}$. Any symmetric bilinear form f divides the space $\mathscr{V}$ to three disjoint subsets (none of which is usually subspaces): $\mathscr{V}_{(+)}$ is the *set* of all the vectors over which f is positive definite (i.e. $f(u, u) > 0$), and $\mathscr{V}_{(-)}$ is the *set* of all the vectors over which f is negative definite (i.e. $f(u, u) < 0$).

Proposition 6.2.3. *The following basis properties hold for $\mathscr{V}_{(-)}, \mathscr{V}_{(0)}$ and $\mathscr{V}_{(+)}$:*

(i) *$\mathscr{V}^{\perp} \subseteq \mathscr{V}_{(0)}$, equality possibly not even holding if f is non-degenerate: $\mathscr{V}^{\perp} = \{0\} \subset \mathscr{V}_{(0)}$.*

(ii) *$\mathscr{V} = \mathscr{V}_{(-)} \cup \mathscr{V}_{(0)} \cup \mathscr{V}_{(+)}$.*

(iii) *$\mathscr{V}_{(-)}$, $\mathscr{V}_{(0)}$, and $\mathscr{V}_{(+)}$ are disjoint collections of lines (0 excluded from each line) in $\mathscr{V}$ and hence are possibly not subspaces.*

(iv) *If f is (non-)positive(negative)-definite over some set of k mutually f-perpendicular vectors $\{u_i\}_{i=1}^k$, then so is it all over* $\text{span}\{u_i\}_{i=1}^k$.

(v) *Any set of three (and no more) vectors, each from one of the $\mathscr{V}_{(.)}$ are linearly independent.*

(vi) *Despite* (v) *since $\mathscr{V}_{(.)}$s are not subspaces, they are not linearly independent of each other as sets, and hence* $\text{span}\{\mathscr{V}_{(.)}\}$*s intersect non-trivially.*

Proof.

(i) Any degeneracy is obviously isotropic, as a result the radical is contained in $\mathscr{V}_{(0)}$ (but there exist isotropic vectors that are not degeneracies). To see how a non-degenerate bilinear form might have a non-trivial $\mathscr{V}_{(0)}$ take $\mathscr{V}$ to be $\mathbb{R}^2$ and define f to be the following:

$$f(x,y) = x^\intercal \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} y$$

Expanding the above we get:

$$f(x,y) = x_1y_2 + x_2y_1$$

Obviously this is a non-degenerate symmetric bilinear form, but $\mathscr{V}_{(0)}$ consists of the lines $x_1 = 0$ and $x_2 = 0$.

(ii) trivial.

(iii) To show that $\mathscr{V}_{(.)}$s are all collection of lines (subspaces of dimensionality 1), we just need to prove that if u is in any of them so is αu for any scalar α, which is straightforward since $\alpha^2 \geq 0$ and hence can not change the sign of $f(\alpha u, \alpha u)$ from whatever sign $f(u,u)$ has. Also notice that if $u, v \in \mathscr{V}_{(.)}$ we have:

$$\text{sgn}\{f(u+v,u+v)\} = \text{sgn}\{f(u,u) + f(v,v) + 2f(u,v)\}$$

and although $f(u,u)$ and $f(v,v)$ have the same sign, without any knowledge about $f(u,v)$, $u+v$ could possibly fall out of that $\mathscr{V}_{(.)}$ that u, v belong to.

(iv) Let u be some linear combination of the u_i. We have:

$$f(u,u) = f\left(\sum_{i=1}^k \alpha_i u_i, \sum_{i=1}^k \alpha_i u_i\right) = \sum_{i=1}^k \alpha_i^2 f(u_i,u_i) + \sum_{i \neq j} \alpha_i \alpha_j f(u_i,u_j)$$

from the assumption that u_is are all mutually f-perpendicular, the last term vanishes:

$$f(u,u) = \sum_{i=1}^k \alpha_i^2 f(u_i,u_i)$$

and since all α_i^2s are positive and for all i, $\text{sgn}\{f(u_i,u_i)\}$ is the same we will get:

$$\text{sgn}\{f(u,u)\} = \text{sgn}\{f(u_i,u_i)\}$$

(v) Let $u \in \mathscr{V}_{(+)}, v \in \mathscr{V}_{(-)}$ and $w \in \mathscr{V}_{(0)}$. Since all $\mathscr{V}_{(\cdot)}$s are collections of lines, for any vanishing linear combination of u, v, w, we could reselect u, v, w such that $u + v + w = 0$. Now notice this results in:

$$f(u, u + v + w) = 0$$
$$f(v, u + v + w) = 0$$

expanding the two we get:

$$f(u, u) + f(u, v) = 0$$
$$f(v, v) + f(u, v) = 0$$

now notice that if $f(u, v) \geq 0$ the first equality results in $f(u, u) \leq 0$ which is a contradiction, and if $f(u, v) \leq 0$ the second results in $f(v, v) \geq 0$ which again is a contradiction.

□

Definition 6.2.3. We have seen that for any symmetric bilinear form f, the sets $\mathscr{V}_{(+)}, \mathscr{V}_{(-)}$ and $\mathscr{V}_{(0)}$ are collections of 0-excluded lines and not subspaces. However these do not need to be just a bunch of streaks, and at some parts of $\mathscr{V}$ the respective lines of each set could form together a contiguous subspace of dimensionality larger than 1. In the case of $\mathscr{V}_{(0)}$ we can easily prove that $\mathscr{V}^{\perp}$ is the unique largest subspace therein. We define $\mathscr{V}^{+}$ to be *any* of the subspaces with maximal dimensionality lying completely inside $\mathscr{V}_{(+)} \cup \{0\}$, i.e. all over which f is positive definite (except for 0). We refer to $\mathscr{V}^{+}$ by any of the **largest f-PD subspaces** in $\mathscr{V}$. In the same fashion we define $\mathscr{V}^{-}$ to refer to *any* of the subspaces with maximal dimensionality lying completely within $\mathscr{V}_{(-)} \cup \{0\}$, i.e. all over which f is negative definite (except for 0). We refer to $\mathscr{V}^{+}$ by any of the **largest f-ND subspaces** in $\mathscr{V}$. Although the choice of $\mathscr{V}^{+}$ and $\mathscr{V}^{-}$ is not unique but their respective dimensionalities are unique. We define:

$$n_{+}(f) := \dim(\mathscr{V}^{+})$$
$$n_{-}(f) := \dim(\mathscr{V}^{-})$$

and for consistency we define $n_0(f)$ to be $\eta(f)$:

$$n_0(f) = \dim(\mathscr{V}^{\perp})$$

Since the just defined subspaces mutually intersect only trivially, the following obviously holds:

Proposition 6.2.4. *The subspaces $\mathscr{V}^{+}, \mathscr{V}^{-}$ and $\mathscr{V}^{\perp}$ are all mutually linearly independent.*

The tuple $\langle n_{+}(f), n_{-}(f), n_0(f)\rangle$ is referred to by the *signature* of a symmetric bilinear form. We will now prove the last proposition of this section, the Sylvester's law of inertia, that asserts that in any diagonalization of a symmetric bilinear form, the diagonal has positive, negative and zero entries in accordance with the signature:

Proposition 6.2.5 (Sylvester's Law of Inertia)**.** *Let f by a symmetric bilinear form on the vector space $\mathscr{V}$. Let $\mathscr{B}$ be an f-orthogonal basis for $\mathscr{V}$, the existence of which is assured by proposition 6.2.1. For any such $\mathscr{B}$ the number of positive, negative, and zero diagonal entries of $[f]_{\mathscr{B}}$ are exactly $n_{+}(f)$, $n_{-}(f)$, and $n_0(f)$. Consequently for any symmetric bilinear form we will have $n_{+}(f) + n_{-}(f) + n_0(f) = \dim(\mathscr{V})$, and obviously we would have:*

$$\mathscr{V} = \mathscr{V}^{+} \oplus \mathscr{V}^{-} \oplus \mathscr{V}^{\perp}$$

Proof. We name the members of $\mathscr{B}$ to be $b_1, b_2, \ldots, b_n$. Without loss of generality we will assume that the first t diagonal entries in $[f]_{\mathscr{B}}$ are positive, the next s entries are negative, and the rest are zero. We know that the number of zero diagonals is exactly $n_0(f) = \eta(f)$. We know from part (iv) of proposition 6.2.3 that f is positive definite all over $\text{span}\{b_i\}_{i=1}^{t}$. From the definition of $n_+(f)$ it would immediately follow that $t \leq n_+(f)$. With the same argument we would have $s \leq n_-(f)$. Now by looking at the following:

$$n = t + s + n_0(f) \leq n_+(f) + n_-(f) + n_0(f) \leq n$$

it immediately follows that for both inequalities equality should hold and thus the proof is complete. □

Corollary 6.2.2. In other words for any symmetric bilinear form f, any f-good basis for $\mathscr{V}$, i.e. weird-free and f-orthogonal, has exactly $n_+(f)$ vectors on which f is positive definite, exactly $n_-(f)$ vectors on which f is negative definite, and exactly $n_0(f)$ degeneracies.

Corollary 6.2.3. If f is a positive definite symmetric bilinear form over $\mathbb{R}^n$, i.e. $n_+(f) = \dim(\mathscr{V})$ there exists a basis $\mathscr{B}$ such that $[f]_{\mathscr{B}}$ is the identity matrix. As a result for any symmetric positive definite real matrix A there exists a non-singular matrix P such that:

$$\mathrm{A} = \mathrm{P}^{\intercal}\mathrm{P}$$

6.3 Inner Product Spaces and applications in numerical linear algebra

Definition 6.3.1. Let ϕ be a positive definite symmetric bilinear form over the vector space $\mathscr{V}$. We call ϕ an **inner product** over the space $\mathscr{V}$.

Proposition 6.3.1. *The following properties immediately follow from the definition of an inner product ϕ:*

(i) *ϕ is non-degenerate.*

(ii) *ϕ does not allow for any non-zero isotropic vectors.*

(iii) $n_+(\phi) = \dim(\mathscr{V})$.

(iv) *There exists a basis $\mathscr{B}$ for $\mathscr{V}$ such that $[\phi]_{\mathscr{B}} = \mathrm{I}_{n\times n}$, where n is the dimensionality of $\mathscr{V}$. We would say ϕ is a dot product with respect to some basis $\mathscr{B}$:*

$$\phi(u, v) = [u]_{\mathscr{B}}^{\intercal}[v]_{\mathscr{B}}$$

and that $\mathscr{B}$ is an orthonormal basis w.r.t ϕ.

(v) *For an arbitrary basis $\mathscr{B}'$ we would have:*

$$\phi(u, v) = [u]_{\mathscr{B}'}^{\intercal}\mathrm{A}[v]_{\mathscr{B}'}$$

for some SPD matrix A *which is $[\phi]_{\mathscr{B}'}$.*

(vi) *For any inner product ϕ, if $\mathscr{B}_1$ and $\mathscr{B}_2$ are two ϕ-orthonormal bases, we will have:*

$$[\mathscr{B}_1]_{\mathscr{B}_2}^{\intercal}[\mathscr{B}_1]_{\mathscr{B}_2} = \mathrm{I}_{n\times n}$$

and we would say that the matrix $\mathrm{A} = [\mathscr{B}_1]_{\mathscr{B}_2}$ is an orthogonal matrix: $\mathrm{A}^{\intercal}\mathrm{A} = \mathrm{A}\mathrm{A}^{\intercal} = \mathrm{I}$.

If ϕ is an inner product we usually use the following notation:

$$\langle u, v\rangle_{\phi} = \phi(u, v)$$

We saw that any inner product is in fact the dot product with respect to any basis. As a result one could think of defining orthogonality in inner product spaces not as we have done by now, with respect to an inner product, but with respect to a basis:

Definition 6.3.2. Let $\mathscr{V}$ be a vector space and $\mathscr{B}$ be any arbitrary basis for $\mathscr{V}$. We say that two vectors are $\mathscr{B}$-orthogonal if we have:

$$[u]_{\mathscr{B}}^{\intercal}[v]_{\mathscr{B}} = 1$$

Obviously the members of $\mathscr{B}$ are themselves mutually $\mathscr{B}$-orthogonal, for any choice of $\mathscr{B}$. But if a basis $\mathscr{B}'$ is such that its members are $\mathscr{B}$-orthogonal, we say that $\mathscr{B}'$ is a $\mathscr{B}$-**orthonormal** basis. Of course in that case $\mathscr{B}$ is a $\mathscr{B}'$-orthonormal basis as well.

Remark. In the case of $\mathscr{V} = \mathbb{F}^n$, as always, the existence of a natural basis leads us to define a natural orthogonality criteria as well, and hence we say that a basis, and consequently a matrix, is orthogonal, with no reference to anything, if it is $\mathscr{E}$-orthogonal where $\mathscr{E} = \{e_1, \ldots, e_n\}$ is the natural basis for $\mathbb{F}^n$.

6.3.1 Projectors, Reflectors, Gram Schmidt, and the QR decomposition

6.3.2 *** Linear Least Squares Problem

6.3.3 *** Singular Value Decomposition

Chapter 7

Singularity checks for linear maps: Determinants and Condition numbers

In this chapter first I prove, borrowing the rough blueprint from the 5^{th}chapter of [HK71], that by setting simple and intuitive requirements for a function to be an alarm for singularity, existence and uniqueness of such function (the determinant) can be proved. Second, I will go through the notion of condition numbers, and see how they are more reliable alarms for singularity of linear systems.

7.1 Derivation of the determinant

The terminology I have used in this section does not match those of the references, since my derivation is quite different anyway. In fact, there is no unanimously agreed upon flow of arguments towards building the notion of determinants; some even push the diversity to a really confusing level by just introducing it as a magical divine rule bestowed on us from history, and not as an innate property of linear transformations but as a property of matrices. I do not find defining determinants with formulae at all useful, since as we will see, they are not that useful of a tool for their historically primitive goals. For a tool designed specifically for computational reasons, the determinant is neither as computationally efficient, nor structurally informative as other indices (such as say the condition number). One might therefore ask, reasonably, "why is there at least a chapter of discussion about determinants in any linear algebra and matrix analysis textbook?". The turns of events, the answer is that, engraved them in the theory of linear algebra and matrices due to the critical role they ended up playing in the theoretical treatise of canonical forms, and *not* because they are useful indicators of singularity of linear systems.

Determinants were originally an extension of the "indicator"s of singularity of small linear systems. We already familiar with bilinear functions of the form $f : \mathscr{V} \times \mathscr{V} \to \mathbb{F}$. The idea would be to build a scalar function which takes n vectors in an n-dimensional vector space and sets out an alarm (by being zero) if only if the set of vectors are linearly dependent. This would be an indicator of singularity of a linear system of size n. A natural extension of the tools we already have built would be to define an *n-linear* function of n vectors in the fashion we are just about to describe. It is important to notice that, although this is definitely not the standard formulation of determinants, it the mind-setting we have been working in, and fairly easy to derive.

Definition 7.1.1. Let $\mathscr{V}$ be an n-dimensional vector space defined over a field $\mathbb{F}$. We define a function $g\colon \mathscr{V}^n \to \mathbb{F}$ to be n-**linear** if it by fixing any $n-1$ of its arguments it becomes a linear functional over the remaining argument. In other words the following should hold:

$$\left.\begin{array}{c} \forall u_1, u_2, \ldots, u_n \in \mathscr{V} \\ \forall i \leq n \\ \forall u_i' \in \mathscr{V} \\ \forall \alpha \in \mathbb{F} \end{array}\right\} \begin{array}{rl} g(u_1, \ldots, \alpha u_i, \ldots, u_n) = & \alpha g(u_1, \ldots, u_i, \ldots, u_n) \\ g(u_1, \ldots, u_i + u_i', \ldots, u_n) = & g(u_1, \ldots, u_i, \ldots, u_n) + \\ & g(u_1, \ldots, u_i', \ldots, u_n) \end{array}$$

This is nothing specific until now. Our requirement of g building and indicator of linear independence, is too complex. We break this property down, in light of n-linearity, to simpler criteria. Then little by little we will inject our requirements as properties of g and will, in the end, observe that there would exist only a unique n-linear function satisfying those properties.

Definition 7.1.2. Let $\mathscr{V}$ be an n-dimensional vector space defined over a field $\mathbb{F}$ and $g\colon \mathscr{V}^n \to \mathbb{F}$ be any n-linear function over $\mathscr{V}$. We say that g is a **order-independent alarm** if for any $u_1, \ldots, u_n$ that $g(u_1, \ldots, u_n)$ is zero, so is $g(u_{\pi_1}, u_{\pi_2}, \ldots, u_{\pi_n})$ for any permutation π of $1, 2, \ldots, n$.

Remark. Notice that indifference of g towards permutations of arguments, does not mean g should be insensitive to permutation of the arguments in general, and hence it would be wrong to say g is order-independent. We only expect g to be consistent with respect to it setting the "alarm" which is when its arguments are such that its value vanishes.

We now prove that our requirement of an n-linear to be zero if and only if the arguments are linearly dependent, translates, in light of the n-linearity property, to a much simpler criterion.

Proposition 7.1.1. *Let $\mathscr{V}$ be an n-dimensional vector space defined over a field $\mathbb{F}$ and $g\colon \mathscr{V}^n \to \mathbb{F}$ be a non-trivial n-linear order-independent-alarm function over $\mathscr{V}$. The following are equivalent:*

(i) $g(u_1, \ldots, u_n) = 0$ *if and only if* $\{u_1, \ldots, u_n\}$ *is linearly dependent.*

(ii) *Whenever* $u_i = u_j$ *for some* i, j *we have* $g(u_1, \ldots, u_n) = 0$.

Proof. The fact (ii) follows if (i) holds is trivial. We prove the other direction. Let $g(u_1, \ldots, u_n) = 0$ whenever some u_i, u_j coincide. Now assume we have a bunch of linearly dependent vectors $\{v_1, \ldots, v_n\}$. There exists a vanishing linear combination of v_i:

$$\alpha_1 v_1 + \alpha_2 v_2 + \ldots + \alpha_n v_n = 0$$

Without loss of generality we can assume $\alpha_1 \neq 0$. We will write:

$$g\left(\sum_{i=1}^{n} \alpha_i v_i, v_2, v_3, \ldots, v_n\right) = \sum_{i=1}^{n} \alpha_i g(v_i, v_2, \ldots, v_n)$$

The summands on the right hand side all vanish by assumption, except for the one for $i = 1$, thus we will get:

$$0 = g(\sum_{i=1}^{n} \alpha_i v_i, v_2, v_3, \ldots, v_n) = \alpha_1 g(v_1, v_2, \ldots, v_n)$$

from which it immediately follows $g(v_1, \ldots, v_n) = 0$. Now assume $g(v_1, \ldots, v_n) = 0$ holds, we will prove that v_is have to be linearly dependent if g has to be non-trivial. Assume $v_1, \ldots, v_n$ are linearly independent. They would span all of $\mathscr{V}$, as a result for any $\{u_1, \ldots, u_n\}$ in $\mathscr{V}$ they would all have representations in terms of $v_1, \ldots, v_n$. Lets say the representation of u_i in terms of $v_1, \ldots, v_n$ is $[\alpha_{i,1}, \alpha_{i,2}, \ldots, \alpha_{i,n}]^\intercal$, we could then write:

$$g(u_1, \ldots, u_n) = g\left(\sum_{j=1}^{n} \alpha_{1,j} v_j, \sum_{i=1}^{n} \alpha_{2,j} v_j, \ldots, \sum_{j=1}^{n} \alpha_{n,j} v_j\right)$$

which by n-linearity of g decomposes to a giant sum of terms all of which contain some $g(v_{\pi_1}, \ldots, v_{\pi_n})$. Any term on the right hand side whose its corresponding π_is are not distinct, thus by the assumption that g is an order-independent alarm, would vanish. And for terms whose π_is are distinct we will get a valid permutation of the arguments in $g(v_1, \ldots, v_n)$ which is again zero due to the fact that g is an order-independent alarm. As a result if $v_1, \ldots, v_n$ are linearly independent g would be the zero function, and hence the proof is complete. □

And now we are ready to define what we mean by a determinant function. Obviously we will need to prove the existence of such a function, and then fix normalization criteria to force the choice of the determinant become unique.

Definition 7.1.3. Let $\mathscr{V}$ be an n-dimensional vector space defined over a field $\mathbb{F}$. We call a function $g\colon \mathscr{V}^n \to \mathbb{F}$ a **determinant** function if it satisfies all the following:

(i) g is non-trivial.

(ii) g is n-linear.

(iii) g is an order-independent alarm.

(iv) whenever any two arguments of g coincide, g assigns 0 to the whole n-tuple (or the equivalent criterion in proposition 7.1.1).

Proposition 7.1.2 (Primitive properties of the determinant)**.** *Let $\mathscr{V}$ be an n-dimensional vector space defined over a field $\mathbb{F}$ and $g\colon \mathscr{V}^n \to \mathbb{F}$ be any determinant function over $\mathscr{V}$. The following properties hold for g:*

(i) *Flipping two arguments of g would negate its value: For any $u_1, \ldots, u_n$ we have:*

$$\begin{aligned} \forall i < j: \quad & g(u_1, \ldots, u_{i-1}, u_i, u_{i+1}, \ldots, u_{j-1}, u_j, u_{j+1}, \ldots, u_n) = \\ & -1 \times g(u_1, \ldots, u_{i-1}, u_j, u_{i+1}, \ldots, u_{j-1}, u_i, u_{j+1}, \ldots, u_n) \end{aligned}$$

(ii) *Adding any scalar multiple of any of the arguments to any other argument does not change the value of g:*

$$\begin{aligned} \forall \alpha \in \mathbb{F},\ \forall i \neq j: \quad & g(u_1, \ldots, u_i, \ldots, \quad u_j, \ldots, u_n) = \\ & g(u_1, \ldots, u_i, \ldots, \quad \alpha u_i + u_j, \ldots, u_n) = \end{aligned}$$

(iii) *If any of the arguments is zero, so is g.*

Proof.

(i) We notice that:

$$\begin{aligned}\forall i < j: \qquad & g(u_1, \ldots, u_{i-1}, u_i, u_{i+1}, \ldots, u_{j-1}, u_j, u_{j+1}, \ldots, u_n) + \\ & g(u_1, \ldots, u_{i-1}, u_j, u_{i+1}, \ldots, u_{j-1}, u_i, u_{j+1}, \ldots, u_n) = \\ & g(u_1, \ldots, u_{i-1}, u_i + u_j, u_{i+1}, \ldots, u_{j-1}, u_i + u_j, u_{j+1}, \ldots, u_n) = 0\end{aligned}$$

(ii) We see that:

$$\begin{aligned}\forall \alpha \in \mathbb{F},\ \forall i \neq j: \quad & g(u_1, \ldots, u_i, \ldots, \quad \alpha u_i + u_j, \ldots, u_n) = \\ & \alpha g(u_1, \ldots, u_i, \ldots, \quad u_i, \ldots, u_n) + \\ & g(u_1, \ldots, u_i, \ldots, \quad u_j, \ldots, u_n)\end{aligned}$$

where the summand on the left obviously vanishes.

(iii) trivial.

□

Remark. Notice that in (ii) one could not put $\alpha u_i + u_j$ in the place of u_i and expect g to not change. One could just add a multiple of u_i to any of the u_j (with no coefficient on u_j).

Corollary 7.1.1. An important observation can be made here. We know from combinatorics that any permutation can be reached by a the combination of a number of "flip"s. Furthermore all the "flip sequence"s that produce the same permutation either all consist of an even number of flips, or all consist of an odd number of flips. This simple theorem neatly identifies any permutation to be *odd* or *even* depending on the type of its equivalent flip sequences. Now if we define the *sign* of a permutation in the following fashion:

$$\operatorname{sgn}(\pi) = \begin{cases} +1 & \pi \text{ is a even} \\ -1 & \pi \text{ is a odd} \end{cases}$$

it would immediately follow from (i) that for any determinant function, any vectors $u_1, \ldots, u_n$ and any permutation $\pi = \langle \pi_1, \ldots, \pi_n \rangle$ of $1, \ldots, n$ one could write:

$$g(u_{\pi_1}, u_{\pi_2}, \ldots, u_{\pi_n} = (-1)^{\operatorname{sgn}(\pi)} g(u_1, u_2, \ldots, u_n) = \operatorname{sgn}(\pi) \times g(u_1, u_2, \ldots, u_n)$$

Here we prove a crucial result. We here prove the existence and uniqueness of a determinant function *with respect to* some basis $\mathscr{B}$, the choice of which is arbitrary, by adding a single requirement. We will then use this to build the determinants of linear transformations. Through the course of the proof we derive a closed formula for the computation of the determinant, although this is almost never the way the determinant is computed efficiently.

Proposition 7.1.3. *Let $\mathscr{V}$ be an n-dimensional vector space defined over a field $\mathbb{F}$ and $\mathscr{B} = \{b_1, \ldots, b_n\}$ be any arbitrary basis for $\mathscr{V}$. There exists a unique determinant function $g_{\mathscr{B}}$ satisfying:*

$$g_{\mathscr{B}}(b_1, \ldots, b_n) = 1$$

Furthermore any other choice but 1 for the above, would change the unique $g_{\mathscr{B}}$ only by a scalar factor.

Proof. We first assume such a determinant function exists, and see how the value of $g_{\mathscr{B}}$ would be uniquely identified by the assumption $g_{\mathscr{B}}(\mathscr{B}) = 1$. Then we could easily check that the closed form derivation we get is in fact a determinant function, i.e. satisfies properties in the definition of determinants we provided.
Let $u_1, u_2, \ldots, u_n$ be any arbitrary set of vectors in $\mathscr{V}$. Each have a representation in terms of $\mathscr{B}$. Let:

$$[u_i]_{\mathscr{B}} = \begin{bmatrix} \alpha_{i,1} \\ \alpha_{i,2} \\ \vdots \\ \alpha_{i,n} \end{bmatrix}$$

It is now easy to see in the same fashion as we saw in 7.1.1 that $g_{\mathscr{B}}(u_1, \ldots, u_n)$ rewritten as the following:

$$g_{\mathscr{B}}(u_1, \ldots, u_n) = g_{\mathscr{B}}\left(\sum_{j=1}^{n} \alpha_{1,j} b_j, \sum_{i=1}^{n} \alpha_{2,j} b_j, \ldots, \sum_{j=1}^{n} \alpha_{n,j} b_j\right)$$

would decompose into a giant sum over different permutations of $g_{\mathscr{B}}(\mathscr{B})$. The uniqueness of this sum is ensured from uniqueness of representation in terms of bases. We here can easily see that any other choice for $g_{\mathscr{B}}(\mathscr{B})$ will multiply all the summands of this giant sum by a same scalar and hence modifying g only by a scalar factor. Let's now go through the decomposition:

$$\begin{aligned}
g_{\mathscr{B}}(u_1, \ldots, u_n) &= g_{\mathscr{B}}\left(\sum_{j=1}^{n} \alpha_{1,j} b_j, \sum_{i=1}^{n} \alpha_{2,j} b_j, \ldots, \sum_{j=1}^{n} \alpha_{n,j} b_j\right) \\
&= g_{\mathscr{B}}\left(\alpha_{1,1} b_1 + \alpha_{1,2} b_2 + \ldots, \alpha_{1,n} b_n, \sum_{i=1}^{n} \alpha_{2,j} b_j, \ldots, \sum_{j=1}^{n} \alpha_{n,j} b_j\right) \\
&= \sum_{\pi_1=1}^{n} \alpha_{1,j} g_{\mathscr{B}}\left(b_{\pi_1}, \sum_{i=1}^{n} \alpha_{2,j} b_j, \ldots, \sum_{j=1}^{n} \alpha_{n,j} b_j\right) \\
&= \sum_{\pi_1=1}^{n} \sum_{\pi_2=1}^{n} \alpha_{1,\pi_1} \alpha_{2,\pi_2} g_{\mathscr{B}}\left(b_{\pi_1}, b_{\pi_2}, \ldots, \sum_{j=1}^{n} \alpha_{n,j} b_j\right)
\end{aligned}$$

We can carry on this process all the way, until we get the following:

$$g_{\mathscr{B}}(u_1, \ldots, u_n) = \sum_{\pi_1=1}^{n} \sum_{\pi_2=1}^{n} \ldots \sum_{\pi_n=1}^{n} \left[\alpha_{1,\pi_1} \alpha_{2,\pi_2} \ldots \alpha_{n,\pi_n} \times g_{\mathscr{B}}(b_{\pi_1}, b_{\pi_2}, \ldots, b_{\pi_n})\right]$$

Notice that of all the terms in the final sum, the only ones that do not vanish are those whose π_i are all distinct, i.e the corresponding $\pi = \langle \pi_1, \ldots, \pi_n \rangle$ is a valid permutation of $1, \ldots, n$. As a result we can rewrite this result in the following form:

$$g_{\mathscr{B}}(u_1, \ldots, u_n) = \sum_{\pi} \left[\alpha_{1,\pi_1} \alpha_{2,\pi_2} \ldots \alpha_{n,\pi_n} \times g_{\mathscr{B}}(b_{\pi_1}, b_{\pi_2}, \ldots, b_{\pi_n})\right]$$

where the summation is a over all the $n!$ permutations $\pi = \langle \pi_1, \ldots, \pi_n \rangle$ of $1, \ldots, n$. Also we know by corollary 7.1.1 that:

$$g_{\mathscr{B}}(b_{\pi_1}, b_{\pi_2}, \ldots, b_{\pi_n}) = \text{sgn}(\pi) \times g_{\mathscr{B}}(b_1, b_2, \ldots, b_n) = \text{sgn}(\pi)$$

which gets us to the final closed form formula for the determinant:

$$g_{\mathscr{B}}(u_1, \ldots, u_n) = \sum_{\pi} \left[\text{sgn}(\pi) \prod_{i=1}^{n} \alpha_{i,\pi_i} \right]$$

Let's now prove that this function is a determinant function:

(i) The function derived above is n-linear:

$$g_{\mathscr{B}}(u_1, \ldots, \alpha u_k, \ldots, u_n) = \sum_{\pi} \left[\text{sgn}(\pi) \alpha \prod_{i=1}^{n} \alpha_{i,\pi_i} \right] = \alpha g_{\mathscr{B}}(u_1, \ldots, u_k, \ldots, u_n)$$

To see why the first derivation holds, notice that for any permutation π exactly one of the terms in the product term has something to do with u_k.

$$\begin{aligned} g_{\mathscr{B}}(u_1, \ldots, u_k + u'_k, \ldots, u_n) &= \sum_{\pi} \left[\text{sgn}(\pi) \alpha \prod_{i \neq k} \alpha_{i,\pi_i} (\alpha_{k,\pi_k} + \alpha'_{k,\pi_k}) \right] \\ &= \sum_{\pi} \left[\text{sgn}(\pi) \alpha \prod_{i \neq k} \alpha_{i,\pi_i} \alpha_{k,\pi_k} \right] + \sum_{\pi} \left[\text{sgn}(\pi) \alpha \prod_{i \neq k} \alpha_{i,\pi_i} \alpha'_{k,\pi_k} \right] \\ &= g_{\mathscr{B}}(u_1, \ldots, u_k, \ldots, u_n) + g_{\mathscr{B}}(u_1, \ldots, u'_k, \ldots, u_n) \end{aligned}$$

(ii) The function derived above is obviously an order-independent alarm.

(iii) If $u_i = u_j$, it would follow that for any permutation π there exists another permutation π' resulting from flipping π_i and π_j that satisfies $\text{sgn}(\pi) = -1 \times \text{sgn}(\pi')$. Since $u_i = u_j$ obviously we would have:

$$\prod_{i=1}^{n} \alpha_{i,\pi_i} = \prod_{i=1}^{n} \alpha_{i,\pi'_i}$$

from this it would be obvious that all the $n!$ terms in the big sum decomposes to $\frac{n!}{2}$ couples of equal but negative values, and hence g would be zero.

□

Corollary 7.1.2. If instead of expanding u_i in terms of the basis over which we chose our determinant, we expand the u_i in terms of some other basis $\mathscr{B}'$ we could, with the same flow of arguments, get the following:

$$g_{\mathscr{B}}(u_1, u_2, \ldots, u_n) = g_{\mathscr{B}'}(u_1, \ldots, u_n) \times g_{\mathscr{B}}(b'_1, \ldots, b'_n)$$

This is exactly the scalar factor that $g_{\mathscr{B}'}$ and $g_{\mathscr{B}}$ differ in:

$$g_{\mathscr{B}}(\cdot) = g_{\mathscr{B}}(\mathscr{B}') \times g_{\mathscr{B}'}(\cdot)$$

It would immediately follow that:

$$g_{\mathscr{B}}(\mathscr{B}') \times g_{\mathscr{B}'}(\mathscr{B}) = 1$$

Corollary 7.1.3. For any non-singular linear transformation $\mathrm{T}\colon \mathscr{V} \to \mathscr{V}$ obviously the representation of $u_1, \ldots, u_n$ is the same as the representation of $\mathrm{T}u_1, \ldots, \mathrm{T}u_n$ in terms of $\mathrm{T}b_1, \ldots, \mathrm{T}b_n$ for any basis $\mathscr{B} = \{b_1, \ldots, b_n\}$ and again with the same argument as in the derivation of the closed form formula for the determinant we would get:

$$g_{\mathscr{B}}(\mathrm{T}u_1, \ldots, \mathrm{T}u_n) = g_{\mathscr{B}}(u_1, \ldots, u_n) \times g_{\mathscr{B}}(\mathrm{T}b_1, \ldots, \mathrm{T}b_n)$$

We finally are ready to define the determinant of a linear transformation:

Proposition 7.1.4. *Let* $\mathrm{T}\colon \mathscr{V} \to \mathscr{V}$ *be an endomorphism over the* n*-dimensional vector space* $\mathscr{V}$*. For any basis* $\mathscr{B} = \{b_1, b_2, \ldots, b_n\}$ *for* $\mathscr{V}$ *the following scalar value is the same:*

$$\det(\mathrm{T}) = g_{\mathscr{B}}(\mathrm{T}b_1, \mathrm{T}b_2, \ldots, \mathrm{T}b_n)$$

and hence we define this constant scalar to be the determinant of T.

Proof. If T is singular, for any basis $\mathscr{B}$ the set $\{\mathrm{T}b_1, \mathrm{T}b_2, \ldots, \mathrm{T}b_n\}$ is linearly dependent and hence $g_{\mathscr{B}}(\mathrm{T}b_1, \mathrm{T}b_2, \ldots, \mathrm{T}b_n)$ would be zero by the definition of $g_{\mathscr{B}}$. So we assume T is non-singular. Let $\mathscr{B}'$ and $\mathscr{B}$ be two bases for $\mathscr{V}$, we use the corollaries of the last proposition:

$$\begin{aligned} g_{\mathscr{B}'}(\mathrm{T}b'_1, \mathrm{T}b'_2, \ldots, \mathrm{T}b'_n) &= g_{\mathscr{B}'}(\mathscr{B}) \times g_{\mathscr{B}}(\mathrm{T}b'_1, \ldots, \mathrm{T}b'_n) \\ &= g_{\mathscr{B}'}(\mathscr{B}) \times g_{\mathscr{B}}(\mathscr{B}') \times g_{\mathscr{B}}(\mathrm{T}b_1, \ldots, \mathrm{T}b_n) \\ &= g_{\mathscr{B}}(\mathrm{T}b_1, \ldots, \mathrm{T}b_n) \end{aligned}$$

and the proof is complete. □

Corollary 7.1.4 (Properties of the determinant). For any endomorphism $\mathrm{T}\colon \mathscr{V} \to \mathscr{V}$ we have:

(i) $\det(\mathrm{T})$ is zero if and only if T is singular.

(ii) The matrix determinant of any of the similar matrices representing T are equivalent, in other words for any non-singular matrix P we have:

$$\det(\mathrm{A}) = \det(\mathrm{PAP}^{-1})$$

(iii) $\det(\mathrm{T} \circ \mathrm{S}) = \det(\mathrm{T}) \times \det(\mathrm{S})$ where S is any other endomorphism on $\mathscr{V}$. As a special case we would have:

$$\det(\mathrm{T}^m) = \det(\mathrm{T})^m$$

(iv) If T is non-singular $\det(\mathrm{T}) \times \det(\mathrm{T}^{-1}) = 1$.

(v) If A is an $n \times n$ triangular matrix we will have:

$$\det(\mathrm{A}) = \prod_{i=1}^{n} a_{i,i}$$

(vi) If $\mathscr{V}$ has the following T-invariant direct sum decomposition:

$$\mathscr{V} = \mathscr{W}_1 \oplus \mathscr{W}_2 \oplus \ldots \oplus \mathscr{W}_k$$

then each of the $T_{\mathscr{W}_i}$s are endomorphisms again. From proposition 5.2.5 we will have:

$$\det(T) = \prod_{i=1}^{k} \det(T_{\mathscr{W}_i})$$

(vii) Let A be a square block diagonal matrix:

$$A = \begin{bmatrix} A_{1,1} & 0 & \ldots & 0 \\ 0 & A_{2,2} & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \ldots & A_{k,k} \end{bmatrix}$$

Then the determinant of A is equivalent to the product of the determinants of its diagonal blocks:

$$\det(A) = \prod_{i=1}^{k} \det(A_{i,i})$$

7.2 Condition numbers and numerical stability

Chapter 8

Structural Breakdown of Endomorphisms: Rational and Jordan Canonical Forms

We remember from section 5.2 that invariant direct sum decompositions, if existing, provide a nice breakdown of an endomorphism, which not only renders a simple explanation of its behavior possible, but also provides us with zero-rich matrix representations, in which, when talking about numerical linear algebra, is what we are ultimately interested.
In this chapter we focus on a single arbitrary endomorphism $\mathrm{T}: \mathscr{V} \to \mathscr{V}$ over an arbitrary vector space $\mathscr{V}$, whose carrier field is *totally ordered* and provide a breakdown of its behavior by introducing a unique minimal T-invariant direct sum decomposition. The core issue through which we are going to attack the problem is the powers of a transformation, that are obviously closely tied to the notion of *polynomials*. We mentioned earlier that the elements in the space $\mathscr{L}(\mathscr{V}, \mathscr{V})$ together with transformation composition used as multiplication, form a non-commutative ring. We will now put to use the results from section 1.4 to open up the door to deriving powerful results leading to canonical forms.

8.1 Polynomials of linear maps

The first thing I will do in this section is to answer the question that "How would anyone evaluate a polynomial with scalar coefficients over a linear transformation?"; a question that every single textbook I referred to assumed trivial, and in fact, is not! Of course translating a linear map to a matrix reduces some (but not all) of the ambiguity, but even those of the references that insisted on canonical (basis-free, and hence matrix free) derivations did not work through this issue. We recall that given two arbitrary rings $(R, \odot_r, +)$ and $(S, \odot_s, +)$, we understood in which cases, and how, we can feed members of S to polynomials defined over R, i.e members of $R[x]$; and in simpler terms *evaluate* polynomials with coefficients in ring R (*the coefficient domain*) over objects from ring S (*the evaluation domain*). We did so through assuming some homomorphism between the two rings, namely $\varphi: R \to S$, and then defining the evaluation map with respect to φ to be a mapping $\mathrm{ev}_\varphi: R[x] \times S \to S$ that assigns to an arbitrary element $s \in S$ and a polynomial $p(x) \in R[x]$, namely the following polynomial:

$$p(x) = \alpha_0 + \alpha_1 \odot_r x + \alpha_2 \odot_r x^2 + \ldots + \alpha_n \odot_r x^n = \sum_{i=0}^{n} \alpha_i \odot_r x^n$$

a member of S, namely the following:

$$\mathrm{ev}_\varphi\Big(p(x), s\Big) = \sum_{i=0}^{n} \varphi(\alpha_i) \odot_s s^i \in S$$

Now let the coefficient domain be $\mathbb{F}$, where $\mathbb{F}$ is the carrier field of $\mathscr{V}$, and the evaluation domain be $\mathscr{L}(\mathscr{V}, \mathscr{V})$. The standard way, that all authors assume to be evident, of feeding a linear map to a polynomial with scalar coefficients is, in fact, the evaluation map with respect to the following homomorphism $\varphi\colon \mathbb{F} \to \mathscr{L}(\mathscr{V}, \mathscr{V})$:

$$\alpha \xrightarrow{\quad\varphi\quad} \alpha \mathrm{I}$$

Now let's look at what, in fact, the polynomial form:

$$\alpha_0 + \alpha_1 \mathrm{T} + \alpha_2 \mathrm{T}^2 + \ldots + \alpha_m \mathrm{T}^m$$

means. According to our evaluation map, if $p(x) = \sum \alpha_i x^i$ what we mean by the above is:

$$\mathrm{ev}_\varphi\Big(p(x), \mathrm{T}\Big) = \sum_{i=0}^{m} \varphi(\alpha_i) \times \mathrm{T}^i \in \mathscr{L}(\mathscr{V}, \mathscr{V})$$

where $\times$ is the multiplication of the ring $\mathscr{L}(\mathscr{V}, \mathscr{V})$ which is, as we have defined, the composition operator. Thus we can write:

$$\mathrm{ev}_\varphi\Big(p(x), \mathrm{T}\Big) = \sum_{i=0}^{m} \alpha_i \mathrm{I} \circ \mathrm{T}^i = \sum_{i=0}^{m} \alpha_i \mathrm{T}^i$$

and this is what we refer to when we write $p(\mathrm{T})$ for a polynomial $p(x)$ and a linear map T.

Proposition 8.1.1. *Although $\mathscr{L}(\mathscr{V}, \mathscr{V})$ is not a commutative ring and hence does not satisfy the sufficient condition given in proposition 1.4.4 for the following to hold:*

$$p(\mathrm{T})q(\mathrm{T}) = pq(\mathrm{T})$$

but this property does hold for any two polynomials in $\mathbb{F}[x]$ and any linear transformation T, and as a result:

$$p(\mathrm{T}) = 0 \;\text{ and } p(x) | f(x) \Rightarrow f(\mathrm{T}) = 0$$

*However since the ring $\mathscr{L}(\mathscr{V}, \mathscr{W})$ does not admit universal inverses, the following does **not** hold:*

$$p(\mathrm{T})q(\mathrm{T}) = 0 \Rightarrow p(\mathrm{T}) = 0 \text{ or } q(\mathrm{T}) = 0$$

Proof. Let $p(x)$, $q(x)$ and $pq(x)$ be the following:

$$p(x) = \sum_{i=1}^{m} \alpha_i x^i$$

$$q(x) = \sum_{i=1}^{n} \beta_i x^i$$

$$pq(x) = \sum_{i=1}^{m+n} \gamma_i x^i$$

where:

$$\gamma_i = \sum_{j=0}^{i} \alpha_i \beta_{i-j}$$

We will get:

$$\begin{aligned} p(\mathrm{T})q(\mathrm{T}) &= \left(\sum_{i=0}^{m} \alpha_i \mathrm{I} \circ \mathrm{T}^i\right) \circ \left(\sum_{i=0}^{n} \beta_i \mathrm{I} \circ \mathrm{T}^i\right) \\ &= \sum_{k=1}^{m+n} \sum_{i=0}^{k} \alpha_i \mathrm{I} \circ \mathrm{T}^i \circ \beta_{k-i} \mathrm{I} \circ \mathrm{T}^{k-i} \end{aligned}$$

Now notice that although linear transformations do not commute through composition, but the term:

$$\alpha_i \mathrm{I} \circ \mathrm{T}^i \circ \beta_{k-i} \mathrm{I} \circ \mathrm{T}^{k-i}$$

is a special case in which all terms commute through composition. The whole term equals $\alpha_i \beta_{k-i} \mathrm{T}^k$ and thus we get:

$$p(\mathrm{T})q(\mathrm{T}) = \sum_{k=1}^{m+n} \sum_{i=0}^{k} \alpha_i \beta_{k-i} \mathrm{T}^k = \sum_{k=1}^{m+n} \left[\left(\sum_{i=0}^{k} \alpha_i \beta_{k-i} \right) \mathrm{T}^k \right] = \sum_{k=1}^{m+n} \gamma_k \mathrm{T}^k$$

and thus the proof is complete. □

8.1.1 Building blocks of canonical forms: Cyclic subspaces

Now we mention the basic construct of the rational canonical form: T-cyclic subspaces. We have seen in chapter 5.2 that T-invariant subspaces can provide us with valuable properties. But up to now we have no means of actually finding such subspaces. The idea of a T-cyclic subspace is to build the most trivial T-invariant subspace: pick an arbitrary vector u and include all vectors $\mathrm{T}^k u$ in the subspace:

Definition 8.1.1. Let $\mathrm{T}: \mathscr{V} \to \mathscr{V}$ be a linear transformation and $u \in \mathscr{V}$ be an arbitrary vector. We define the T-cyclic subspace induced by u, denoted by $\langle u \rangle_{\mathrm{T}}$, to be the following:

$$\langle u \rangle_{\mathrm{T}} = \mathrm{span}\{u, \mathrm{T}u, \mathrm{T}^2 u, \ldots\}$$

Notice that in the above sequence of vectors, as soon as the first $\mathrm{T}^m u$ is observed that does not increase the dimensionality of $\mathrm{span}\{u, \mathrm{T}u, \mathrm{T}^2 u, \ldots, \mathrm{T}^{m-1} u\}$, i.e. falls within the aforementioned subspace, all the future $\mathrm{T}^{t \geq m} u$ s will also fall in the same subspace, and thus we will not meet any

other "interesting" vector again. As a result if $k = \dim\langle u\rangle_{\mathrm{T}}$, then k is the smallest integer such that T^k is linearly dependent on $u, \mathrm{T}u, \mathrm{T}^2, \ldots, \mathrm{T}^{k-1}u$:

$$\mathrm{T}^k = \sum_{i=0}^{k-1} \alpha_i \mathrm{T}^i u$$

which can be seen as:

$$p(T)u = 0$$

for some polynomial $p(x) \in \mathbb{F}[x]$. Notice that obviously this $p(x)$ is the lowest degree such polynomial:

Definition 8.1.2. Let $\mathrm{T}\colon \mathscr{V} \to \mathscr{V}$ be a linear transformation and $u \in \mathscr{V}$ be an arbitrary vector. If $k = \dim\langle u\rangle_{\mathrm{T}}$ there exist $\{\alpha_i\}_{i=0}^{k-1}$ such that:

$$\mathrm{T}^k = \sum_{i=0}^{k-1} \alpha_i \mathrm{T}^i u$$

we define the u-**minimal polynomial**[1] of T to be the following:

$$m_{\mathrm{T}}^{(u)}(x) = x^k - \alpha_{k-1}x^{k-1} - \alpha_{k-2}x^{k-2} - \ldots - \alpha_1 x - \alpha_0$$

Obviously $m_{\mathrm{T}}^{(u)}(x)$ is the smallest degree polynomial that satisfies:

$$m_{\mathrm{T}}^{(u)}(\mathrm{T})u = 0$$

Proposition 8.1.2. *If $m_{\mathrm{T}}^{(u)}(x)$ is the u-minimal polynomial of* T *over $\langle u\rangle_{\mathrm{T}}$ then it will divide any polynomial $f(x) \in \mathbb{F}[x]$ that satisfies $f(\mathrm{T})u = 0$.*

Proof. Let $f(x)$ be as mentioned and $q(x)$ and $r(x)$ be the quotient and remainder of dividing it by $m_{\mathrm{T}}^{(u)}(x)$:

$$f(x) = q(x)m_{\mathrm{T}}^{(u)}(x) + r(x)$$

where $\deg(r) < \deg(m_{\mathrm{T}}^{(u)})$. We have:

$$f(\mathrm{T})u = q(\mathrm{T})m_{\mathrm{T}}^{(u)}(\mathrm{T})u + r(\mathrm{T})u$$

and since the left hand side is the zero vector we get:

$$r(\mathrm{T})u = 0$$

which is a contradiction since $\deg(r) < \deg(m_{\mathrm{T}}^{(u)})$. □

[1] In the literature this is usually referred to by the *order* and not the u-minimal polynomial. Order is a term borrowed from abstract algebra, and since we do not have an abstract algebraic viewpoint to the problem of decomposition of $\mathscr{V}$, we pick a more conveying term for this concept.

Corollary 8.1.1. Let $\mathscr{W}$ be the T-cyclic subspace $\langle u\rangle_{\mathrm{T}}$ for some u. Then since $\mathscr{W}$ is T-invariant, $\mathrm{T}_{\mathscr{W}}\colon \mathscr{W}\to\mathscr{W}$ is well-defined. If $k=\dim(\mathscr{W})=\deg(m_{\mathrm{T}}^{(u)})$ then $\mathscr{B}=\{u,\mathrm{T}u,\dots,\mathrm{T}^{k-1}u\}$ is a basis for $\mathscr{W}$ and the representation of $\mathrm{T}_{\mathscr{W}}$ in this representation is the following:

$$[\mathrm{T}_{\mathscr{W}}]_{\mathscr{B}}=\begin{bmatrix} 0 & & & & \dots & 0 & \alpha_0\\ 1 & 0 & & & \dots & 0 & \alpha_1\\ & 1 & 0 & & & & \vdots\\ & & 1 & 0 & & & \\ & & & & \ddots & & \\ & & & & & 0 & \alpha_{k-2}\\ & & & & & 1 & \alpha_{k-1}\end{bmatrix}$$

where α_i are the coefficients of $m_{\mathrm{T}}^{(u)}(x)$. This what is usually referred to by the *companion matrix of a polynomial* in the literature.

It can easily be seen that for any $v\in\langle u\rangle_{\mathrm{T}}$ there exists a polynomial $f(x)\in\mathbb{F}[x]$ such that:

$$v=f(\mathrm{T})u$$

and also for *any* polynomial $f(x)$ the vector $v=f(\mathrm{T})u$ lies in $\langle u\rangle_{\mathrm{T}}$ (in a very small special case this results in T-invariance of T-cyclic subspaces). Also notice that $m_{\mathrm{T}}^{(u)}(\mathrm{T})$ is zero all over $\langle u\rangle_{\mathrm{T}}$ since:

$$m_{\mathrm{T}}^{(u)}(\mathrm{T})v=m_{\mathrm{T}}^{(u)}(\mathrm{T})f(\mathrm{T})u=0$$

In fact $m_{\mathrm{T}}^{(u)}(x)$ is the smallest degree polynomial that has the property of being zero all over $\langle u\rangle_{\mathrm{T}}$ (this obviously follows from proposition 8.1.2).

Definition 8.1.3. As we just saw the u-minimal polynomial is such that $m_{\mathrm{T}}^{(u)}(\mathrm{T})$ is zero all over the subspace $\langle u\rangle_{\mathrm{T}}$, or equivalently $\langle u\rangle_{\mathrm{T}}\subseteq\mathrm{Ker}[m_{\mathrm{T}}^{(u)}(\mathrm{T})]$. In the same fashion as we defined the u-minimal polynomial, one can think of the $\mathscr{W}$-minimal polynomial $m_{\mathrm{T}}^{\mathscr{W}}(x)$ for any T-invariant subspace $\mathscr{W}$:

$$m_{\mathrm{T}}^{\mathscr{W}}(x)=\mathrm{l.c.m}_{u\in\mathscr{W}}\left[m_{\mathrm{T}}^{(u)}(x)\right]$$

where l.c.m denotes the least common multiple of polynomials. And again for any polynomial $f(x)$ that makes $f(\mathrm{T})$ be zero all over $\mathscr{W}$ (i.e. $\mathscr{W}\subseteq\mathrm{Ker}[f(\mathrm{T})]$) we will have $m_{\mathrm{T}}^{\mathscr{W}}(x)|f(x)$.

Definition 8.1.4. Getting back to the notion of $\mathscr{W}$-minimal polynomial of T that we just constructed, we now look at the special case of $\mathscr{W}=\mathscr{V}$. We call the $\mathscr{V}$-minimal polynomial of T its **minimal polynomial** $m_{\mathrm{T}}(x)$. Of course, as before, $m_{\mathrm{T}}(x)$ is the lowest degree polynomial that makes $m_{\mathrm{T}}(\mathrm{T})$ become zero all over its respective invariant subspace, and since this subspace is the whole of $\mathscr{V}$, this rewrites as the following:
"Being the **minimal polynomial** of $\mathrm{T}\colon\mathscr{V}\to\mathscr{V}$, $m_{\mathrm{T}}(x)$ is the lowest degree polynomial that makes $m_{\mathrm{T}}(\mathrm{T})$ equal to the zero transformation." Again we can easily prove that for any other polynomial $f(x)\in\mathbb{F}[x]$ making $f(\mathrm{T})$ become the zero transformation, we should have $m_{\mathrm{T}}(x)|f(x)$."

After the introduction of T-cyclic subspaces and $*$-minimal polynomials, there arise a few important questions:

1- Does for any T-invariant subspace $\mathscr{W}$ exist a vector u such that $\mathscr{W} = \langle u \rangle_{\mathrm{T}}$? In other words, are T-cyclic subspaces the only type of invariant subspaces? The answer is No! To see why, simply take $\mathscr{W} = \mathrm{Ker}[\mathrm{T}]$, which is obviously an invariant subspace of dimensionality $\eta(\mathrm{T})$. But there is no u in Ker[T] such that $\langle u \rangle_{\mathrm{T}}$ has dimensionality more than 1.

2- What if T is non-singular over $\mathscr{W}$? The answer is again no! To see this, just look at the identity transformation which is any sense the most "well behaved" transformation, and notice that for *any* vector $u \in \mathscr{V}$ the subspace $\langle u \rangle_{\mathrm{I}}$ has dimensionality 1.

3- Does u have a central role in $\langle u \rangle_{\mathrm{T}}$ or the T-cyclic subspace of any vector $v \in \langle u \rangle_{\mathrm{T}}$ would again yield the same subspace? The answer here is that u, together with some other special vectors in $\langle u \rangle_{\mathrm{T}}$, *does* have a central role in $\langle u \rangle_{\mathrm{T}}$, in the sense they are the only ones whose T-cyclic subspaces result in the whole subspace. Let's say $v = f(\mathrm{T})u$ for some polynomial $f(x) \in \mathbb{F}[x]$. The fact that

$$\langle v \rangle_{\mathrm{T}} \subseteq \langle u \rangle_{\mathrm{T}}$$

is obvious, from which it immediately follows that $m_{\mathrm{T}}^{(v)}(x) | m_{\mathrm{T}}^{(u)}(x)$. In fact we can prove the following easily:

Proposition 8.1.3. *Let* $\mathrm{T}: \mathscr{V} \to \mathscr{V}$ *be any linear transformation, and* $v = f(\mathrm{T})u$ *be any member of* $\langle u \rangle_{\mathrm{T}}$ *for some* $u \in \mathscr{V}$*. Then the* v*-minimal polynomial of* T *is related to the* u*-minimal polynomial of* T *in the following manner:*

$$m_{\mathrm{T}}^{(v)}(x) = \frac{m_{\mathrm{T}}^{(u)}(x)}{\mathrm{g.c.d}\big[m_{\mathrm{T}}^{(u)}(x), f(x)\big]}$$

An obvious consequence of the above is that:

$$\begin{aligned} \dim\langle v \rangle_{\mathrm{T}} &= \deg(m_{\mathrm{T}}^{(v)}) \\ &= \deg(m_{\mathrm{T}}^{(u)}) - \deg(\mathrm{g.c.d}\big[m_{\mathrm{T}}^{(u)}(x), f(x)\big]) \\ &= \dim\langle u \rangle_{\mathrm{T}} - \deg(\mathrm{g.c.d}\big[m_{\mathrm{T}}^{(u)}(x), f(x)\big]) \end{aligned}$$

And thus one can immediately see that:

Proposition 8.1.4. *Let* $\mathrm{T}: \mathscr{V} \to \mathscr{V}$ *be any linear transformation, and* $v = f(\mathrm{T})u$ *be any member of* $\langle u \rangle_{\mathrm{T}}$ *for some* $u \in \mathscr{V}$*. The* T*-cyclic subspace of* v*, which is by definition a subset of* $\langle u \rangle_{\mathrm{T}}$*, will coincide with the whole subspace* $\langle u \rangle_{\mathrm{T}}$ *if and only if* $f(x)$ *and* $m_{\mathrm{T}}^{(u)}(x)$ *are relatively prime:*

$$\mathrm{g.c.d}\big[m_{\mathrm{T}}^{(u)}, f\big] = 1$$

This result could have been derived in another fashion: It is easy to see that a necessary and sufficient condition for $\langle v \rangle_{\mathrm{T}}$ coinciding with $\langle u \rangle_{\mathrm{T}}$ is that $u \in \langle v \rangle_{\mathrm{T}}$. So the necessary and sufficient condition is that there exists a polynomial $p(x)$ such that:

$$p(\mathrm{T})v = p(\mathrm{T})f(\mathrm{T})u = u$$

which by proposition 8.1.2 happens if and only if we have:

$$m_{\mathrm{T}}^{(u)}(x) \big| p(x)f(x) - 1$$

which again, with some small manipulations, can be seen to hold if and only if $f(x)$ and $m_{\mathrm{T}}^{(u)}(x)$ are relatively prime.

As it can be seen moving through u-minimal to $\mathscr{W}$-minimal and to $\mathscr{V}$-minimal polynomials, we get from a completely vector dependent property of T to an innate property of T. Looking again at the perfect example of $\mathrm{T} = \mathrm{I}$, some extremely important observations can be made. The minimal polynomial of T is $x - 1$, and so is it for *any* $\mathscr{W}$-minimal polynomial, and *any* u-minimal polynomial of T. We know want to prove some results relating linearly independence (of vectors or of subspaces) to $*$-minimal polynomials.

Proposition 8.1.5. *Any polynomial $f(x) \in \mathbb{F}[x]$ induces a* T*-invariant subspace $\mathscr{W} = \mathrm{Ker}[f(\mathrm{T})]$, with respect to which the minimal polynomial of* T *divides (and is not necessarily equal to) $f(x)$:*

$$m_{\mathrm{T}}^{\mathrm{Ker}[f(\mathrm{T})]}(x) \Big| f(x)$$

On the other hand any T*-invariant subspace $\mathscr{W}$ induces a respective minimal polynomial of* T*, namely $f(x) = m_{\mathrm{T}}^{\mathscr{W}}(x)$, which when evaluated over* T*, i.e $f(\mathrm{T})$, the kernel of the resulting linear map, i.e $\mathrm{Ker}[f(\mathrm{T})]$, is a superset of (and not necessarily equal to) $\mathscr{W}$:*

$$\mathscr{W} \subseteq \mathrm{Ker}[m_{\mathrm{T}}^{\mathscr{W}}(\mathrm{T})]$$

Proposition 8.1.6. *Let $\mathscr{Y}$ and $\mathscr{W}$ be two* T*-invariant subspaces of $\mathscr{V}$. $\mathscr{Y}$ is a subspace of $\mathscr{W}$ if (and not only if) the $\mathscr{Y}$-polynomial of* T *divides the $\mathscr{W}$-polynomial:*

$$m_{\mathrm{T}}^{\mathscr{Y}}(x) | m_{\mathrm{T}}^{\mathscr{W}}(x) \;\Rightarrow\; \mathscr{Y} \subseteq \mathscr{W}$$

Proof. For any arbitrary vector $u \in \mathscr{Y}$ we, by definition, have:

$$m_{\mathrm{T}}^{\mathscr{Y}}(\mathrm{T})u = 0$$

Since $m_{\mathrm{T}}^{\mathscr{Y}}(x)$ divides $m_{\mathrm{T}}^{\mathscr{W}}(x)$ we get:

$$m_{\mathrm{T}}^{\mathscr{W}}(\mathrm{T})u = 0$$

□

Proposition 8.1.7. *Two* T*-invariant subspaces $\mathscr{W}_1$ and $\mathscr{W}_2$ are linearly independent (intersect only trivially) if (and not only if) their respective $\mathscr{W}_i$-minimal polynomials of* T *are relatively prime.*

Proof. Let $d(x)$ be the greatest common divisor of $m_{\mathrm{T}}^{\mathscr{W}_1}(x)$ and $m_{\mathrm{T}}^{\mathscr{W}_2}(x)$. Take the subspace:

$$\mathscr{Y} = \mathrm{Ker}[d(\mathrm{T})]$$

Now let $\mathscr{Z}$ be the intersection of $\mathscr{W}_1$ and $\mathscr{W}_2$. Obviously we have:

$$m_{\mathrm{T}}^{\mathscr{Z}}(x) | m_{\mathrm{T}}^{\mathscr{W}_1}(x)$$
$$m_{\mathrm{T}}^{\mathscr{Z}}(x) | m_{\mathrm{T}}^{\mathscr{W}_2}(x)$$

which results in:

$$m_{\mathrm{T}}^{\mathscr{Z}}(x) | d(x) = m_{\mathrm{T}}^{\mathscr{Y}}(x)$$

□

8.2 Stabilizing powers and superdiagonalization of nilpotent transformations

We remember from corollary 3.1.1 that for an endomorphism $\mathrm{T}: \mathscr{V} \to \mathscr{V}$ we have the following:

(i) $\mathrm{Ker}[\mathrm{T}] \subseteq \mathrm{Ker}[\mathrm{T} \circ \mathrm{T}] \subseteq \mathrm{Ker}[\mathrm{T} \circ \mathrm{T} \circ \mathrm{T}] \subseteq \ldots$.

(ii) $\mathrm{Im}[\mathrm{T}] \supseteq \mathrm{Im}[\mathrm{T} \circ \mathrm{T}] \supseteq \mathrm{Im}[\mathrm{T} \circ \mathrm{T} \circ \mathrm{T}] \supseteq \ldots$

It can easily be seen that, by regarding $\mathrm{Im}[\mathrm{T}^m]$ as a subset of $\mathscr{V}$, for any m we have:

$$\mathrm{Im}[\mathrm{T}^{m+1}] = \mathrm{TIm}[\mathrm{T}^m]$$

and that:

$$\mathrm{Ker}[\mathrm{T}^m] = \mathrm{TKer}[\mathrm{T}^{m+1}]$$

Obviously both these chains of subspaces would stop changing at some point. Assume:

$$\mathrm{Ker}[\mathrm{T}] \subset \mathrm{Ker}[\mathrm{T}^2] \subset \mathrm{Ker}[\mathrm{T}^3] \subset \ldots \subset \mathrm{Ker}[\mathrm{T}^k] = \mathrm{Ker}[\mathrm{T}^{k+1}]$$

We can easily see that $\mathrm{Ker}[\mathrm{T}^m]$ would stop changing *for ever* after k. We know that for all u:

$$\mathrm{T}^{k+1}u = 0 \Rightarrow \mathrm{T}^k u = 0$$

As a result for any u if $\mathrm{T}^{k+2}u = 0$ we write $\mathrm{T}^{k+1}(\mathrm{T}u) = 0$. Using the above rule for $v = \mathrm{T}u$, we get $\mathrm{T}^k v = 0$ which means $\mathrm{T}^{k+1}u = 0$, and therefore $\mathrm{Ker}[\mathrm{T}^{k+1}] = \mathrm{Ker}[\mathrm{T}^{k+2}]$. We also know that in both chains the drop (jump) in dimensionalities at each increment in the power of T should be in accord with the Rank-Nullity theorem. In other words for every i, since we have:

$$\eta(\mathrm{T}^i) + \rho(\mathrm{T}^i) = n$$

$$\eta(\mathrm{T}^{i+1}) + \rho(\mathrm{T}^{i+1}) = n$$

subtracting the two yields:

$$\eta(\mathrm{T}^{i+1}) - \eta(\mathrm{T}^i) = \rho(\mathrm{T}^i) - \rho(\mathrm{T}^{i+1})$$

As a result the moment the chain of kernels stops changing, so does the chain of images. Thus we have proved the following:

Proposition 8.2.1. *For any endomorphism* $\mathrm{T} : \mathscr{V} \to \mathscr{V}$ *there exists a unique positive integer* $k \leq \dim(\mathscr{V})$ *such that:*

$$\begin{array}{cccccccc} \mathrm{Ker}[\mathrm{T}] & \subset \mathrm{Ker}[\mathrm{T}^2] & \subset \mathrm{Ker}[\mathrm{T}^3] & \subset \ldots & \subset \mathrm{Ker}[\mathrm{T}^k] & = \mathrm{Ker}[\mathrm{T}^{k+1}] & = \mathrm{Ker}[\mathrm{T}^{k+2}] & = \ldots \\ \mathrm{Im}[\mathrm{T}] & \supset \mathrm{Im}[\mathrm{T}^2] & \supset \mathrm{Im}[\mathrm{T}^3] & \supset \ldots & \subset \mathrm{Im}[\mathrm{T}^k] & = \mathrm{Im}[\mathrm{T}^{k+1}] & = \mathrm{Im}[\mathrm{T}^{k+2}] & = \ldots \end{array}$$

Definition 8.2.1. We define the unique integer defined in 8.2.1 to be the **stabilizing power** of T. The fact that the stabilizing power can not exceed $\dim(\mathscr{V})$ is obvious by looking at the chain of kernels.

Proposition 8.2.2. *Let* $\mathrm{T}\colon \mathscr{V}\to\mathscr{V}$ *be a linear transformation and* k *be its stabilizing power. Then we have:*

$$\mathrm{Im}[\mathrm{T}^k]\cap\mathrm{Ker}[\mathrm{T}^k]=\{0\}$$

and hence we can write:

$$\mathscr{V}=\mathrm{Im}[\mathrm{T}^k]\oplus\mathrm{Ker}[\mathrm{T}^k]$$

Furthermore both $\mathrm{Im}[\mathrm{T}^k]$ *and* $\mathrm{Ker}[\mathrm{T}^k]$ *are* T*-invariant subspaces.*

Proof. Let $u\in\mathrm{Im}[\mathrm{T}^k]\cap\mathrm{Ker}[\mathrm{T}^k]$. It follows that $u=\mathrm{T}^m v$ for some v and also $\mathrm{T}^m u=0$, and thus $\mathrm{T}^{2k}v=0$. Since $\mathrm{Ker}[\mathrm{T}^{2k}]=\mathrm{Ker}[\mathrm{T}^k]$ it follows that $\mathrm{T}^k v=0$ and hence $u=0$. Now $\mathrm{Im}[\mathrm{T}^k],\mathrm{Ker}[\mathrm{T}^k]$ are two linearly independent subspaces of $\mathscr{V}$ whose dimensionalities by Rank-Nullity theorem add up to $\dim(\mathscr{V})$, and hence we can write:

$$\mathscr{V}=\mathrm{Im}[\mathrm{T}^k]\oplus\mathrm{Ker}[\mathrm{T}^k]$$

Also notice that by definition of the stabilizing power we have:

$$\mathrm{Im}[\mathrm{T}^k]=\mathrm{Im}[\mathrm{T}^{k+1}]=\mathrm{TIm}[\mathrm{T}^k]$$

which translates to $\mathrm{Im}[\mathrm{T}^k]$ being T-invariant. The case of $\mathrm{Ker}[\mathrm{T}^k]$ is different: $\mathrm{Ker}[\mathrm{T}^m]$ is T-invariant for *any* m. □

Corollary 8.2.1. If T is non-singular it can easily be seen that its stabilizing power is $k=0$.

Corollary 8.2.2. From the fact that the stabilizing power can not exceed $n=\dim(\mathscr{V})$ we can write in general that for any linear transformation $\mathrm{T}\colon\mathscr{V}\to\mathscr{V}$, the following is a T-invariant direct sum decomposition of $\mathscr{V}$:

$$\mathscr{V}=\mathrm{Im}[\mathrm{T}^n]\oplus\mathrm{Ker}[\mathrm{T}^n]$$

We now turn our attention to the special case where the chain of kernels goes on to swallow up the whole space $\mathscr{V}$, which is of extreme importance.

Definition 8.2.2. Let $\mathrm{T}:\mathscr{V}\to\mathscr{V}$ be a linear transformation and k be its stabilizing power. If $\eta(\mathrm{T}^k)=\dim(\mathscr{V})$ we say that T is **nilpotent** of degree k, or k-nilpotent. In other words a k-nilpotent transformation is one that for any $u\in\mathscr{V}$ we have $\mathrm{T}^k u=0$ but there do exist vectors such as v such that $\mathrm{T}^{k-1}v\neq 0$. Obviously the degree of nilpotency of a transformation can not exceed $\dim(\mathscr{V})$.

Corollary 8.2.3. Let $\mathrm{T}:\mathscr{V}\to\mathscr{V}$ be a linear transformation and k be its stabilizing power. Then $\mathrm{T}_{\mathrm{Im}[\mathrm{T}^k]}$ is non-singular and $\mathrm{T}_{\mathrm{Ker}[\mathrm{T}^k]}$ is k-nilpotent.

Proposition 8.2.3. *Let* $\mathrm{T}:\mathscr{V}\to\mathscr{V}$ *be a* k*-nilpotent linear transformation. For any* $u\in\mathscr{V}-\mathrm{Ker}[\mathrm{T}^{k-1}]$ *the following* k *vectors form a linearly independent set:*

$$\{u,\mathrm{T}u,\mathrm{T}^2u,\ldots,\mathrm{T}^{k-1}u\}$$

Proof. Let there be a vanishing linear combination of these vectors:

$$\sum_{i=0}^{k-1}\alpha_i\mathrm{T}^i u=0$$

Let t be the smallest integer for which α_t is not zero:

$$0 = \sum_{i=t}^{k-1} \alpha_i \mathrm{T}^i u = \sum_{i=0}^{k-t-1} \alpha_{t+i} \mathrm{T}^{t+i} u = \mathrm{T}^t \sum_{i=0}^{k-t-1} \alpha_{t+i} \mathrm{T}^i u = \alpha_t \mathrm{T}^t u + \sum_{i=1}^{k-t-1} \alpha_{t+i} \mathrm{T}^{t+i} u$$

thus defining β_i to be $-\frac{\alpha_i}{\alpha_t}$ we get:

$$\mathrm{T}^t u = \sum_{i=1}^{k-t-1} \beta_{t+i} \mathrm{T}^{t+i} u$$

we now apply T^{k-t-1} to both sides:

$$\mathrm{T}^{k-1} u = \sum_{i=1}^{k-t-1} \beta_{t+i} \mathrm{T}^{k-1+i} u$$

Now in all the summands of the left hand side the power of T is at least k and hence they all vanish to zero, but the right hand side is non-zero by the choice of u which is a contradiction. □

Notice that from this result if the degree of nilpotency of T coincides with $\dim(\mathscr{V})$ (i.e. $k = n$) then one can readily use the basis $\mathscr{B} = \{\mathrm{T}^{n-1}u, \mathrm{T}^{k-2}u, \dots \mathrm{T}u, u\}$ and $[\mathrm{T}]_{\mathscr{B}}$ would be superdiagonal matrix[2]:

$$[\mathrm{T}]_{\mathscr{B}} = \begin{bmatrix} 0 & 1 & & & & \\ & 0 & 1 & & 0 & \\ & & \ddots & \ddots & & \\ & & & \ddots & 1 & \\ & 0 & & & & 1 \\ & & & & & 0 \end{bmatrix}$$

But what if $k < n$ we can not get such a superdiagonalization immediately, but the thing to notice is that the fact that $k < n$ means T is a *simpler* transformation than an n-nilpotent transformation. We would guess that T could still be superdiagonalized except for some of the 1s being struck out. In fact this intuition leads to a true proposition:

Proposition 8.2.4. *Let* $\mathrm{T}\colon \mathscr{V} \to \mathscr{V}$ *be a* k*-nilpotent linear transformation. There exist a* T*-cyclic direct sum decomposition* $\mathscr{V} = \bigoplus_{i=1}^{p} \langle u_i \rangle_{\mathrm{T}}$ *such that for any* i, $\mathrm{T}_{\langle u_i \rangle_{\mathrm{T}}}$ *is exactly* d_i*-nilpotent, where* $d_i = \dim(\langle u_i \rangle_{\mathrm{T}})$*. It would immediately follow that* T *has a superdiagonal block diagonalization with* p *of the ones being struck out.*

Proof. We prove our claim with induction, and follow the proof given by John Armstrong. We pick an arbitrary vector $u_1 = u \in \mathscr{V} - \mathrm{Ker}[\mathrm{T}^{k-1}]$ and pick the first component to be $\langle u \rangle_{\mathrm{T}}$. Obviously $\langle u \rangle_{\mathrm{T}}$ is a k-dimensional T-invariant subspace of $\mathscr{V}$ and $\mathrm{T}_{\langle u \rangle_{\mathrm{T}}}$ is obviously k-nilpotent. We now will introduce, in a way, a T-invariant subspace $\mathscr{W}^{(1)}$, linearly independent of $\langle u \rangle_{\mathrm{T}}$ such that:

$$\mathscr{V} = \langle u \rangle_{\mathrm{T}} \oplus \mathscr{W}^{(1)}$$

[2]It would be more consistent if we chose the reverse order for the basis so that the matrix becomes the transpose of this matrix. That would have been more in accord with the structure of companion matrices we mentioned earlier, but here I have followed the unanimous way textbooks superdiagonalize nilpotent matrices

and then decompose $\mathscr{W}^{(1)}$ in the same fashion (this is not where we are performing induction). We construct $\mathscr{W}^{(1)}$ inductively: Note that by the mere fact that $\mathscr{W}^{(1)} \subseteq \mathscr{V}$, T is nilpotent on $\mathscr{W}^{(1)}$ with a degree at most k.
First notice that u is some where outside $\mathrm{Ker}[\mathrm{T}^{k-1}]$. As a result $\mathrm{Ker}[\mathrm{T}^{k-1}] \oplus \mathrm{span}\{u\}$ is direct, but does not necessarily cover all $\mathscr{V}$. Define $\mathscr{W}_{(0)}$ to be such that:

$$\mathscr{V} = \mathrm{Ker}[\mathrm{T}^{k-1}] \oplus \mathrm{span}\{u\} \oplus \mathscr{W}_{(0)}$$

Obviously $\mathscr{W}_{(0)}$ too lies completely in $\mathscr{V} - \mathrm{Ker}[\mathrm{T}^{k-1}]$. Now we use the induction hypothesis for the confinement of T to the subspace $\mathrm{Ker}[\mathrm{T}^{k-1}]$, which is a $(k-1)$-nilpotent endomorphism. The vector $\mathrm{T}u$ plays the same role for $\mathrm{T}_{\mathrm{Ker}[\mathrm{T}^{k-1}]}$ that u plays for T. There, thus, exists a T-invariant subspace $\mathscr{W}_{(1)} \subseteq \mathrm{Ker}[\mathrm{T}^{k-1}]$, linearly independent of $\mathrm{span}\{\mathrm{T}u, \mathrm{T}^2u, \ldots, \mathrm{T}^{k-1}u\}$, such that:

$$\mathrm{Ker}[\mathrm{T}^{k-1}] = \mathrm{span}\{\mathrm{T}u, \mathrm{T}^2u, \ldots, \mathrm{T}^{k-1}u\} \oplus \mathscr{W}_{(1)}$$

Notice that since $\mathscr{W}_{(1)}$ lies completely in $\mathrm{Ker}[\mathrm{T}^{k-1}]$, it follows from its linear independence of

$$\mathrm{span}\{\mathrm{T}u, \mathrm{T}^2u, \ldots, \mathrm{T}^{k-1}u\}$$

that it is as well linearly independent of $\langle u \rangle_{\mathrm{T}}$. We now choose the desired $\mathscr{W}^{(1)}$ to be the following:

$$\mathscr{W}^{(1)} = \mathscr{W}_{(0)} \oplus \mathscr{W}_{(1)}$$

This is depicted in the figure below. We have to show that $\mathscr{W}^{(1)}$ intersects $\langle u \rangle_{\mathrm{T}}$ only trivially and

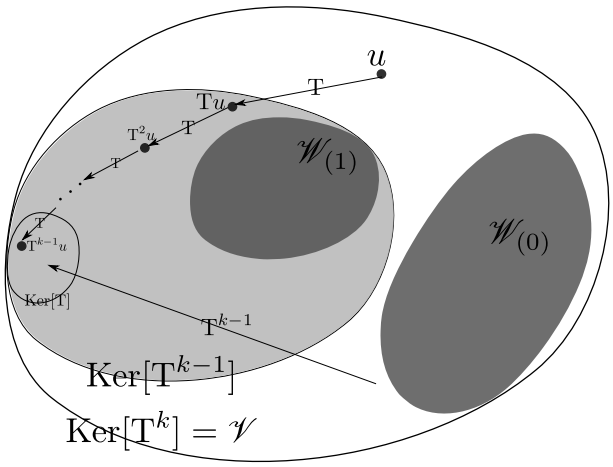

together with $\langle u \rangle_{\mathrm{T}}$, it builds the whole $\mathscr{V}$. We already know that:

$$\begin{aligned} \mathscr{V} = \mathrm{Ker}[\mathrm{T}^{k-1}] \oplus \mathrm{span}\{u\} \oplus \mathscr{W}_{(0)} &= \Big(\mathrm{span}\{\mathrm{T}u, \mathrm{T}^2u, \ldots, \mathrm{T}^{k-1}u\} \oplus \mathscr{W}_{(1)}\Big) + \mathrm{span}\{u\} + \mathscr{W}_{(0)} \\ &= \mathrm{span}\{u, \mathrm{T}u, \mathrm{T}^2u, \ldots, \mathrm{T}^{k-1}u\} + \mathscr{W}_{(1)} + \mathscr{W}_{(0)} \\ &= \langle u \rangle_{\mathrm{T}} + \mathscr{W}^{(1)} \end{aligned}$$

So we just need to prove that $\mathscr{W}^{(1)} \cap \langle u \rangle_{\mathrm{T}} = \{0\}$. Assume there exists $u_1 \in \langle u \rangle_{\mathrm{T}}$, $u_{(0)} \in \mathscr{W}_{(0)}$, and $u_{(1)} \in \mathscr{W}_{(1)}$ such that:

$$u_1 = u_{(0)} + u_{(1)}$$

where we have:

$$u_1 = \alpha_0 u + \alpha_1 \mathrm{T}u + \ldots + \alpha_{k-1}\mathrm{T}^{k-1}u$$

The fact that this can not happen unless these are zero would follow by applying T^{k-1} to both sides:

$$\mathrm{T}^{k-1}u_{(1)} + \mathrm{T}^{k-1}u_{(0)} = \mathrm{T}^{k-1}u_1 = \alpha_0\mathrm{T}^{k-1}u$$

The left most term is obviously zero, thus we get:

$$\mathrm{T}^{k-1}(u_{(0)} - \alpha u) = 0$$

which means $u_{(0)} - \alpha u \in \mathrm{Ker}[\mathrm{T}^{k-1}]$ which is a contradiction since we chose $\mathscr{W}_{(0)}$ to be such that:

$$\mathscr{V} = \mathrm{Ker}[\mathrm{T}^{k-1}] \oplus \mathrm{span}\{u\} \oplus \mathscr{W}_{(0)}$$

and hence the proof is complete. □

8.3 Rational Canonical Form

We now turn back to where we left $*$-minimal polynomials and T-cyclically irreducible direct sum decompositions. As we have probably created the impression by now, finding a T-*cyclic d.s.d.* for the whole of $\mathscr{V}$ is an alluring temptation. Firstly because the matrix representation of T over T-cyclic subspaces is, as we mentioned, fairly simple (with lots and lots of zero, as we always love it!) and thus the matrix representation of T would become block diagonal with the diagonal blocks being companion matrices. Pondering upon this idea, an immediate question would be that "when would a T-cyclic subspace $\langle u\rangle_{\mathrm{T}}$ be such that we can not further decompose it?". The answer would be easy: If $\langle u\rangle_{\mathrm{T}}$ is such that we have:

$$\langle u\rangle_{\mathrm{T}} = \langle v_1\rangle_{\mathrm{T}} \oplus \langle v_2\rangle_{\mathrm{T}}$$

it would follow that:

$$m_{\mathrm{T}}^{(u)}(x) = \text{l.c.m}\left[m_{\mathrm{T}}^{(v_1)}(x), m_{\mathrm{T}}^{(v_2)}(x)\right]$$

As a result a *sufficient* condition for $\langle u\rangle_{\mathrm{T}}$ to be T-*cyclically undecomposable* is that the u-minimal polynomial of T is a prime polynomial. But would that be a necessary condition as well? We will see later that something quite similar is a necessary condition.
Let's first rewrite: *if* $\mathscr{V}$ has a T-cyclic d.s.d as the following:

$$\mathscr{V} = \langle u_1\rangle_{\mathrm{T}} \oplus \langle u_2\rangle_{\mathrm{T}} \oplus \ldots \oplus \langle u_s\rangle_{\mathrm{T}}$$

(which is obviously a T-invariant d.s.d.) it will immediately follow that:

$$m_{\mathrm{T}}(x) = \text{l.c.m}\left[m_{\mathrm{T}}^{(u_1)}(x), m_{\mathrm{T}}^{(u_2)}(x), \ldots, m_{\mathrm{T}}^{(u_t)}(x)\right]$$

and since we are looking for undecomposable decompositions, *if* we manage to find u_i such that $m_{\mathrm{T}}^{(u_t)}(x)$s are *distinct* prime polynomials it would follow that:

$$m_{\mathrm{T}}(x) = m_{\mathrm{T}}^{(u_1)}m_{\mathrm{T}}^{(u_2)}\ldots m_{\mathrm{T}}^{(u_t)}(x)$$

As a result *if* there exists an irreducible decomposition it seems that it must have something to do with factorization of the minimal polynomial of T. But as the reader might have noticed, we swoop through a serious twist here: the factorization of $m_{\mathrm{T}}(x)$ to prime factors does not simply result in distinct prime polynomials, but in *powers* of distinct primes; and we have not discussed T-cyclic irreducibility of $\langle u\rangle_{\mathrm{T}}$ if the u-minimal polynomial is not a prime factor in $\mathbb{F}[x]$ but a power of a prime!

We now will formally tackle the problem we just verbally described. We first make the observation that if the $\mathscr{W}$-minimal polynomial of T is a power of a prime, for any T-invariant subspace $\mathscr{Y}$ of $\mathscr{W}$, the $\mathscr{Y}$-minimal polynomial of T has to be a power of the same prime factor. The first thing we will show here, is that any such subspace as $\mathscr{W}$ has a T-cyclically *irreducible* T-cyclic d.s.d that is also *unique* in a sense.

Definition 8.3.1. Let $\mathrm{T}\colon \mathscr{V}\to\mathscr{V}$ be a linear transformation and $\mathscr{W}$ be a subspace of $\mathscr{V}$. We call the following a T**-cyclic direct sum decomposition** of $\mathscr{W}$:

$$\mathscr{W} = \langle v_1\rangle_{\mathrm{T}} \oplus \langle v_2\rangle_{\mathrm{T}} \oplus \ldots \oplus \langle v_t\rangle_{\mathrm{T}}$$

Obviously this is a T-invariant direct sum decomposition, and furthermore $\mathscr{W}$ does not get to have such a decomposition unless it, itself, is T-invariant.

Definition 8.3.2. We call a T-invariant subspace $\mathscr{W}$ of $\mathscr{V}$ T**-cyclically undecomposable** if one can not find a T-cyclic direct sum decomposition for $\mathscr{W}$ to subspaces of strictly smaller dimensionalities.

Definition 8.3.3. Let $\mathrm{T}\colon \mathscr{V}\to\mathscr{V}$ be a linear transformation and $\mathscr{W}$ be a subspace of $\mathscr{V}$. We call the following a T-cyclic **irreducible** direct sum decomposition of $\mathscr{W}$:

$$\mathscr{W} = \langle v_1\rangle_{\mathrm{T}} \oplus \langle v_2\rangle_{\mathrm{T}} \oplus \ldots \oplus \langle v_t\rangle_{\mathrm{T}}$$

if all the subspaces $\langle v_i\rangle_{\mathrm{T}}$ are T-cyclically undecomposable.

Remark. Notice that a decomposition being reducible, does not result in the fact that there is no other decomposition with smaller dimensionality subspaces. It just means that *this* specific decomposition can not be further decomposed to T-cyclic subspaces.

We now use what we did in proposition 8.2.4, which was the case where the minimal polynomial of T was actually $m_{\mathrm{T}}(x) = x^k$, and prove that if the $\mathscr{W}$-minimal polynomial of T over is a power of a prime, $\mathscr{W}$ will decompose irreducibly to a T-cyclic direct sum.

Proposition 8.3.1. *Let* $\mathrm{T}\colon \mathscr{V}\to\mathscr{V}$ *be a linear transformation and let* $\mathscr{W}$ *be a* T*-invariant subspace such that the* $\mathscr{W}$*-minimal polynomial of* T *is* $p(x)^a$ *for some prime polynomial* $p(x)\in\mathbb{F}$ *and some integer* a*. There exists an irreducible* T*-cyclic direct sum decomposition* $\mathscr{W} = \bigoplus_{i=1}^{r}\langle u_i\rangle_{\mathrm{T}}$*. The* $\langle u_i\rangle_{\mathrm{T}}$*-minimal polynomial of* T *is obviously* $p(x)^{s_i}$ *for some integers* s_i*.*

Proof. We prove our claim by simultaneous induction on the power a of the prime factor, and the dimensionality of ????? and follow the same pattern as in the proof of proposition 8.2.4, except for instead of $\mathrm{Ker}[\mathrm{T}^m]$s we will be talking about $\mathrm{Ker}[p(\mathrm{T})^m]$s. We first define the degree of $p(x)$ to be d, and immediately observe that since $p(x)^a$ is the $\mathscr{W}$-minimal polynomial of T, then:

$$\dim(\mathscr{W}) = ad$$

We pick an arbitrary vector:

$$u \in \mathscr{W} - \mathrm{Ker}[p(\mathrm{T})^{a-1}]$$

and name its T-cyclic subspace $\mathscr{W}_1$. Since $p(\mathrm{T})^a = 0$ all over $\mathscr{W}_1 = \langle u_1 \rangle_{\mathrm{T}}$, it immediately follows (since $p(x)$ is prime) that the $\mathscr{W}_1$-minimal polynomial of T is $p(x)^{s_1}$ for some integer $s_1 \leq a$. If $s_1 = a$ since the dimensionality of $\mathscr{W}_1$ which is $s_1 d$ is equal to the dimensionality of $\mathscr{W}$ which is ad, then $\mathscr{W} = \mathscr{W}_1$. We now will introduce, in a way, subspace $\mathscr{W}^{(1)}$, linearly independent of $\langle u \rangle_{\mathrm{T}}$ such that:

$$\mathscr{W} = \langle u \rangle_{\mathrm{T}} \oplus \mathscr{W}^{(1)}$$

and then since $p(\mathrm{T})^a = 0$ all over $\mathscr{W}^{(1)}$ as well, and thus the $\mathscr{W}^{(1)}$-minimal polynomial of T is $p(x)^b$ for some integer b, we will continue in the same fashion (this is not where we are performing induction). We prove the existence of $\mathscr{W}^{(1)}$ by induction:
First notice that u is somewhere outside $\mathrm{Ker}[p(\mathrm{T})^{a-1}]$. Also notice that since $p(x)^a$ is the *smallest degree* polynomial that makes $p(\mathrm{T})u$ get to zero, all the following vectors all completely lie outside of $\mathrm{Ker}[p(\mathrm{T})^{a-1}]$:

$$u, \mathrm{T}u, \mathrm{T}^2 u, \ldots, \mathrm{T}^{d-1} u$$

where $d = \deg p(x)$, since for all these vectors $\mathrm{T}^i u$ with $i < d$, we have:

$$\deg(x^i p(x)^{a-1}) = ad + a - d < ad = \deg(m_{\mathrm{T}}^{(u)})$$

As a result $\mathrm{Ker}[p(\mathrm{T})^{a-1}] \oplus \mathrm{span}\{u, \mathrm{T}u, \ldots, \mathrm{T}^{d-1}u\}$ is direct, but does not necessarily cover all $\mathscr{W}$. Define $\mathscr{W}_{(0)}$ to be such that:

$$\mathscr{W} = \mathrm{Ker}[p(\mathrm{T})^{a-1}] \oplus \mathrm{span}\{u, \mathrm{T}u, \ldots, \mathrm{T}^{d-1}u\} \oplus \mathscr{W}_{(0)}$$

Obviously $\mathscr{W}_{(0)}$ too lies completely in $\mathscr{W} - \mathrm{Ker}[p(\mathrm{T})^{a-1}]$. Now we use the induction hypothesis for the confinement of T to the subspace $\mathrm{Ker}[p(\mathrm{T})^{a-1}]$, which has a subspace minimal polynomial of again a power of $p(x)$. The vector $p(\mathrm{T})u$ plays the same role for $\mathrm{T}_{\mathrm{Ker}[p(\mathrm{T})^{a-1}]}$ that u plays for T. It easy to see that the $p(\mathrm{T})u$-minimal polynomial of T is $p(x)^{a-1}$. There, by induction hypothesis, exists a T-invariant subspace $\mathscr{W}_{(1)} \subseteq \mathrm{Ker}[p(\mathrm{T})^{a-1}]$, linearly independent of $\langle p(\mathrm{T})u \rangle_{\mathrm{T}}$, such that:

$$\mathrm{Ker}[p(\mathrm{T})^{a-1}] = \langle p(\mathrm{T})u \rangle_{\mathrm{T}} \oplus \mathscr{W}_{(1)}$$

Notice that since $\mathscr{W}_{(1)}$ lies completely in $\mathrm{Ker}[p(\mathrm{T})^{a-1}]$, it follows from its linear independence of $\langle p(\mathrm{T})u \rangle_{\mathrm{T}}$ that it is as well linearly independent of $\langle u \rangle_{\mathrm{T}}$. We now choose the desired $\mathscr{W}^{(1)}$ to be the following:

$$\mathscr{W}^{(1)} = \mathscr{W}_{(0)} \oplus \mathscr{W}_{(1)}$$

We have to show that $\mathscr{W}^{(1)}$ intersects $\langle u \rangle_{\mathrm{T}}$ only trivially and together with $\langle u \rangle_{\mathrm{T}}$, it builds the whole $\mathscr{W}$. We already know that:

$$\begin{aligned}
\mathscr{W} &= \mathrm{Ker}[p(\mathrm{T})^{a-1}] \oplus \mathrm{span}\{u, \mathrm{T}u, \ldots, \mathrm{T}^{d-1}u\} \oplus \mathscr{W}_{(0)} \\
&= \left(\langle p(\mathrm{T})u \rangle_{\mathrm{T}} \oplus \mathscr{W}_{(1)}\right) + \mathrm{span}\{u, \mathrm{T}u, \ldots, \mathrm{T}^{d-1}u\} + \mathscr{W}_{(0)} \\
&= \left(\langle p(\mathrm{T})u \rangle_{\mathrm{T}} + \mathrm{span}\{u, \mathrm{T}u, \ldots, \mathrm{T}^{d-1}u\}\right) + \mathscr{W}_{(1)} + \mathscr{W}_{(0)} \\
&= \left(\langle p(\mathrm{T})u \rangle_{\mathrm{T}} + \mathrm{span}\{u, \mathrm{T}u, \ldots, \mathrm{T}^{d-1}u\}\right) + \mathscr{W}^{(1)}
\end{aligned}$$

It is fairly easy to see that:

$$\langle p(\mathrm{T})u\rangle_{\mathrm{T}} + \mathrm{span}\{u, \mathrm{T}u, \ldots, \mathrm{T}^{d-1}u\} = \langle u\rangle_{\mathrm{T}}$$

and thus we have got:

$$\mathscr{W} = \langle u\rangle_{\mathrm{T}} + \mathscr{W}^{(1)}$$

We will now prove that $\mathscr{W}^{(1)} \cap \langle u\rangle_{\mathrm{T}} = \{0\}$. Assume there exists $u_1 \in \langle u\rangle_{\mathrm{T}}$, $u_{(0)} \in \mathscr{W}_{(0)}$, and $u_{(1)} \in \mathscr{W}_{(1)}$ such that:

$$u_1 = u_{(0)} + u_{(1)}$$

where we have:

$$u_1 = f(\mathrm{T})u$$

for some polynomial $f(x)$. The fact that this can not happen unless these are zero would follow by applying $p(\mathrm{T})^{a-1}$ to both sides:

$$p(\mathrm{T})^{a-1}u_{(1)} + p(\mathrm{T})^{a-1}u_{(0)} = p(\mathrm{T})^{a-1}u_1 = p(\mathrm{T})^{a-1}f(\mathrm{T})u$$

The leftmost term is obviously zero, and also notice that if we divide $f(x)$ by $p(x)$:

$$f(x) = p(x)q(x) + r(x)$$

for some polynomial $r(x)$ with smaller degree than $p(x)$ we will get:

$$p(\mathrm{T})^{a-1}(u_{(0)} - r(\mathrm{T})u) = 0$$

which means $u_{(0)} - r(\mathrm{T})u \in \mathrm{Ker}[p(\mathrm{T})^{a-1}]$. Now notice that whatever $r(x)$ is, since it has smaller degree than d we get:

$$r(\mathrm{T})u \in \mathrm{span}\{u, \mathrm{T}u, \ldots, \mathrm{T}^{d-1}u\}$$

which is where the contradiction lies, since we chose $\mathscr{W}_{(0)}$ to be such that:

$$\mathscr{W} = \mathrm{Ker}[p(\mathrm{T})^{a-1}] \oplus \mathrm{span}\{u, \mathrm{T}u, \ldots, \mathrm{T}^{d-1}u\} \oplus \mathscr{W}_{(0)}$$

and we have now got a member of $\mathscr{W}_{(0)}$ being added to a member of $\mathrm{span}\{u, \mathrm{T}u, \ldots, \mathrm{T}^{d-1}u\}$ and then the result falling inside $\mathrm{Ker}[p(\mathrm{T})^{a-1}]$. □

Proposition 8.3.2. *Let* $\mathrm{T}\colon \mathscr{V} \to \mathscr{V}$ *be a linear transformation and let* $\mathscr{W}$ *be a* T*-invariant subspace such that the* $\mathscr{W}$*-minimal polynomial of* T *is* $p(x)^a$ *for some prime polynomial* $p(x) \in \mathbb{F}$ *and some integer* a*. The following* T*-cyclic direct sum decomposition:*

$$\mathscr{W} = \langle u_1\rangle_{\mathrm{T}} \oplus \langle u_2\rangle_{\mathrm{T}} \oplus \ldots \oplus \langle u_r\rangle_{\mathrm{T}}$$

the existence of which is ensured by 8.3.1, is unique in the sense that r *is unique and the degrees of* $m_{\mathrm{T}}^{(u_i)}(x)$ *are unique, up to a rearrangement.*

Proof. We here mention the neat proof given by [Cur84]. Let s_i be the same as in proposition 8.3.1. Let's investigate the dimensionality of $\text{Ker}[p(\text{T}_{\mathscr{W}})] \subseteq \mathscr{W}$. For any vector v in there according to the T-cyclic decomposition we have:

$$v = f_1(\text{T})u_1 + \ldots + f_r(\text{T})u_r$$

and hence:

$$0 = p(\text{T})v = p(\text{T})f_1(\text{T})u_1 + \ldots + p(\text{T})f_r(\text{T})u_r$$

which by assumption of *direct sum decomposition* means that for every $i \leq r$ we have:

$$p(\text{T})f_i(\text{T})u_i = 0$$

which immediately results in the fact that $p(x)^{s_i-1}|f_i(x)$ and thus we conclude that for any $v \in \text{Ker}[p(\text{T}_{\mathscr{W}})]$, its corresponding $f_i(x)$ are multiples of $p(x)^{s_i-1}$. Thus we can write:

$$\text{Ker}[p(\text{T}_{\mathscr{W}})] = \langle p(\text{T})^{s_1-1}u_1\rangle_{\text{T}} \oplus \langle p(\text{T})^{s_2-1}u_2\rangle_{\text{T}} \oplus \ldots \oplus \langle p(\text{T})^{s_r-1}u_r\rangle_{\text{T}}$$

Taking dimensions of both sides we get:

$$\eta[p(\text{T}_{\mathscr{W}})] = r\deg(p)$$

and hence r is unique and equal to:

$$r = \frac{\eta[p(\text{T}_{\mathscr{W}})]}{\deg(p)}$$

Now let us investigate the dimensionality of $\text{Ker}[p(\text{T}_{\mathscr{W}})^2]$. In the same fashion we conclude that $p(x)^{s_i-1}|f_i(x)$. Now if the number of s_is that are equal to 1 is x_1, then for those is, the requirement that $p(x)^{s_i-1}|f_i(x)$ is satisfied for any $f(x)$. As a result we get:

$$x_1\deg(p) + 2(r - x_1)\deg(p) = \eta[p(\text{T}_{\mathscr{W}})^2]$$

and hence x_1 is unique and equal to:

$$x_1 = -\frac{\eta[p(\text{T}_{\mathscr{W}})^2]}{\deg(p)} + 2r$$

following in the same fashion, investigating the dimensionality of $\text{Ker}[p(\text{T}_{\mathscr{W}})^i]$ we will prove that the unordered content of $\langle s_1, \ldots, s_r\rangle$ is unique. □

And finally the last proposition that finishes the treatise of rational canonical forms is the following:

Proposition 8.3.3. *Let* $\text{T}\colon \mathscr{V} \to \mathscr{V}$ *be a linear transformation with minimal polynomial* $m_{\text{T}}(x)$ *that has the following prime factorization:*

$$m_{\text{T}} = p_1^{e_1}p_2^{e_2}\ldots p_s^{e_s}$$

There exists an invariant direct sum decomposition of T *to* s *subspaces of* $\mathscr{V}$*, namely* $\mathscr{W}_1, \ldots, \mathscr{W}_s$ *such that the* $\mathscr{W}_i$*-minimal polynomial of* T *is a power of* $p_i(x)$ *(this power is not equal to* e_i*), and that:*

Proof. This proof is again from [Cur84] with some simplifications. We simply pick $\mathscr{W}_i$ to be the following:

$$\mathscr{W}_i = \mathrm{Ker}[p_i(\mathrm{T})^{e_i}]$$

Obviously the $\mathscr{W}_i$-minimal polynomial of T has to divide $p_i(x)^{e_i}$, and hence is $p_i(x)^{f_i}$. For all i, we build the polynomial $q_i(x)$ to be exactly the same as $m_{\mathrm{T}}(x)$ except for all the powers of $p_i(x)^{e_i}$ struck out of the factorization. Formally:

$$q_i(x) = \prod_{j \neq i} p_j(x)^{e_j}$$

The great thing about $q_i(x)$ is that when we apply $p_i(\mathrm{T})^{e_i}$ to the vector $q_i(\mathrm{T})u$ for any $u \in \mathscr{V}$, we will get $p_i(\mathrm{T})q_i(\mathrm{T})u = m_{\mathrm{T}}(\mathrm{T})u = 0$, and thus $q_i(\mathrm{T})$ sends all the vectors in $\mathscr{V}$ to the subspace $\mathscr{W}_i$ (it is not a projection though). Using this tool, we can easily observe that the $\mathscr{W}_i$ are linearly independent. Let there be $u_i \in \mathscr{W}_i$ such that:

$$u_1 + u_2 + \ldots + u_s = 0$$

Now for any i applying $q_i(\mathrm{T})$ to this sum the only survivor of the left hand side would be $q_i(\mathrm{T})u_i$, since for all the other u_j, $q_i(\mathrm{T})$ is a multiple of the respective u_i-minimal polynomial (which is a power of $p_j(x)^{e_j}$). Since for all i we have $q_i(\mathrm{T})u_i = 0$ we get that for all i the u_i-minimal polynomial which is $p_i(x)^{g_i}$ for some g_i should divide $q_i(x)$ which can not happen, and thus all u_i should be zero.

Now we will prove that:

$$\mathscr{W}_1 + \mathscr{W}_2 + \ldots + \mathscr{W}_s = \mathscr{V}$$

and finish the proof. To prove this we are going to use Bézout's Theorem. It can easily be seen that:

$$\text{g.c.d}\,[q_1(x), \ldots, q_s(x)] = 1$$

where 1 means the 1 polynomial (i.e. $l(x) = 1 + 0x + 0x^2 + \ldots$). By Bézout's we know that there exists polynomials $a_1(x), \ldots, a_s(x)$ such that:

$$a_1(x)q_1(x) + \ldots + a_s(x)q_s(x) = 1$$

Now pick an arbitrary vector $u \in \mathscr{V}$ and notice that:

$$a_1(\mathrm{T})q_1(\mathrm{T})u + \ldots + a_s(\mathrm{T})q_s(\mathrm{T})u = u$$

As we mentioned before, for every i we know that $q_i(\mathrm{T})u$ lies inside $\mathscr{W}_i$ and since $\mathscr{W}_i$ is obviously T-invariant, there lies also $a_i(\mathrm{T})q_i(\mathrm{T})u$. So for any vector $u \in \mathscr{V}$ we have found vectors $v_i = a_i(\mathrm{T})q_i(\mathrm{T})u \in \mathscr{W}_i$ such that $u = \sum v_i$ and the proof is complete. □

Remark. It is easy to see that with the choice of $q_i(x)$ and $a_i(x)$ as in the proof, $a_i(\mathrm{T})q_i(\mathrm{T})$ is the projection operator on $\mathscr{W}_i$.

Corollary 8.3.1. It can easily be seen that the only way a T-invariant subspace can be T-cyclically undecomposable is when the $\mathscr{W}$-minimal polynomial of T does not have more than prime factor, and hence is a power of a prime.

The result of the propositions of this section is the following:

Proposition 8.3.4 (Rational Canonical Form)**.** *Let* $\mathrm{T} : \mathscr{V} \to \mathscr{V}$ *be a linear transformation with minimal polynomial* $m_{\mathrm{T}}(x)$ *that has the following prime factorization:*

$$m_{\mathrm{T}} = p_1^{e_1} p_2^{e_2} \dots p_s^{e_s}$$

where we define d_i *to be the degree of* $p_i(x)$*. There exists a* T*-cyclic direct sum decomposition for* $\mathscr{V}$*:*

$$\mathscr{V} = \langle u_1 \rangle_{\mathrm{T}} \oplus \langle u_2 \rangle_{\mathrm{T}} \oplus \dots \oplus \langle u_k \rangle_{\mathrm{T}}$$

such that the u_i*-minimal polynomial of* T *is, for all* i*, the following:*

$$m_{\mathrm{T}}^{(u_i)} = p_{\pi_i}(x)^{a_i}$$

for some:

$$1 \leq \pi_i \leq s$$

$$1 \leq a_i \leq e_i$$

Furthermore 1) the sequence $\{\pi_1, \pi_2, \dots, \pi_k\}$ *is unique up to a rearrangement. Also 2) the subsequence* $\{a_j | \pi_j = i\}$ *(the powers of* $p_i(x)$ *in the* u_j*-minimal polynomial of those of* u_i *whose minimal polynomial is a power of* p_i*) is unique up to a rearrangement, for any* i*.*
As a result of these, any linear transformation has a unique (up to rearrangement) Rational Canonical Form in which in terms of an appropriate basis the matrix representation would is block diagonal, with k *blocks. Furthermore, along the diagonal, for any* $i = 1, \dots, k$*, there exist a block of size* $d_i a_i \times d_i a_i$*. Having chosen an appropriate basis this block gets the following form:*

$$\mathrm{A}_{i,i} = \begin{bmatrix} \mathrm{P}_i & \mathrm{C} & & & & \\ & \mathrm{P}_i & \mathrm{C} & & 0 & \\ & & \ddots & \ddots & & \\ & & & & \mathrm{C} & \\ & 0 & & \ddots & \mathrm{P}_i & \mathrm{C} \\ & & & & & \mathrm{P}_i \end{bmatrix}$$

with a_i *copies of the* $d_i \times d_i$ *matrix* P_i*, the companion matrix of* $p_i(x) = x^{d_i} - \alpha_{d_i-1} x^{d_i-1} - \dots - \alpha_1 x - \alpha_0$*, on the diagonal:*

$$\mathrm{P}_i = \begin{bmatrix} 0 & & & \dots & 0 & \alpha_0 \\ 1 & 0 & & \dots & 0 & \alpha_1 \\ & 1 & 0 & & & \vdots \\ & & 1 & 0 & & \\ & & & \ddots & & \\ & & & & 0 & \alpha_{d_i-2} \\ & & & & 1 & \alpha_{d_i-1} \end{bmatrix}$$

and C *is the all zero matrix except for a one on the top right corner:*

$$\mathrm{C} = \begin{bmatrix} 0 & \dots & & 1 \\ \vdots & 0 & & \vdots \\ & & \ddots & 0 \\ & & \dots & 0 \end{bmatrix}$$

Proposition 8.3.5. *In the case where* $\mathbb{F} = \mathbb{R}$ *we know that all prime factors of* $\mathbb{R}[x]$ *are quadratic or linear factors, and thus the blocks of the block diagonalization of the Rational Canonical form becomes has a maximum of block size equal to 2. The case where* $\mathbb{F} = \mathbb{C}$ *is the subject of the next section.*

8.4 λ-fixed subspaces and Jordan Normal Form

In the special case where $\mathbb{F} = \mathbb{C}$, or any other algebraically closed field for that matter, we know from the fundamental theorem of algebra 1.5.6 that all prime factors of $\mathbb{C}[x]$ are linear polynomials. As a result no matter what endomorphism T we are studying, its prime factors are of the following form:

$$p_i(x) = x - \lambda_i$$

and thus the linear transformations $p_i(\mathrm{T})$ get the following form:

$$p_i(\mathrm{T}) = \mathrm{T} - \lambda_i \mathrm{I}$$

and the linear transformations $p_i(\mathrm{T})^a$ would be:

$$p_i(\mathrm{T}) = (\mathrm{T} - \lambda_i \mathrm{I})^a$$

In the rational canonical form of T, all the companion matrices become 1×1 entries lying on the diagonal, and the C matrix becomes a 1×1 entry containing 1. The ones on the superdiagonal vanish when there is a change between the diagonal value.

Proposition 8.4.1 (Jordan Normal Form)**.** *Let* $\mathscr{V}$ *be a vector space defined over the algebraically closed field* $\mathbb{F}$*. Let* $\mathrm{T}\colon \mathscr{V} \to \mathscr{V}$ *be any linear transformation. There exists a basis* $\mathscr{B}$ *for* $\mathscr{V}$ *such that* $[\mathrm{T}]_{\mathscr{B}}$ *becomes bidiagonal with the superdiagonal being 0 or 1.*

We put specific names to the concepts we already have defined. Let $\mathrm{T}\colon \mathscr{V} \to \mathscr{V}$ be a linear transformation, and $p(x)$ be some prime factor in $\mathbb{C}[x]$ defined as:

$$p(x) = x - \lambda$$

for any arbitrary $\lambda \in \mathbb{C}$. We call the subspace $\mathrm{Ker}[p(\mathrm{T})]$, which now becomes $\mathrm{Ker}[\mathrm{T} - \lambda\mathrm{I}]$, the λ-**fixed subspace** of $\mathscr{V}$. If this subspace is non-trivial, which happens only if $(x - \lambda) | m_{\mathrm{T}}(x)$, we call λ an **eigenvalue** of T, and we refer to $\mathrm{Ker}[p(\mathrm{T})] = \mathrm{Ker}[\mathrm{T} - \lambda\mathrm{I}]$ by the λ-**eigenspace**:

$$\mathscr{W}_{\lambda_i} = \mathrm{Ker}[\mathrm{T} - \lambda\mathrm{I}]$$

and call the dimensionality of this subspace the **dimensionality** of the eigenvalue λ:

$$e_i = \dim\big[\mathrm{Ker}[\mathrm{T} - \lambda\mathrm{I}]\big] = \eta(\mathrm{T} - \lambda\mathrm{I})$$

Obviously the minimal polynomial of T would factor as the following:

$$m_{\mathrm{T}}(x) = \prod_{i=1}^{k}(x - \lambda_i)^{e_i}$$

The linear transformation $p(\mathrm{T})$, as any other linear transformation, has a stabilizing power. Without getting into the technicality of introducing stabilizing powers, we use that fact that whatever it is, it is no larger than n, the dimensionality of $\mathscr{V}$. Now we call the stabilized kernel $\mathrm{Ker}[p(\mathrm{T})^n]$, which is now $\mathrm{Ker}[(\mathrm{T} - \lambda\mathrm{I})^n]$, the λ-**generalized eigenspace**:

$$\mathscr{W}^{\lambda_i} = \mathrm{Ker}[(\mathrm{T} - \lambda\mathrm{I})^n]$$

and call the dimensionality of this subspace the **multiplicity** of the eigenvalue λ.

$$d_i = \dim\big[\mathrm{Ker}[(\mathrm{T} - \lambda\mathrm{I})^n]\big] = \eta\big[\mathrm{T} - \lambda\mathrm{I})^n\big]$$

Trivially the generalized eigenspace is a superset of the ordinary eigenspace:

$$\mathscr{W}_{\lambda_i} \subseteq \mathscr{W}^{\lambda_i}$$

Following immediately from proposition 8.3.3, we know that if $\lambda_1, \lambda_2, \ldots, \lambda_k$ are *all* the distinct eigenvalues of T, $\mathscr{V}$ gets directly decomposed to the respective *generalized* eigenspaces:

$$\mathscr{V} = \mathscr{W}^{\lambda_1} \oplus \mathscr{W}^{\lambda_2} \oplus \ldots \oplus \mathscr{W}^{\lambda_k}$$

And as a result the *multiplicities* of the eigenvalues of any linear transformation T adds up to n, the dimensionality of $\mathscr{V}$. We will see later that the characteristic polynomial of T factors as the following:

$$c_{\mathrm{T}}(x) = \prod_{i=1}^{k}(x - \lambda_i)^{d_i}$$

An important theorem in the case where $\mathbb{F}$ is algebraically closed is Schur's thoerem:

Proposition 8.4.2 (Schur's Theorem)**.** *Let $\mathscr{V}$ be a vector space defined over the algebraically closed field $\mathbb{F}$. For any arbitrary linear transformation $\mathrm{T}\colon \mathscr{V} \to \mathscr{V}$ there exists a basis $\mathscr{B}$ for $\mathscr{V}$ such that $[\mathrm{T}]_{\mathscr{B}}$ is upper (or equivalently lower) triangular.*

Translating the rational canonical form discussions to matrix terminology, neat observations can be made through Schur's theorem.

Proposition 8.4.3. *Let $\mathrm{T} : \mathscr{V} \to \mathscr{V}$ be a linear transformation with its Schur form being the triangular matrix* A*. There exists* $\dim\big(\mathrm{Ker}[\mathrm{T}^n]\big)$ *zeros on the diagonal of* A*.*

This is of course the *multiplicity* of the 0 eigenvalue. It would immediately follow from this proposition that the number of occurrences of any λ on the diagonal of the Schur form of T is the multiplicity of that eigenvalue.

8.5 Computational Issues of canonical forms

8.5.1

Computation of Eigenvalues
The QR algorithm

8.5.2

Algorithms for Jordan Normal Form

8.5.3

Algorithms for Rational Canonical Form

Bibliography

[Cur84] Charles W. Curtis, *Linear algebra: an introductory approach*, Springer-Verlag, New York, 1984.

[Fin66] Daniel T. Finkbeiner, *Introduction to matrices and linear transformations*, W.H. Freeman, San Francisco, 1966.

[HK71] Kenneth Hoffman and Ray Alden Kunze, *Linear algebra*, Prentice-Hall, Englewood Cliffs, N.J., 1971.

[Hun74] Thomas W. Hungerford, *Algebra*, Holt, Rinehart and Winston, New York, 1974.