

A primer of linear algebra

Introduction:p.3. Summary of contents.

Chapter One: Linear spaces and linear maps, theory of dimension.p.5.

linear spaces and subspaces;
linear maps;
product spaces;
quotient spaces;
linear combinations;
independent sets;
spanning sets;
bases;
dimension;
kernel, rank, image;
dual spaces, orthogonal complements;
dual of a map, double duals of spaces, maps;
left and right equivalence.

Chapter Two: Computations using matrices. p.24.

Elementary row operations;
row reduced echelon form; pivot variables; pivot columns;
row space, column space, null space;
existence and uniqueness of row reduced echelon form;
Excursion: the Grassmannian variety of projective lines in P^3 ;
matrices and linear maps;
dot products and matrix multiplication;
row equivalent matrices have the same null space;
column rank = row rank;
parametric versus implicit representation of a subspace;
representing a linear map by a matrix, using a basis;
representing composition of maps by matrix multiplication;
matrix inverses;
matrix of transpose = transpose of matrix;
elementary matrices for elementary row operations;
calculating matrix inverses by row operations;
row equivalence of matrices = left equivalence of maps;

Chapter Three: Decomposing a space with operator T, into T-cyclic subspaces. p.59.

Polynomials satisfied by an operator;
minimal polynomial at a vector; minimal polynomial on a subspace;
degree of minimal polynomial \leq dimension of subspace;
cyclic bases and companion matrices;
Invariant factor theorem: (a space with operator T, is a product of maximal T-cyclic subspaces, & the minimal polynomials divide each other);
Computing the invariant factor decomposition from a presentation matrix;
the “characteristic presentation”, characteristic polynomial;
diagonalizing the characteristic matrix;
the minimal polynomial divides the characteristic polynomial.

Chapter Four: General and classical Jordan form. p.80.

Decomposing a space with operator into minimal T-cyclic subspaces;
Generalized eigenspaces; eigenvalues; Jordan blocks;
Diagonalizable operators;
Computing Jordan bases, (given the eigenvalues);
Nilpotent operators;

Chapter Five: “Spectral theorems”, orthogonally diagonalizable operators.p.107.

Real symmetric operators are orthogonally diagonalizable;
Structure of real isometries;
Complex Hermitian operators are orthogonally diagonalizable;
Linear differential equations.

Appendix: Review of determinants.p.125.

Introduction:

In the first chapter of these notes we discuss the concept of dimension of a vector space over a field k , in particular what a basis is, as well as the constructions of product spaces and quotient spaces. After those fundamentals, the next concept is that of a linear transformation, or linear map, between two vector spaces, for example the operation of differentiation from the space of smooth real valued functions to itself. Our principal goal is to understand the structure of an arbitrary linear transformation of a finite dimensional space to itself, the subject of chapters 3 and 4. The key to answering this question is the “minimal polynomial” satisfied by a linear operator T on a finite dimensional space, i.e. the unique monic polynomial f of least degree over k such that $f(T) =$ the zero operator. (The fact that on a finite dimensional space, the space of linear operators is also finite dimensional, implies there must be a linear relation among the infinitely many powers T^m of T , hence T satisfies some polynomial relation f .)

The model example of a linear operator in finite dimensions is the transformation T defined by multiplication by X , on the quotient space $k[X]/(f)$ where $f = a_0 + a_1X + \dots + a_{n-1}X^{n-1} + X^n$ is a monic polynomial of degree n . This space has natural basis $1, X, X^2, \dots, X^{n-1}$, and the map T permutes these basis vectors “cyclically”, i.e. 1 is sent to X , X is sent to X^2 , ..., until the last one X^{n-1} , is sent to X^n , which in the quotient space equals the following linear combination of the basis vectors: $-a_0 - a_1X - a_2X^2 - \dots - a_{n-1}X^{n-1}$. The minimal polynomial of this operator is just f , since f is the polynomial of lowest degree such that $f(T)$, i.e. multiplication by $f(X)$, is identically zero. Remarkably, this simple example is universal for all linear transformations in finite dimensions. Indeed the structure theorem proved in chapter 2 implies that every linear transformation T on a finite dimensional space V is isomorphic to a product of these models, where the polynomials f are factors of the minimal polynomial of T on all of V . Thus determining the minimal polynomial of a given transformation is a crucial step in understanding that operator.

The question of how to find the polynomials f associated to the decomposition of an operator T , in particular the minimal polynomial, is answered by the concept of a “characteristic matrix” for T . Briefly, given a basis for the n dimensional space V , the operator $T:V \rightarrow V$ is represented by an $n \times n$ matrix A with entries in the field k , and the characteristic matrix in this basis, is the matrix of polynomials $[X.I-A]$.

Diagonalizing this matrix yields the polynomials f associated to the model examples above describing T .

If the diagonal entries are chosen to be monic and to successively divide each other, as may be done, then they are unique and the last one, the one of largest degree, is the minimal polynomial. In particular this diagonalized matrix is independent of the basis chosen to form A . I.e. although choosing a different basis for V will usually give a different matrix A for the operator T , and hence also a different characteristic matrix $[X.I-A]$, after diagonalizing so that the diagonal entries are monic and successively divide each other, we always get the same diagonal matrix. In particular, the product of the diagonal entries, which equals the determinant of the characteristic matrix $[X.I-A]$, is independent of choice of basis, and is called $\text{ch}_A(X)$ = the “characteristic polynomial” of T . By the divisibility condition, $\text{ch}_A(X)$ has the same irreducible factors as the minimal polynomial, but possibly with higher multiplicities; in particular the characteristic polynomial is a multiple of the minimal polynomial (“Cayley - Hamilton theorem”).

If we can factor $\text{ch}_A(X)$ completely into irreducible factors, we can describe T in terms of models where each polynomial f is a power of an irreducible polynomial. This factorization, although always theoretically possible, may not be feasible in practice. But at least in theory, over an algebraically closed field like the complex numbers, we can thus assume each polynomial f is $(X-c)^n$, for c a constant. Then the preferred basis for the model space $k[X]/(f)$ is the set $\{1, (X-c), \dots, (X-c)^{n-1}\}$, i.e. powers of $(X-c)$ instead of powers of X .

Since we have set $f(T) = (T-c)^n = 0$, in this basis the map $T = X = c + (X-c)$, is the sum of the scalar multiplier c , plus the “nilpotent” operator multiplication by $(X-c)$. T thus sends each basis vector, except the last, to c times itself plus the next basis vector, and sends the last vector just to c times itself. In this basis the matrix is said to be in “Jordan form”. It is often of theoretical value to know that a Jordan form exists even if it cannot be explicitly computed.

The simplest Jordan matrix is a diagonal matrix. This simplest case occurs when the minimal polynomial not only factors completely into linear factors like $(X-c)$, but when every such factor occurs in the minimal polynomial with exponent $n = 1$. The map T is then a product of copies of multiplication by X on spaces of form $k[X]/(X-c)$. Thus the space V is a product of subspaces on each of which the map T is just multiplication by some constant c . These are called diagonalizable maps, and fortunately there exist useful criteria to recognize these maps in some special cases, without having to actually carry out the factorization of the minimal

polynomial. The simplest of these criteria, over the real number field, is when the matrix for the map is “symmetric” about the main diagonal. This phenomenon occurs because such matrices not only have eigenvectors, but also preserve perpendicularity, hence are orthogonally diagonalizable. Matrices that are not symmetric but are “length preserving”, the so called “rigid motions”, also preserve perpendicularity but may not have eigenvectors. These may not be diagonalizable, but can be expressed as combinations of mutually orthogonal rotations and reflections. These criteria, called "spectral theorems", are discussed in chapter 5.

To summarize, given a linear operator T on a finite dimensional k -vector space V , V has a decomposition into a product of subspaces on each of which T is isomorphic to the action of multiplication by X on a quotient space $k[X]/(f)$. These subspaces can be chosen so that the corresponding sequence of monic polynomials f_1, \dots, f_r successively divide each other, and when this is done the sequence of polynomials is uniquely determined by T . This, the so called “invariant factor decomposition”, can be computed by hand from any matrix for T . Two operators S, T on the same space are “similar”, i.e. $T = (U^{-1})SU$ for some invertible operator U , if and only if S, T have the same invariant factors.

A second standard decomposition exists where the polynomials f in the model spaces are all powers of irreducible polynomials. For this decomposition, the sequence of polynomials is again uniquely determined by T , except for a chosen ordering of the irreducible polynomials. This decomposition, called the “generalized Jordan decomposition”, always exists in theory, but can be computed in practice only for those examples where the irreducible factors of the minimal polynomial of T can actually be found, e.g. for a “triangular” matrix.

A special case of the Jordan decomposition occurs precisely when the minimal polynomial factors completely into distinct linear factors. Then the Jordan form, which may or may not be effectively computable, is a diagonal matrix. This is always the case when the matrix consists of real entries which are symmetric about the main diagonal.

Finally, in the appendix we recall some properties of determinants.

Chapter One: Linear spaces and linear maps, theory of dimension

Linear algebra is about linear spaces or vector spaces, and linear maps, (also called linear operators, or linear transformations), between them. The first topic is therefore linear spaces. Vector spaces make sense over any “field” of scalars, i.e. any reasonable number system with commutative addition and multiplication, and where division is possible by all non zero elements. For simplicity we discuss mostly vector spaces defined over the field \mathbb{R} of real numbers. In some proofs we will exploit the fact that if k is any field and f is an irreducible polynomial over k , then the quotient space $k[X]/(f)$, defined below, is a field larger than k and to which a root of the polynomial $f(X)$ has been adjoined. E.g. the field $\mathbb{R}[X]/(X^2+1)$ is essentially the field of complex numbers, since in this quotient field $X^2 + 1 = 0$, so that X plays the role of i , a square root of -1 .

Defn: A (real) **vector space** V , is a set V of "vectors" closed under addition, and closed under “scalar” multiplication of vectors by real numbers, such that V is an “abelian group” under addition (the usual properties hold, like associativity, commutativity and existence of a zero and negatives), and has the usual properties under scalar multiplication (multiplication by 1 acts as the identity, multiplication distributes over addition, and $a(bv) = (ab)v$ if a, b , are numbers and v is a vector).

Eg: The basic example is \mathbb{R}^n , ordered n - tuples (v_1, \dots, v_n) of real numbers v_i , with component - wise addition and multiplication: $(v_1, \dots, v_n) + (w_1, \dots, w_n) = (v_1+w_1, v_2+w_2, \dots, v_n+w_n)$, $a(v_1, \dots, v_n) = (av_1, \dots, av_n)$.

Other important examples are the space $\mathbb{R}[X]$ of polynomials in one variable with real coefficients, and the space $\text{Diff}(\mathbb{R})$ of differentiable real valued functions defined on the real line.

Defn: A “**subspace**” of V is a non empty subset W of V , which is closed under addition and scalar multiplication. In particular W is also a vector space.

E.g. If $V = \mathbb{R}^n$, the subspace W might be the subset of all (v_1, \dots, v_n) such that $v_1 + \dots + v_n = 0$, and if $V = \text{Diff}(\mathbb{R})$, W might be those f with $f(0) = f'(0) = 0$.

The space $\text{Diff}(\mathbb{R})$ is a subspace of the space $C(\mathbb{R})$ of continuous functions on \mathbb{R} . The set of solutions of the differential equation $f'' + f = 0$ is the subspace of the

space $\text{Diff}(\mathbb{R})$, consisting of functions of form $a.\sin(t) + b.\cos(t)$.

Defn: Given a subspace W of V , we define a new vector space V/W , the “quotient” of V by W , by identifying two vectors u, v in V provided $u-v$ lies in W . Addition is defined by setting $[u] + [v] = [u+v]$, and $c[v] = [cv]$, where $[v]$ denotes the equivalence class of the vector v .

E.g. if $V = \mathbb{R}^n$, and W is those (v_1, \dots, v_n) with $v_1 + \dots + v_n = 0$, then two vectors are equivalent iff their difference has coefficients which sum to zero. Thus two vectors are equivalent if their coefficients have the same sum, so each equivalence class corresponds to a real number, the common sum of the coefficients of vectors in that class. I.e. V/W is essentially \mathbb{R} , where the equivalence class of (v_1, \dots, v_n) corresponds to the real number $v_1 + \dots + v_n$.

Similarly, if $V = \text{Diff}(\mathbb{R})$, and W is those f with $f(0) = f'(0) = 0$, then V/W is essentially \mathbb{R}^2 , where the equivalence class $[f]$ corresponds to the pair of real numbers $(f(0), f'(0))$.

Rmk: The precise meaning of “essentially” above, is given by the notion of isomorphism below.

Defn: For any two vector spaces V, W we define a new space $V \times W$, the “direct product” or simply “product” of V and W , consisting of all ordered pairs (v, w) with v in V and w in W . Addition and multiplication are defined on components separately. Similarly $V_1 \times \dots \times V_n$ is defined as the set of n tuples whose i th element is a vector in V_i .

E.g. $\mathbb{R} \times \mathbb{R}$ is precisely \mathbb{R}^2 = the product of two copies of the real numbers, and $\mathbb{R} \times \dots \times \mathbb{R}$ (n factors), is just \mathbb{R}^n .

We define next a way to compare two vector spaces to see when they are essentially the same.

Defn: A map $T: V \rightarrow W$ from V to W , (vector spaces over the same field F), is **linear** if $T(x+y) = T(x)+T(y)$ for all x, y , in V , and if also $T(ax) = aT(x)$ for all x in V and all real numbers a .

Defn: The composition of two linear maps $T: V \rightarrow W$, and $S: W \rightarrow U$, is the map $(S \circ T): V \rightarrow U$, such that $(S \circ T)(x) = S(T(x))$ for all x in V .

Ex. The composition of two linear maps is also linear.

Defn: An **isomorphism** is a linear map $T:V \rightarrow W$ with a (2-sided) linear inverse $S:W \rightarrow V$. I.e. S and T are both linear and $S(T(x)) = x$ for all x in V , and $T(S(y)) = y$ for all y in W . S is then also an isomorphism with T as its inverse.

Thus if we define the identity map of a space as that map taking each vector to itself, then a linear map T is an isomorphism if there exists some linear map S such that both compositions $S \circ T$ and $T \circ S$ are the corresponding identity maps.

Defn: A left inverse for a linear map $T:V \rightarrow W$ is a linear map $S:W \rightarrow V$ such that $S(T(x)) = x$ for all x in V . S is a right inverse for T if $T(S(y)) = y$ for all y in W .

Ex. Show that if a linear map T has a left inverse then T is injective, and if T has a right inverse then T is surjective.

We will consider two vector spaces as essentially the same, i.e. as having the same linear structure, when there exists an isomorphism between them. To actually identify their elements, we will need to choose a specific isomorphism.

Ex: 1) Given a space V and subspace W , addition and scalar multiplication are well defined on V/W independent of choice of representatives of equivalence classes, and the map $V \rightarrow V/W$ sending x to $[x]$, is a linear map sending just the vectors in W to zero, i.e. $[x] = [0]$ iff x is in W .

2) If $V = \mathbb{R}^n$, and W is those (v_1, \dots, v_n) with $v_1 + \dots + v_n = 0$, then the map $V/W \rightarrow \mathbb{R}$ taking (v_1, \dots, v_n) to $v_1 + \dots + v_n$, is a linear map, with inverse $\mathbb{R} \rightarrow V/W$ taking t to $[(t, \dots, 0)]$.

3) Every bijective linear map is an isomorphism, i.e. if a linear map has an inverse map, that inverse is also linear.

4) The set $\text{Hom}(V, W)$ of all linear maps $V \rightarrow W$ is closed under addition and scalar multiplication, where $(S+T)(v) = S(v) + T(v)$ and $(cS)(v) = c(S(v))$, hence $\text{Hom}(V, W)$ is also a vector space. The special case $\text{Hom}(V, \mathbb{R}) = V^*$ is called the "dual" vector space of V .

5) $\text{Hom}(\mathbb{R}, V)$ is isomorphic to V , where the linear map f corresponds to the vector

$f(1)$ in V .

The concept of linear combinations

Since non zero vector spaces contain infinitely many vectors, it is useful to be able to represent all vectors in a given space in terms of only a finite subset of them, or if that is not possible, in terms of some distinguished subset. This is done by taking "linear combinations".

Defn: A "linear combination" of the vectors $\{v_1, \dots, v_m, \dots\}$ is a finite sum of scalar multiples of the given ones, i.e. an expression of form $w = a_1v_1 + \dots + a_mv_m$, where the a 's are real numbers. If $w = a_1v_1 + \dots + a_mv_m$, i.e. if w equals such a linear combination, we also call w a linear combination of the $\{v_i\}$.

Eg: In \mathbb{R}^3 , $(4, -5, 1)$ is a linear combination of $(2, 2, 3)$ and $(8, -1, 7)$, since $(4, -5, 1) = (8, -1, 7) - 2(2, 2, 3)$.

Defn: A set S of vectors "**spans**" or "**generates**" a vector space V iff every non zero vector in V is a linear combination of vectors in S , or equivalently if the set S is not contained in any proper subspace of V . In particular, the empty set spans the space $\{0\}$.

Rmk: Although we were careful to avoid saying it, it actually makes good sense to assert that the zero vector is a linear combination of the vectors in the empty set, since this is compatible with the way summations behave. I.e. if we have an indexed collection of vectors and we partition the index set into two disjoint subsets, we would expect that summing over each subset separately and then adding the two results, would give the same result as summing over the whole index set. For this to be true also for the partition into the empty subset and the whole subset, we therefore need the sum over the empty set to be zero. For the same reason, any operation carried out over the empty index set should give the identity element for that operation. E.g. the product of the empty set of real numbers should equal 1, and the union of an empty collection of subsets should equal the empty set, while the intersection of an empty collection of subsets of a space should equal the whole space. This convention is very helpful, since by using it, whenever we partition an index set, we do not have to stop and check whether all the subsets are non empty. With this understanding we could simplify the definition above and just say that S spans V if every vector in V is a linear combination of vectors in S .

Eg: The set $\{(1,0), (0,1)\}$ spans \mathbb{R}^2 since every vector (a,b) can be written as the linear combination $a(1,0) + b(0,1) = (a,b)$. More generally \mathbb{R}^n is spanned by the set of n standard vectors $\{e_1 = (1,0,\dots,0), e_2 = (0,1,0,\dots,0), \dots, e_n = (0,\dots,0,1)\}$.

Ex: For any subset $S = \{v_1, \dots, v_m, \dots\}$ of a vector space V , the set of all finite linear combinations of vectors in S (which includes 0 even if S is empty), is a subspace $L(S)$ of V spanned by S .

Ex: If $V = \mathbb{R}^n$ and W is the subspace spanned by e_n , then V/W is isomorphic to \mathbb{R}^{n-1} .

Two natural problems arise involving linear combinations: given a set of vectors in V , we want to know what linear subspace of V they span; and dually, given a subspace W of V , we want to be able to find a nice set of vectors, e.g. as small a set as possible, that span W . Spaces that can be spanned by a finite set of vectors are especially important.

Defn: A space V is **finite dimensional** iff V has a finite spanning set.

For example \mathbb{R}^n is finite dimensional since it can be spanned by the set of n standard vectors $\{(1,0,\dots,0), (0,1,0,\dots,0), \dots, (0,\dots,0,1)\}$. The space of real polynomials of degree $\leq d$ in t , is finite dimensional since it can be spanned by the $d+1$ monomials $\{1, t, t^2, \dots, t^d\}$. As we will learn later, the subspace W of $\text{Diff}(\mathbb{R})$ consisting of solutions of $f'' + f = 0$, is spanned by $\sin(t)$ and $\cos(t)$, hence is finite dimensional.

Linear independence, bases and the concept of dimension

Next we will see how to assign a specific dimension to each finite dimensional vector space. This is the most important concept in the whole subject. In fact the notion of dimension completely determines a vector space up to isomorphism.

Since we will see that \mathbb{R}^n has dimension n , this will mean that the spaces \mathbb{R}^n give a complete list of equivalence classes of finite dimensional vector spaces. Thus to see how to define dimension, we naturally look at \mathbb{R}^n .

Recall that we draw the spaces \mathbb{R} , \mathbb{R}^2 , \mathbb{R}^3 , as spanned by axis systems: one axis for \mathbb{R} , two axes for \mathbb{R}^2 , three axes for \mathbb{R}^3 , and it is not too great a stretch to imagine n axes for \mathbb{R}^n . So the dimension might be thought of as the number of axes needed

to represent the space, and we want to express this algebraically. Since each axis consists of vectors whose coordinates have form $(0, \dots, 0, t, 0, \dots, 0)$, each axis is a line spanned by a standard vector of form $(0, \dots, 0, 1, 0, \dots, 0)$. So the number of axes needed to span a space geometrically, should be equivalent to the number of vectors needed to span the space algebraically. Notice \mathbb{R} can be spanned by the single vector 1, but no fewer (why not?), and \mathbb{R}^2 can be spanned by $(1,0)$ and $(0,1)$, but no fewer, since the multiples of a single vector (a,b) would always have entries proportional to (a,b) . We use this idea to define dimension as follows:

Definition: The dimension of a finite dimensional vector space is the minimum cardinality of a spanning set. I.e. a vector space V has finite dimension n , if V has a spanning set of n vectors but no spanning set with fewer than n vectors.

It is not so obvious that for every n , \mathbb{R}^n has dimension n . Of course the n standard vectors do span \mathbb{R}^n , so the dimension is at most n , but it is harder to prove that \mathbb{R}^n cannot be spanned by fewer than n vectors. (If you can prove this from scratch on your own, or come pretty close, even for $n = 3$, you are a potential mathematician).

So the problem is how to detect when a spanning set has the minimum possible number of vectors. Certainly a spanning set does not have the minimum number of elements if one or more of its vectors are superfluous, in the sense that after throwing some of them out the remaining set still spans, since that remaining spanning set would have fewer vectors. So when does this happen? If w can be eliminated from the spanning set $\{v_1, \dots, v_m, w\}$ and the remaining vectors $\{v_1, \dots, v_m\}$ still span, then w must be a linear combination of the remaining vectors $\{v_1, \dots, v_m\}$, since everything is. We give this situation a name:

Provisional definition: A set of vectors is called (linearly) dependent if some vector in it is in the span of the other vectors.

Rmks: Since even the span of the empty set contains 0, any set containing a zero vector is dependent. The empty set is independent, and a set containing only one vector is independent if and only if that vector is non zero.

In practice we want to consider indexed sets of vectors, in which some of the vectors with different indices may be the same vector. Thus we want to consider an indexed set to be dependent if one of the vectors is in the span of the vectors with different indices. Thus any indexed set with a repeated vector, i.e. a set in which the same vector occurs with two different indices, is dependent.

Refined definition: An indexed set of vectors $\{v_1, \dots, v_m, \dots\}$ is dependent if there is some index j , such that v_j is a finite linear combination of vectors v_i with $i \neq j$.

Finally, note that if a vector v in an indexed set, can be written as a linear combination of the others, then by subtracting v from both sides of this linear combination, we have expressed the zero vector as a linear combination of all the vectors, and at least one coefficient is not zero, since v now has coefficient -1 . Conversely, if we can express zero as a finite linear combination of vectors in which some coefficient is non zero, then we can divide through by that coefficient so that it becomes one, and then by putting all the other vectors on the other side, we have expressed that one vector as a linear combination of the others. This gives a more efficient, if less intuitive, way of considering linear dependence of a set without singling out any one vector. In particular, we have proved the following.

Lemma: An indexed set of vectors $\{v_1, \dots, v_m, \dots\}$ is (linearly) **dependent** if there is a finite indexed set of real numbers $\{a_1, \dots, a_m\}$, not all zero, such that $a_1 v_1 + \dots + a_m v_m = 0$.

Equivalently, An indexed set of vectors $\{v_1, \dots, v_m, \dots\}$ is **independent** iff the only scalars a_1, \dots, a_m such that $a_1 v_1 + \dots + a_m v_m = 0$, are $a_1 = a_2 = \dots = a_m = 0$.

Useful remark: Note that a finite or infinite sequence of vectors $\{v_1, \dots, v_m, \dots\}$ is dependent if and only if some vector is a linear combination of earlier vectors in the sequence. To see this, given a finite dependency relation $a_1 v_1 + \dots + a_m v_m = 0$, in which not all coefficients are zero, then after deleting coefficients which are zero we may assume $a_m \neq 0$, and then we can solve for v_m as a linear combination of the previous vectors. I.e. just solve for the last vector occurring with a non zero coefficient in the linear combination.

Eg. In the space of all polynomials in t , the infinite set of all monomials $\{1, t, t^2, \dots, t^n, \dots\}$ is independent, since no monomial is a linear combination of monomials of lower degrees.

Eg. $\{(1,0), (0,1)\}$ is independent since the only way we can have $a(1,0) + b(0,1) = (a,b) = (0,0)$, is to have $a = b = 0$.

Eg. The empty set is independent, since there are no vectors at all, hence none which depend on the others.

Eg. Another argument that the set of monomials $\{1, x, x^2, \dots, x^n\}$ is independent, is to view them as differentiable functions, since if $f = a_0 1 + \dots + a_n x^n = 0 =$ the zero function, then $0 = f(0) = a_0$, and $0 = f'(0) = a_1, \dots, 0 = f^{(n)}(0) = n! a_n = a_n$, so all the coefficients are zero.

Defn: A subset S of V is called a **basis** of V , if S is independent and $L(S) = V$, i.e. a basis is just an independent spanning set.

Thm: Every finite dimensional space V has a basis.

Pf: Choose a finite spanning set $S = \{v_1, \dots, v_n\}$ of V . Throw out all zero vectors. If v_2 is a multiple of v_1 , throw out v_2 , if not keep it. If v_3 is a linear combination of $\{v_1, v_2\}$, throw out v_3 , if not keep it. Continue throwing out vectors which are linear combinations of previous ones, which does not change the span. The ones left are a basis, since no vector depends on previous ones. **QED.**

Rmk: We have actually shown that every finite spanning set contains a basis. The same argument shows that every independent set in a finite dimensional space can be enlarged to a basis. Just start by augmenting the independent set by adding in a spanning set at the end. Then discarding vectors that depend on earlier ones as in the previous proof gives a basis containing the original independent set.

Eg: The set of unit vectors $e_1 = (1, 0, \dots, 0), \dots, e_n = (0, \dots, 0, 1)$, is a basis of \mathbb{R}^n called the “standard basis”. $\{(3, 0), (2, 5)\}$ is another basis of \mathbb{R}^2 . The empty set is a basis of $\{0\}$.

We have remarked above that a finite spanning set with the minimum number of vectors must be independent, i.e. must be a basis. We claim that conversely every basis, i.e. every independent spanning set, does have the minimum number of vectors. For this it suffices to show all bases have the same number of vectors. Then we will be able to determine the dimension of a space just by finding one basis and counting the number of vectors.

We will give two proofs, the second one being the famous “exchange lemma”. Our first proof is not the shortest possible, but will help us learn some important results about linear maps. First you will show that if a space has a basis with n vectors then that space is isomorphic to \mathbb{R}^n . Then we will show that if \mathbb{R}^n is isomorphic to

\mathbb{R}^m then $n = m$. Thus if a space has two bases with n vectors and m vectors respectively, then it is isomorphic to both \mathbb{R}^n and \mathbb{R}^m which are thus isomorphic to each other, so $n = m$; i.e. any two finite bases of the same space have the same number of vectors. Here is your part of the proof:

Ex: 1) If $S = \{v_1, \dots, v_n\}$ is a finite sequence in V , i.e. a function from $\{1, \dots, n\}$ to V , there is a unique linear map $T: \mathbb{R}^n \rightarrow V$ sending (a_1, \dots, a_n) to $a_1 v_1 + \dots + a_n v_n$.

2) The map T in 1) above is injective if and only if $S = \{v_1, \dots, v_n\}$ is independent, and

3) T is surjective if and only if $S = \{v_1, \dots, v_n\}$ spans V , and

4) T is an isomorphism if and only if $S = \{v_1, \dots, v_n\}$ is a basis for V .

Cor (of 1): There is a one to one correspondence between linear maps $T: \mathbb{R}^n \rightarrow V$ and ordered subsets $\{v_1, \dots, v_n\}$ of n vectors in V . In fact this correspondence is linear, so defines an isomorphism of vector spaces $\text{Hom}(\mathbb{R}^n, V) \approx \text{Hom}(\mathbb{R}, V) \times \dots \times \text{Hom}(\mathbb{R}, V) \approx V \times \dots \times V$, n factors.

Remark: In particular, $\text{Hom}(\mathbb{R}^n, \mathbb{R}^m) \approx \mathbb{R}^m \times \dots \times \mathbb{R}^m$, n factors. So a linear map $\mathbb{R}^n \rightarrow \mathbb{R}^m$ is defined by a sequence of n vectors in \mathbb{R}^m . If these vectors are arranged in order as a sequence of n column vectors each of length m , the resulting array is called an “ m by n matrix”. We will study in the next chapter how to use such matrices to compute invariants of the corresponding linear map. In particular they allow us to compute linear maps as a sort of “matrix” multiplication.

Cor (of 4): There is a one to one correspondence between isomorphisms $T: \mathbb{R}^n \rightarrow V$ and ordered bases $\{v_1, \dots, v_n\}$ of V . In particular, since a finite dimensional space V has a basis, it is also isomorphic to some \mathbb{R}^n .

Cor (of 1,4): A basis B of V defines a one - one correspondence between linear maps from V to another space W , and set functions from B to W , i.e. every function $B \rightarrow W$ extends uniquely to a linear map $V \rightarrow W$.

Pf : By 1),4) there is an isomorphism between \mathbb{R}^n and V , taking the basis B for V to the standard basis for \mathbb{R}^n . Then property 1) for \mathbb{R}^n translates into this statement for B and V . **QED.**

Remark: This last corollary is very important, since it tells us how easy it is to define linear maps on a space for which we have a basis, in particular there is always a map sending any given basis to any values we please. Use this fact to do

the next exercises, thus providing converses to some earlier ones.

Ex. Show that an injective linear map $T:V \rightarrow W$ between finite dimensional spaces always has a (linear) left inverse. Then show a surjective linear map $T:V \rightarrow W$ between finite dimensional spaces always has a (linear) right inverse.

Terminology: An isomorphism $\mathbb{R}^n \rightarrow V$ is often called a **parametrization** of V .

Rmk: Dually to the case of maps $\mathbb{R}^n \rightarrow V$, a linear map $V \rightarrow \mathbb{R}^n$ corresponds to an ordered sequence of n linear maps $V \rightarrow \mathbb{R}$, i.e., a sequence of n elements of V^* called coordinate functions. This correspondence is again linear so $\text{Hom}(V, \mathbb{R}^n) \approx \text{Hom}(V, \mathbb{R}) \times \dots \times \text{Hom}(V, \mathbb{R}) \approx V^* \times \dots \times V^*$. This time the map $V \rightarrow \mathbb{R}^n$ is surjective iff the sequence of elements of V^* is independent, and injective iff the sequence of coordinate functions spans V^* . Again an isomorphism $V \rightarrow \mathbb{R}^n$ corresponds to an ordered basis of V^* .

Terminology: An isomorphism $V \rightarrow \mathbb{R}^n$ is often called a **coordinate system** for V .

I.e. since an ordered basis for V gives an isomorphism $\mathbb{R}^n \rightarrow V$, the inverse isomorphism $V \rightarrow \mathbb{R}^n$ gives a way to introduce linear coordinates into V , since each vector in V gets represented by a sequence of numbers, i.e. a coordinate vector in \mathbb{R}^n . More precisely, if $\{v_1, \dots, v_n\}$ is a basis for V , the associated isomorphism $T: \mathbb{R}^n \rightarrow V$ sends (a_1, \dots, a_n) to $a_1 v_1 + \dots + a_n v_n$, while the inverse isomorphism $V \rightarrow \mathbb{R}^n$ sends each v to its coordinate vector (a_1, \dots, a_n) for the expansion of $v = a_1 v_1 + \dots + a_n v_n$ in the given basis.

Given a finite dimensional space V , a fundamental problem is to find as simple, or as “nice”, a basis as possible for V . E.g. once we introduce length and angle measure, it is often useful to have a basis of mutually perpendicular unit length vectors.

Now we have enough skill with linear maps to show the number of vectors in any basis always equals the dimension.

Lemma: A linear surjection $T: \mathbb{R}^n \rightarrow V$ which is not injective, restricts to an isomorphism from some linear subspace \mathbb{R}^m of \mathbb{R}^n to V (where $m < n$).

Pf: T takes the standard basis of \mathbb{R}^n to a generating set for V . Reduce this set to a

basis B , and choose a subset S of standard basis vectors of \mathbb{R}^n mapping bijectively to B , hence an isomorphism from the subspace $L(S)$ of \mathbb{R}^n , to V . $L(S)$ is easily identified with \mathbb{R}^m where m is the number of vectors in the subset S . **QED.**

Thm: If \mathbb{R}^n and \mathbb{R}^m are isomorphic, then $n = m$.

Pf: (induction on n) There is no linear surjection $T: \mathbb{R}^1 \rightarrow \mathbb{R}^m$ if $m > 1$, since all the image vectors of T have proportional entries. Assume T is a linear surjection $T: \mathbb{R}^n \rightarrow \mathbb{R}^m$, with $2 \leq n < m$ and n as small as possible. Let $\{e_1, \dots, e_n\}$, and $\{u_1, \dots, u_m\}$ be the standard bases of \mathbb{R}^n and \mathbb{R}^m . Then the composition $\mathbb{R}^n \rightarrow \mathbb{R}^m / \text{span}(u_m)$ is surjective but not injective, since if $T(v) = u_m$, then $v \neq 0$ maps to $[0]$ in $\mathbb{R}^m / \text{span}(u_m)$. Hence by the previous Lemma, $\mathbb{R}^n \rightarrow \mathbb{R}^m / \text{span}(u_m)$ restricts to a surjection from some subspace \mathbb{R}^k of \mathbb{R}^n , with $k < n$, to $\mathbb{R}^m / \text{span}(u_m) = \mathbb{R}^{m-1}$. Since $k < m-1$, and $k < n$, this contradicts the hypothesis that n is small as possible. **QED.**

Cor: Any two bases of a finite dimensional space have the same number of elements.

Cor: The dimension of a (finite dimensional) vector space is the number of elements of any basis.

Cor: Two finite dimensional spaces are isomorphic iff they have the same dimension, (since then they are both isomorphic to the same \mathbb{R}^n).

Note: This last corollary says we have classified all finite dimensional vector spaces by dimension, namely, up to isomorphism, there is exactly one space of each dimension.

This is our first big theorem, so it is worth giving two proofs. I made the previous proof up because I found the usual ones hard to remember. But I think everyone should see the classic proof, using the method of “exchanging” vectors to show that any spanning set must have at least as many vectors as any independent set. This shows any basis must be at least as large as any other basis, so they are all the same size.

Exchange lemma: If a vector space contains a spanning set $\{w_1, \dots, w_m\}$ and an independent set $\{v_1, \dots, v_n\}$, then $n \leq m$. (As usual, we regard these as indexed sets.)

proof: Suppose $n \geq 1$. Then add v_1 , to the set of w 's placing it first: i.e. form the set $\{v_1, w_1, \dots, w_m\}$. Now since the w 's span the whole space, this set is dependent, and according to an earlier argument, some vector is a linear combination of earlier ones. Since the set of v 's is independent, $v_1 \neq 0$ hence is not a linear combination of earlier ones. So one of the w 's must depend on earlier vector occurring in the set. In particular there are some w 's, and hence $m \geq 1$ also.

We will simply repeat this argument as long as any v 's remain. I.e. we can throw out the w that depends on earlier vectors in the set, and the resulting set will still span. Renumbering the w 's we can assume the one we threw out was w_1 . Now if $n \geq 2$, add v_2 to the set placing it at the beginning, getting the set $\{v_2, v_1, w_2, \dots, w_m\}$. Again, since the set spanned before adding in v_2 , it has now become dependent, and hence some vector depends linearly on earlier ones. Since the v 's are independent however, neither v_2 nor v_1 depend on earlier vectors, so there is at least one more w , and hence $m \geq 2$ also. After we have repeated this process n times, we have shown that $m \geq n$. **QED.**

Note: This proof is more elementary than ours since it does not use any of the relations between bases and isomorphisms with \mathbb{R}^n , hence it could have been given right after the definition of a basis. The argument is often attributed to Steinitz, and he may have been the first to give it in the context of abstract algebra. But the exchange argument occurs decades earlier in Riemann's works, well before Steinitz's birth, to prove essentially that the number of loops in every homology basis for a compact surface is the same.

Again we get the corollaries noted above.

Cor: Any two bases of a finite dimensional vector space have the same number of elements.

proof: If $\{w_1, \dots, w_m\}$ and $\{v_1, \dots, v_n\}$ are both bases, they are both independent and both spanning, hence by the exchange lemma $n \leq m$ and $m \leq n$, so $n = m$. **QED.**

Cor: The dimension of a finite dimensional space equals the number of elements in any basis.

Ex: 1) If v_1, \dots, v_n is a basis B of V , the elements f_1, \dots, f_n of V^* such that $f_i(v_j) = 0$ for $i \neq j$ and, $f_i(v_i) = 1$ for all i , is a basis of V^* called the basis dual to B . Hence $\dim(V) = \dim(V^*)$.

2) If $\{v_1, \dots, v_n\}$ is a basis for V defining the isomorphism $\mathbb{R}^n \rightarrow V$, then the

inverse isomorphism $V \rightarrow R^n$ corresponds to the element of $\text{Hom}(V, R^n) \approx V^* \times \dots \times V^*$ defined by the dual basis.

This exercise says, given a vector v , its sequence of coefficients w.r.t the basis $\{v_1, \dots, v_n\}$ is the sequence of values of the dual basis evaluated on v . I.e. if $\{f_1, \dots, f_n\}$ is the basis dual to $\{v_1, \dots, v_n\}$ then $v = f_1(v) \cdot v_1 + \dots + f_n(v) \cdot v_n$.

Eg: The set of polynomials of degree $\leq d$, has as basis the set of $d+1$ monomials $\{1, X, \dots, X^d\}$. The coordinates of the polynomial $a_0 + a_1X + \dots + a_n X^n$ in this basis, are its polynomial coefficients (a_0, a_1, \dots, a_d) . Another basis is the set of $d+1$ polynomials $\{1, (1+X), (1+X+X^2), \dots, (1+X+\dots+X^d)\}$. In this basis, the coordinate vector of 1 is $(1, 0, \dots, 0)$, the coordinate vector of $(1+X)$ is $(0, 1, 0, \dots, 0), \dots$, and the coordinate vector of $(1+X+\dots+X^d)$ is $(0, \dots, 0, 1)$.

Ex: If $\{v_1, \dots, v_n\}$ is a basis of V , and $\{w_1, \dots, w_m\}$ is a basis of W , then $\{(v_1, 0), \dots, (v_n, 0), (0, w_1), \dots, (0, w_m)\}$ is a basis of $V \times W$. Thus $\dim(V \times W) = \dim(V) + \dim(W)$.

Ex. Use the previous exercise to show that $\text{Hom}(R^n, R^m)$ has dimension nm .

Fundamental invariants of linear maps: kernel and rank

Now that we know all finite dimensional vector spaces over the reals, up to isomorphism, namely there is exactly one for each dimension, the next step is to try to understand all linear maps between finite dimensional spaces.

The most basic questions about a linear map are whether it is injective and/or surjective. A first step is to examine which vectors in the source space are sent to zero, which measures the extent to which the map fails to be injective. Another important goal is to determine which vectors in the target space occur as values, which measures the extent to which the map fails to be surjective. These questions lead to the important concepts of kernel, image, and rank.

Defn: 1) If $T: V \rightarrow W$ is a linear map, the **kernel** of $T = \ker(T) = \{v \text{ in } V: T(v) = 0\}$, and

2) The **image** of $T = \text{Im}(T) = \{w \text{ in } W: w = T(v) \text{ for some } v \text{ in } V\}$. The dimension of $\text{Im}(T)$ is called the **rank** of T .

Rmk: Knowing the kernel of a map tells us whether the map is injective. I.e. if $v \neq$

w and $T(v) = T(w)$, then $T(v-w) = 0$ so $(v-w) \neq 0$ is in $\ker(T)$, hence T is injective if and only if $\ker(T) = \{0\}$.

Knowing the rank tells us whether the map is surjective (in finite dimensions), since if $\dim(W)$ is finite, T is surjective if and only if $\text{rank}(T) = \dim(W)$.

Ex: If $T: V \rightarrow W$ is a linear map then

1) $\ker(T)$ is a subspace of V , and $\text{Im}(T)$ is a subspace of W .

2) T is constant on each equivalence class in $V/\ker(T)$.

3) T defines a linear map $[T]: V/\ker(T) \rightarrow W$ sending $[v]$ to $T(v)$.

4) The induced map $[T]$ in 3) is always injective, and $[T]$ is surjective if and only if T is, hence $[T]$ is an isomorphism if and only if T is surjective.

5) A linear map $[T]$ can be defined in the same way on V/U , for any subspace U contained in $\ker(T)$, but $[T]$ will not be injective unless $U = \ker(T)$.

Lemma: If $\dim(V) < \infty$, every independent set in V is contained in a basis.

Pf: If $\{v_1, \dots, v_n\}$ is independent, and $\{w_1, \dots, w_m\}$ is a spanning set, reduce the generating set $\{v_1, \dots, v_n, w_1, \dots, w_m\}$ to a basis, as in the proof that a finite dim'l space has a basis. QED.

Thm: If $\dim(V) < \infty$, and W is a subspace of V , then $\dim W + \dim(V/W) = \dim V$.

Pf/Ex: Choose a basis w_1, \dots, w_s for W , and extend it to a basis

$\{w_1, \dots, w_s, v_1, \dots, v_t\}$ of V . Then $\{[v_1], \dots, [v_t]\}$ is a basis for (V/W) . QED.

Remark: Note this proof gives a procedure for finding a basis of a quotient space.

Thm: If $\dim(V) < \infty$, and $T: V \rightarrow W$ is a linear surjection, $\dim(\ker(T)) + \dim W = \dim V$.

Pf: T induces an isomorphism from $V/\ker(T)$ to W . QED.

Cor: If $\dim(V) < \infty$, and $T: V \rightarrow W$ is a linear map, then $\text{rank}(T) = \dim(V) - \dim \ker(T)$.

Thus knowing the kernel (and $\dim(V)$) also determines the rank, hence just

knowing the kernel, and the dimensions of the source and target, determines both injectivity and surjectivity of the map. In particular a map between two spaces of the same finite dimension is an isomorphism if and only if the kernel is zero.

Cor: $\dim(V \times W) = \dim V + \dim W$

Pf: The projection taking (x,y) to y is a linear surjection from $V \times W$ to W with kernel V . **QED.**

Ex: 1) If $\dim V > \dim W$, no linear map $V \rightarrow W$ is injective, and no linear map $W \rightarrow V$ is surjective.

2) If $S = \{x_1, \dots, x_k\}$ is a subset of V , and $\dim(V) = n$, then any two of the following implies the third: **a)** S is independent, **b)** S spans V , **c)** $k = n$.

3) If $\dim(V) < \infty$, the map $V \rightarrow (V^*)^* = V^{**}$, taking v to “evaluation at v ” is an isomorphism. I.e. v in V goes to the element of $(V^*)^*$ that takes a linear function f in V^* , to the scalar $f(v)$.

4) If V, W are finite dimensional, $T: V \rightarrow W$ is linear and $\dim \ker(T) \leq \dim(V) - \dim(W)$, then equality holds and T is surjective.

Orthogonal complements of subspaces

There are two mutually complementary, or dual, ways of representing a subspace U of a vector space V . One we have emphasized is to give a spanning set, or better, a basis for the subspace U . Once we have such a spanning set, we can easily produce arbitrarily many elements of U just by taking linear combination of the elements we already have. But there is another problem that arises, that of recognizing elements of U . I.e. if we are presented with an arbitrary vector we can ask whether or not it is an element of our subspace U , and that is not immediately obvious just from having a basis. So given a spanning set for a subspace of V we would also like to have linear equations that vanish precisely on members of that subspace. These linear equations are elements of the dual space V^* . It is traditional, for geometric reasons we will discuss later, if f is a function in V^* that vanishes on a subspace U of V , to say that f is “orthogonal” to U .

Definition: If U is a subspace of V , define the orthogonal complement “ U^{\perp} ” of U as the subspace of V^* consisting of all functions that are identically zero on U .

So $f: V \rightarrow k$ is in U^{\perp} if and only if $f(x) = 0$ for all x in U , if and only if U is a

subspace of $\ker(f)$. It follows from our earlier theory that an element f of U^{\perp} also induces a linear function $[f]$ on V/U , so there is a linear map $U^{\perp} \rightarrow (V/U)^*$, taking f to $[f]$. There is also a map back from $(V/U)^*$ to U^{\perp} by composing a linear function $V/U \rightarrow k$ with the natural projection map $V \rightarrow V/U$. Moreover composing this projection with $[f]$ gives us back f , since by definition the induced map $[f]$ is the unique function on V/U whose composition with the projection equals f . For the same reason, if $f: V \rightarrow k$ is the composition of $V \rightarrow V/U$ with a function $g: (V/U) \rightarrow k$, then $[f] = g$, since both compose with the projection to give f . Thus the subspace U^{\perp} of V^* is naturally isomorphic to $(V/U)^*$. This is just a fancy way of saying that functions on V that vanish on U factor naturally through the quotient space V/U . Thus if V is finite dimensional, then from the known fact that $\dim(V) = \dim(U) + \dim(V/U)$, it follows, for every subspace U of V , that $\dim(V) = \dim(U) + \dim(U^{\perp})$. (When working in \mathbb{R}^n , where we can measure angles, this is essentially the fact that \mathbb{R}^n is the direct product of any subspace U with the subspace of vectors perpendicular to all vectors in U .)

Ex. If U, W are subspaces of V , show the set $U+W$ of all sums $x+y$ of elements, with x in U and y in W , is the smallest subspace of V containing both U and W .

Ex. Show for two subspaces U, W of V , that $(U^{\perp} \cap W^{\perp}) = (U+W)^{\perp}$.

The transpose (or dual) of a map

One natural way to specify a subspace is as the image of a map. In this case the space of equations for that subspace arises as the kernel of a related map.

Definition: If $T: V \rightarrow W$ is a linear map, there is a natural map $T^*: W^* \rightarrow V^*$ defined by “preceding by T ”. I.e. if $f: W \rightarrow k$ is a linear function on W , then $(f \circ T): V \rightarrow k$ is a linear function on V , and we define $T^*(f) = f \circ T$. T^* is called the “transpose” or “adjoint” or “dual” map of T .

Since T^* is defined in terms of T , there are close relations between all subspaces naturally associated to the two maps.

Proposition: If $T: V \rightarrow W$ is a linear map with transpose $T^*: W^* \rightarrow V^*$, the kernel of T^* is the space of equations for $\text{Im}(T)$; i.e. $(\text{Im}(T))^{\perp} = \ker(T^*)$.

proof: If f vanishes on $\text{Im}(T)$, i.e. if f is in $(\text{Im}(T))^{\perp}$, then $(T^*f)(x) = f(T(x)) = 0$ for all x in V , so f is in $\ker(T^*)$. Conversely, if f is in $\ker(T^*)$, then $(T^*f)(x) = f(T(x)) = 0$ for all x in V , so f vanishes on $\text{Im}(T)$. **QED.**

Double duals

In an earlier section we have given as an exercise to prove that the double dual $(V^*)^* = V^{**}$ of a finite dimensional space V is naturally isomorphic to the original space V . This reflects the fact that when we write a function's value as $f(x)$, although we normally think of f as the function and x as the argument on which f is evaluated, we could just as well think of x as the function acting on the argument f . I.e. the pair f , and x , determine a number, and we can think of that number as either $f(x)$ or $x(f)$. But this is potentially confusing and it helps to use different notation for x as a function, as opposed to x as a point. E.g. if x belongs to V , we can denote by ev_x the function "evaluation at x ", which takes an element f in V^* to its value on x , namely $ev_x(f) = f(x)$. This defines a linear map $V \rightarrow V^{**}$, and since when V is finite dimensional, V and V^{**} have the same dimension, it suffices to show this is injective to conclude it is an isomorphism. But if x_1 is any non zero element of V , we can extend it to a basis, hence defining an isomorphism from V to \mathbb{R}^n which takes x_1 to e_1 . Consequently the linear function taking a vector to its first coordinate in this isomorphism equals 1 on x , hence evaluation at x is not zero on all elements of V^* . I.e. if x_1, \dots, x_n is a basis of V , and x_1^*, \dots, x_n^* is the dual basis of V^* , then ev_{x_1} is not zero on x_1^* . Hence $V \rightarrow V^{**}$ is injective.

Remark: If V is infinite dimensional then in general V still injects into V^{**} , but the natural injection need not be surjective. In that setting we usually add some structure to V such as an absolute value or "norm", and then restrict consideration to linear functions that are also continuous in that norm. This makes the dual spaces considered smaller and in some cases does recover the isomorphism.

The double dual of a map

Since in finite dimensions we have $V \approx V^{**}$, the question naturally arises whether under this isomorphism, we have $T^{**} = T$? The answer is yes. I.e. if V, W are finite dimensional, and $T: V \rightarrow W$ is linear, then under the isomorphisms $V \approx V^{**}$ and $W \approx W^{**}$, we do have $T = T^{**}$, in the sense that the element ev_x of V^{**} , is taken by T^{**} to the element $ev_{T(x)}$ of W^{**} . Thus, we claim the compositions $V \rightarrow W \rightarrow W^{**}$ and $V \rightarrow V^{**} \rightarrow W^{**}$ are equal, i.e. for all x in V , $T^{**}(ev_x) = ev_{T(x)}$. To check it choose x in V , and g in W^* , and ask if $(T^{**}(ev_x))(g) = (ev_{T(x)})(g)$. The left side is by definition $((ev_x) \circ T^*)(g) = (ev_x)(g \circ T) = (g \circ T)(x)$. But this equals $g(T(x)) = (ev_{T(x)})(g)$. So much for that.

Structure of linear maps up to isomorphisms of source and target

Just as a vector space is determined up to isomorphism by its dimension, we claim

a linear map $T:V \rightarrow W$ between two finite dimensional spaces V, W , is determined, up to isomorphisms of the source V and the target W , by its rank. I.e. a linear map $V \rightarrow W$ can be changed into any other linear map of the same rank by composing with isomorphisms of V and W . More precisely:

Proposition: If $\dim(V), \dim(W)$ are finite, and the linear maps $f, g: V \rightarrow W$ have the same rank, then there exist isomorphisms $h: V \rightarrow V, k: W \rightarrow W$ such that $g = (k \circ f \circ h): V \rightarrow V \rightarrow W \rightarrow W$.

Proof: Recall that a linear map is determined entirely by what it does to a basis, and that we may define it in any way we like on that basis. Moreover, a linear map between spaces of the same dimension, e.g. from a space to itself, is an isomorphism if and only if it takes a basis to a basis.

So assume $\dim(V) = n, \dim(W) = m$, and that f and g both have rank r , so that $\ker(f)$ and $\ker(g)$ both have dimension $n-r$. Choose a basis $\{v_{r+1}, \dots, v_n\}$ of $\ker(f)$, and extend it to a basis $\{v_1, \dots, v_r, v_{r+1}, \dots, v_n\}$ of V . Then do the same for g , i.e. choose a basis $\{u_{r+1}, \dots, u_n\}$ for $\ker(g)$ and extend to a basis $\{u_1, \dots, u_r, u_{r+1}, \dots, u_n\}$ of V . Then sending the u 's to the v 's defines an isomorphism h from V to V . Since this sends the kernel of g to the kernel of f , this already guarantees that every element of $\ker(g)$ will go to zero both under g and $(f \circ h)$.

Claim: The sets $\{f(v_1), \dots, f(v_r)\}$ and $\{g(u_1), \dots, g(u_r)\}$ are both independent in W .

Proof: If we could express zero as a non trivial linear combination of $\{f(v_1), \dots, f(v_r)\}$, the same linear combination of the v 's would map by f to zero. But this would say some linear combination of $\{v_1, \dots, v_r\}$ belongs to $\ker(f)$, hence to the span of the vectors $\{v_{r+1}, \dots, v_n\}$. That would imply the basis $\{v_1, \dots, v_r, v_{r+1}, \dots, v_n\}$ is actually dependent, a contradiction. Thus $\{f(v_1), \dots, f(v_r)\}$ is independent. By a similar argument $\{g(u_1), \dots, g(u_r)\}$ is independent.

Now extend each of these independent sets to bases $\{f(v_1), \dots, f(v_r), x_{r+1}, \dots, x_m\}$, and $\{g(u_1), \dots, g(u_r), w_{r+1}, \dots, w_m\}$ of W . Next define the isomorphism k of W sending these bases to each other in the order given. Then compare what the maps $k \circ f \circ h$ and g do to the basis $\{u_1, \dots, u_r, u_{r+1}, \dots, u_n\}$ of V . Both maps send the second half of the basis, namely $\{u_{r+1}, \dots, u_n\}$ to zero, and both send the vectors $\{u_1, \dots, u_r\}$ to the vectors $\{g(u_1), \dots, g(u_r)\}$. Since both maps agree on a basis they agree everywhere.

QED.

This implies that up to isomorphisms of V and W , any linear map $V \rightarrow W$ of rank r is equivalent to the simplest possible example. Namely we can consider V and W to be \mathbb{R}^n and \mathbb{R}^m , and we can consider the rank r map to be defined on the standard basis $\{e_1, \dots, e_r, e_{r+1}, \dots, e_n\}$ by taking the vectors $\{e_{r+1}, \dots, e_n\}$ all to zero, and taking the first r vectors $\{e_1, \dots, e_r\}$ to the first r standard vectors $\{u_1, \dots, u_r\}$ of the standard basis $\{u_1, \dots, u_m\}$ for \mathbb{R}^m . I.e. up to isomorphisms, a linear map of rank r from \mathbb{R}^n to \mathbb{R}^m , simply projects \mathbb{R}^n onto the span of its first r axes, and then includes these as the first r axes of \mathbb{R}^m .

Indeed this statement follows directly from the proof above. Given a map $T: V \rightarrow W$, choose a basis $\{v_{r+1}, \dots, v_n\}$ of $\ker(T)$, and extend it to a basis $\{v_1, \dots, v_r, v_{r+1}, \dots, v_n\}$ of V . Then $\{T(v_1), \dots, T(v_r)\}$ is independent and hence extends to a basis $\{T(v_1), \dots, T(v_r), x_{r+1}, \dots, x_m\}$ of W . Now using these bases of V and W to provide isomorphisms of V and W with \mathbb{R}^n and \mathbb{R}^m , we obtain a map $\mathbb{R}^n \rightarrow \mathbb{R}^m$ that sends the last $n-r$ standard basis vectors of \mathbb{R}^n to zero, i.e. projects \mathbb{R}^n onto the span of the first r standard basis vectors, and then maps them to the first r standard basis vectors of \mathbb{R}^m .

This is a concrete realization of the fact proved earlier in an exercise, that every linear map $f: V \rightarrow W$ factors as a composition $V \rightarrow V/\ker(f) \rightarrow W$, of the natural projection $V \rightarrow V/\ker(f)$ followed by the induced injection $V/\ker(f) \rightarrow W$.

Definition: Two linear maps $S: V \rightarrow W$, and $T: V \rightarrow W$ between finite dimensional spaces V, W , are called *equivalent* if there are isomorphisms $P: V \rightarrow V$ and $Q: W \rightarrow W$ of the source and target spaces such that $Q \circ S \circ P = T$.

We just proved S and T are equivalent if and only if S, T have the same rank.

Definition: We call $S, T: V \rightarrow W$ *left - equivalent* if $T = Q \circ S$ for some isomorphism Q of W . Similarly, S, T are *right - equivalent* if $T = S \circ P$, for some isomorphism P of V .

These concepts will be examined from a computational point of view in the next chapter, but you will learn something useful by trying the following exercises now.

Ex. Prove S, T are left equivalent iff they have the same kernel, and they are right - equivalent iff they have the same image.

Ex. Prove S, T are left equivalent iff $S^* \cdot T^*$ have the same image, and right

equivalent iff S^*, T^* have the same kernel.

Chapter Two: Computations using matrices

In this chapter we augment our theoretical considerations by showing how to actually carry out the procedures and compute the invariants discussed abstractly in the previous chapter. E.g. we have repeatedly invoked the possibility of reducing a spanning set for a space to a basis, thus computing the dimension of the space, by eliminating vectors that depend on others, and also of extending an independent set to a basis by adding in new vectors. To actually do this in concrete examples, we first represent our abstract vectors as coordinate vectors by means of a basis, and then give a computational procedure, in terms of “elementary row operations”, for reducing a finite collection of coordinate vectors to an independent spanning set. The technique also lets us extend an independent set to a basis.

The same method computes the rank of a linear map $V \rightarrow W$ between two vector spaces over a field k . I.e. we will show that choosing bases in both V and W allows us to represent a linear map $V \rightarrow W$ as a matrix with entries in k . Then row operations performed on this matrix allow us to compute a basis for the image of the map and hence to compute its rank. The same procedure also allows one to compute a basis for the kernel of the map. At bottom, this is nothing but the middle school technique of “eliminating variables”.

As a limitation on the procedure, it is only useful over fields of scalars where one can actually calculate field operations effectively. E.g. it is useful over the rational field, but less so over the real field, where even representing a given real number precisely may require an infinite decimal. Of course if one is content with an approximation, one may in some cases use finite decimals, which are rational numbers. The trouble with this is that once we introduce approximations, we may no longer obtain the correct answer to a computation of the dimension of a space. I.e. one is then obliged to assume that a given vector is in the span of other vectors if it is so within a “small” error. We do not wish to entertain this sort of situation, nor discuss the problems it may pose, (mainly from ignorance).

Quotient fields such as $Q[X]/(X^3 - X + 1)$ can also offer significant difficulties even in the calculation of field operations like multiplicative inverses. A few calculations may however be feasible here: e.g. can you show easily that the inverse of X in this field equals $(1 - X^2)$? Can you use this fact to quickly compute the inverse of $(1 - X)$? But the inverse of $(1 + X^2)$ seems much harder. (The inverse of $(1 + X^2)$ in this field is a polynomial $g(X)$ such that $g(X) \cdot (1 + X^2) - 1$ is

divisible by $(X^3 - X + 1)$.) Another case where calculations seem feasible is for finite fields, which we do not discuss.

Remark: You could actually skip much of this chapter if you don't want to know how to calculate anything and just want to read more theory, but I don't recommend this. I myself was educated like that, completely theoretically, in the 1960's, and only learned to my amazement that these concepts can actually be computed after starting to teach. In this regard I have read a quote attributed to Kaplansky, to the effect that "we [he and Halmos] think and write invariantly, but when the chips are down, we close the door and compute furiously with matrices".

Finding linear relations

The basic problem is to determine when a finite set of coordinate vectors is dependent, by finding an explicit linear relation they satisfy.

Worked example: To show that the vectors $(4, -2, 5)$, $(1, 0, -3)$, $(-2, 2, -11)$ are dependent means finding scalars x, y, z such that $x \cdot (4, -2, 5) + y \cdot (1, 0, -3) + z \cdot (-2, 2, -11) = (0, 0, 0)$. Multiplying through the vectors by the unknown scalars and adding gives $(4x + y - 2z, -2x + 2z, 5x - 3y - 11z) = (0, 0, 0)$. Since equality of vectors means equality of every entry, This is exactly the same as solving the three simultaneous linear equations: $4x + y - 2z = 0$; $-2x + 0y + 2z = 0$; $5x - 3y - 11z = 0$. If we write these in a vertical array they look like this:

$$\begin{array}{l} 4x + y - 2z = 0; \\ -2x + 0y + 2z = 0; \\ 5x - 3y - 11z = 0. \end{array}$$

Now we may recall from school that the technique for solving such equations is to "eliminate variables" by adding and subtracting multiples of some equations from others. The basic rationale is that a simultaneous solution of two equations will also solve the sum of those two equations, and a solution of one equation will also solve a multiple of that equation. If we avoid multiplying by zero, these statements are almost true in reverse; i.e. a solution of a non zero multiple of an equation will also solve the original equation, and a simultaneous solution of an equation E_1 , and the sum of equations $E_1 + E_2$, will also solve both the original equations E_1 and E_2 . In particular, if $a \neq 0$, then the equations E_1 and E_2 have the same solutions as do the equations $a \cdot E_1$ and $(b \cdot E_1 + E_2)$. We also allow interchanging the ordering of any of the equations. Using these operations, we attempt to transform the original

equations into ones that have fewer variables, hence are easier to solve, but that still have the same solutions. We call systems of equations that have the same solutions “equivalent” systems.

In the example above, we can divide through the second equation by 2, and then add 4 times the new second equation to the first equation, and then add 5 times the new second equation to the third equation, getting the new but equivalent system of equations:

$$\begin{aligned}0x + y + 2z &= 0; \\ x + 0y - z &= 0; \\ 0x - 3y - 6z &= 0.\end{aligned}$$

We may omit the variables with coefficient zero, leaving (after reordering):

$$\begin{aligned}x \quad -z &= 0 \\ \quad y + 2z &= 0 \\ -3y - 6z &= 0\end{aligned}$$

Now adding 3 times the second equation to the third equation leaves:

$$\begin{aligned}x \quad -z &= 0 \\ \quad y + 2z &= 0 \\ \quad 0 \cdot z &= 0.\end{aligned}$$

which we may write as:

$$\begin{aligned}x \quad -z &= 0 \\ \quad y + 2z &= 0\end{aligned}$$

The third equation is now true for every z , or missing entirely, and the first two equations can be solved for y and x , no matter what z is, so we may take z to be anything, such as $z = 1$. Then we have $x = z = 1$, $y = -2z = -2$, and we can check these work. Any multiple of these also works, so we have as solutions any x,y,z of form $(x,y,z) = (z, -2z, z)$ for any choice of z .

Notice it would have been ok for the first equation also to have involved y , since after finding y from the second equation we could have plugged in the known values of both z and y into the first equation to find x . The essential thing is that the second equation not involve x , and the third equation not involve either x or y ,

but the our final form is preferable, where only one equation involves x and only one involves y . We could still have solved the system if the third equation involved z as well, but then the only solution would have been $x = y = z = 0$. Since we were looking for a non trivial linear relation, that result would have meant that no such relation existed, and hence that our three column vectors were actually independent. I.e. recall that vectors are independent, by definition, if the only linear relation they satisfy is the trivial one.

So our method determines whether or not the column vectors are independent, and when they are dependent, it finds all possible non trivial linear relations among them. I.e., it always finds all possible linear relations among the columns, and the columns are independent if and only if the only such relation is the trivial one.

Taking $z = 1$, hence $(x,y,z) = (1,-2,1)$, our solution thus yields the linear relation: $(4, -2, 5) - 2 \cdot (1, 0, -3) + (-2, 2, -11) = (0, 0, 0)$. Consequently any of these three vectors can be expressed as a linear combination of the other two. Since no one of them is a scalar multiple of any other one, the subspace the three vectors span in \mathbb{R}^3 is 2- dimensional, i.e. a “2-plane” or just a “plane”. Any two of them are independent and form a basis for that plane. E.g. $(4, -2, 5)$ and $(1, 0, -3)$ form a basis for that plane, and $(-2, 2, -11) = 2 \cdot (1, 0, -3) - (4, -2, 5)$ lies in that plane.

It is more efficient in such calculations to display the coefficients in a rectangular array or “matrix”, i.e. without the variables. Since the numbers on the right side of the equations are all zeroes, we do not need to write them. The original coordinate vectors appear as the “columns” of the matrix, so we call them “column vectors”.

$$\begin{array}{ccc|c} 4 & 1 & -2 & \\ -2 & 0 & 2 & \\ 5 & -3 & -11 & \end{array}$$

Then the previous transformations give the following sequence of equivalent matrices:

$$\begin{array}{ccc|c} 4 & 1 & -2 & \\ -1 & 0 & 1 & \\ 5 & -3 & -11 & \end{array} \approx$$

$$\begin{array}{ccc|c} 0 & 1 & 2 & \\ -1 & 0 & 1 & \\ 5 & -3 & -11 & \end{array} \approx$$

$$\begin{array}{ccc|c} 0 & 1 & 2 & \\ -1 & 0 & 1 & \\ 0 & -3 & -6 & \approx \end{array}$$

$$\begin{array}{ccc|c} 0 & 1 & 2 & \\ -1 & 0 & 1 & \\ 0 & -3 & -6 & \approx \end{array}$$

$$\begin{array}{ccc|c} -1 & 0 & 1 & \\ 0 & 1 & 2 & \\ 0 & -3 & -6 & \approx \end{array}$$

$$\begin{array}{ccc|c} 1 & 0 & -1 & \\ 0 & 1 & 2 & \\ 0 & -3 & -6 & \approx \end{array}$$

$$\begin{array}{ccc|c} 1 & 0 & -1 & \\ 0 & 1 & 2 & \\ 0 & 0 & 0 & . \end{array}$$

Now we put back the variables, and the zeroes, and solve the equivalent system

$$\begin{array}{l} x - z = 0 \\ y + 2z = 0, \end{array}$$

with solutions $(x, y, z) = (z, -2z, z)$ for every choice of z .

Note: Since the solutions thus consist of all scalar multiples $z \cdot (1, -2, 1)$ of the vector $(1, -2, 1)$, the solutions form a one dimensional subspace of \mathbb{R}^3 , with basis $(1, -2, 1)$.

Another worked example.

Given this set of four vectors in \mathbb{R}^3 : $(3, -1, 4)$, $(6, 8, -2)$, $(3, 9, -6)$, $(0, -10, 10)$, try to decide whether they are independent, and if not, find a non trivial linear relation.

Matrices associated to this system are:

$$\begin{array}{cccc|c} 3 & 6 & 3 & 0 & \\ -1 & 8 & 9 & -10 & \\ 4 & -2 & -6 & 10 & \approx \text{(subtract row 1 from row 3)} \end{array}$$

$$\begin{array}{cccc|c} 3 & 6 & 3 & 0 & \\ -1 & 8 & 9 & -10 & \end{array}$$

$$\begin{vmatrix} 1 & -8 & -9 & 10 \end{vmatrix} \approx (\text{add row 2 to row 3})$$

$$\begin{vmatrix} 3 & 6 & 3 & 0 \\ -1 & 8 & 9 & -10 \\ 0 & 0 & 0 & 0 \end{vmatrix} \approx (\text{divide row 1 by 3})$$

$$\begin{vmatrix} 1 & 2 & 1 & 0 \\ -1 & 8 & 9 & -10 \\ 0 & 0 & 0 & 0 \end{vmatrix} \approx (\text{add row 1 to row 2})$$

$$\begin{vmatrix} 1 & 2 & 1 & 0 \\ 0 & 10 & 10 & -10 \\ 0 & 0 & 0 & 0 \end{vmatrix} \approx (\text{divide row 2 by 10})$$

$$\begin{vmatrix} 1 & 2 & 1 & 0 \\ 0 & 1 & 1 & -1 \\ 0 & 0 & 0 & 0 \end{vmatrix} \approx (\text{subtract 2 times row 2 from row 1})$$

$$\begin{vmatrix} 1 & 0 & -1 & 2 \\ 0 & 1 & 1 & -1 \\ 0 & 0 & 0 & 0 \end{vmatrix}. \text{ Thus if the variables are } x,y,z,w, \text{ we get:}$$

that z, w can be anything, and $x = z - 2w, y = w - z$.

Note: Since the solutions thus have form $(z-2w, w-z, z, w)$ where z,w can be any real numbers, again the set of solutions forms a subspace since it consists of all linear combinations $z \cdot (1, -1, 1, 0) + w \cdot (-2, 1, 0, 1)$ of the vectors $(1, -1, 1, 0)$, and $(-2, 1, 0, 1)$. I.e. the solutions form a two dimensional subspace of \mathbb{R}^4 with basis $(1, -1, 1, 0), (-2, 1, 0, 1)$.

The key point of course is that the solutions of the reduced system are also solutions of the original system. Thus taking $z = 1, w = 0$, shows the third original column vector depends on the first two since then $x = 1, y = -1$ gives the relation $(3, -1, 4) - (6, 8, -2) + (3, 9, -6) = (0, 0, 0)$ so $(3, 9, -6) = (6, 8, -2) - (3, -1, 4)$.

Taking $z = 0, w = 1$ shows similarly that the 4th original column vector also depends linearly on the first two. To be sure let's check that. I.e. then $x = -2, y = 1$, so we should get $-2 \cdot (3, -1, 4) + (6, 8, -2) + (0, -10, 10) = (0, 0, 0)$, which does check, so solving gives $(0, -10, 10) = 2 \cdot (3, -1, 4) - (6, 8, -2)$, thus expressing the fourth vector linearly in terms of the first two.

So the subspace spanned by all four original column vectors is 2 dimensional, with the first two vectors $(3, -1, 4)$ and $(6, 8, -2)$ as basis. Of course there are other choices of basis for this subspace, (in fact probably any choice of two of these four vectors should work, as long as neither is a multiple of the other), but this gives a way of choosing one. I.e. this method essentially throws out column vectors that depend on earlier ones, hence chooses the “earliest two” independent vectors among the four, in their original ordering.

Note: Since the rows of the reduced form of a matrix are linear combinations of the original rows, and vice versa, the rows of the reduced matrix, which are obviously independent, form a basis of the space spanned by the original rows.

Thus given any matrix, the computation we have illustrated gives a way to find a basis for the space spanned by its columns, (its “column space”), the space of its solutions, (its “null space”), and the space spanned by its rows, (its “row space”). Notice also that although a matrix and its reduced form have the same row space and null space, they usually do *not* have the same column space. Indeed the column space of the original n by m matrix, can be any subspace of \mathbb{R}^n , but if the reduced form has r non zero rows, then its column space is just the coordinate subspace of \mathbb{R}^n spanned by the first r standard basis vectors.

Remark: Don’t feel bad if you did these exercises and got them wrong. So did I, repeatedly, even though they are about as simple as they come. These pesky real life computations are notoriously prone to error, especially if you yield to the temptation to take shortcuts and do steps in your head. They look so easy, but it is really easy also to divide 6 by 3 mentally and write down 3 instead of 2, as I originally did in reducing the first row of the second example above. The best advice I can give is to force yourself to write out every step, no matter how boring it seems, or risk having to do it all over again. Even if you are careful, you may need to repeat it several times. In my case I knew I had rigged the examples above to have 2 dimensional spans, but I kept getting 3 dimensional answers over and over, until I just plodded through every little step one at a time. Of course once you get an answer you can check it, and see if it does give you a valid relation among your original vectors, and you should always do this. The way I knew I was wrong even sooner, was that the number of dimensions should equal the number of non zero rows that are left when you are done, but I kept getting something like this, which has three non zero rows:

$$| 1 \ 0 \ 0 \ 1 |$$

$$\begin{array}{l} | 0 \quad 1 \quad 0 \quad 0 | \\ | 0 \quad 0 \quad 1 \quad -1 | \end{array}$$

The wonderful fact that we now have computers to do our calculations means that once you understand how to do these procedures, you can program a computer to carry them out accurately. Of course it doesn't hurt to practice your own accuracy. I once did a complicated partial fractions integral by hand that my \$800 copy of Mathematica had failed to do for some reason. Computers sometimes also get the wrong answer because we typed in an extra space or left out a comma, or did something else quite invisible, so we need to have some intuition as to when the answer they give us is absurdly wrong. So please do some of these computations yourself until you get them right. It will make you feel good and also give you a better grasp on the abstract ideas. So the fact these computations are hard to get right by hand does not invalidate their usefulness.

It also helps to give a precise description of the computational process we have illustrated, learn the standard language for discussing it, and restate the basic facts.

Elementary row operations:

Given a matrix, here are the allowed "row operations":

- 1) interchange any two rows.
- 2) multiply through any row by a non zero scalar.
- 3) add a multiple of any row to any other row.

Reduced echelon form

Here is the type of matrix you want to end up with:

- 1) all zero rows are together at the bottom of your matrix.
- 2) in all non zero rows the first (leftmost) non zero entry is a '1'
- 3) that first '1' in a non zero row is the only non zero entry in its column.
- 4) the initial '1' in a non zero row occurs further to the right than the initial '1' in the previous row.

A matrix in this form is said to be in "reduced echelon form". We will see that this form depends uniquely on the original matrix. I.e. no matter what sequence of steps you use to reduce a matrix, the final result of this type will always be the same. In particular if you rework the two examples above you should get the same final result, no matter what different intermediate stages you go through.

Terminology: In each non zero row of a reduced echelon matrix, the column where the first non zero entry appears is called a "pivot" column. If the matrix is n

by m , and the associated system of equations has variables x_1, \dots, x_m , the variable assigned to a pivot column is called a “pivot variable”.

Note: In the worked examples, we saw that the set of solutions of the associated homogeneous linear system of equations, was equal to the number of non - pivot variables.

It is useful to specify some subspaces associated to a matrix as follows.

Row space:

Definition: Given an n by m matrix A , the subspace of \mathbb{R}^m spanned by the rows is called the “row space of A ”, denoted $R(A)$.

Since there are n rows, the row space has dimension at most n . In fact we claim the rows of the reduced echelon form are a basis of the row space, so the row space of a matrix has dimension equal to the number of non - zero rows in its reduced echelon form. We can see this as follows.

Looking at a reduced echelon matrix, it should be obvious that the non zero rows are independent, so it suffices to prove that the row space of a matrix is the same as the row space of its reduced echelon form. For this it suffices to see that the row space of a matrix is unchanged by performing an elementary row operation. By definition of the elementary row operations, each row in the transformed matrix is a linear combination of rows in the previous matrix. But since each row operation can be reversed by another opposite row operation, also every row in the previous matrix is a linear combination of those in the transformed matrix. E.g. we can reverse the action of multiplying by a non zero scalar by dividing by that non zero scalar, i.e. by multiplying by its inverse. And if we replace a row E_j by the linear combination $E_j + c.E_k$ where $j \neq k$, then we can recover the original row E_j by adding to $E_j + c.E_k$, the multiple $-c.E_k$. We have proved the following:

Proposition: The row space of a matrix has as basis the set of non zero rows in its reduced echelon form. We denote the row space of A as $R(A)$.

Terminology: The “row rank” of a matrix is the dimension of its row space.

As was probably clear from the examples, any matrix has a reduced echelon form, and we want to state that explicitly, as well as the fact that it is unique.

Theorem: Any matrix can be put into reduced echelon form, by a finite sequence of elementary row operations. The final reduced echelon form is uniquely determined by the original matrix and does not depend on the particular choice of

the sequence of operations used.

proof sketch:

Existence: One need only convince oneself the procedure does work. E.g. starting from the left, locate the first column having a non zero entry and exchange rows until that entry is in the top row. Multiply by the inverse of that entry to get a '1' as the upper left most entry. Now subtract suitable multiples of that row from all other rows until there are only zeroes elsewhere in that column. This the first pivot column. If the only non zero entries remaining in the matrix now are all in the first row, you are finished, and the row rank is one. Notice also that projection of the first row into the coordinate axis associated to the first pivot variable sends the first row to the standard basis vector spanning that coordinate axis. Hence if the row rank is one, the row space projects isomorphically to that coordinate axis and the first row vector corresponds under that isomorphism to the standard basis of that coordinate axis. Note also that the first row does not project isomorphically to any earlier coordinate axis, since all earlier entries in the first row are zero.

If there are some non zero entries below the first row, then looking only at entries below the first row, find the first column in which such a non zero entry occurs.

This is the second pivot column and must lie to the right of the first one.

Interchange rows so that new non zero entry is in the second row. Now divide through by the inverse to make that entry a '1' and then subtract multiples of that second row from all other rows until every other entry in that second pivot column is equal to zero. If there are no non zero entries remaining below the second row, you are finished, and the row rank is two. Notice that if the row rank *is* two, then projection onto the coordinate plane corresponding to the two pivot variables is an isomorphism and the first two rows project under that isomorphism to the two standard basis vectors for that coordinate plane. Note also that the row space does not project isomorphically to any earlier coordinate plane since all earlier entries in the second row are zero.

If there are some non zero entries below the second row, continue.....

Uniqueness:

The reduced echelon form reduces the original sequence of rows to a sequence which is independent, hence forms a basis for the row space. So finding the reduced echelon form amounts to choosing an especially nice or simple basis for the row space. From this perspective the process is very simple. If the rank is r , then there must exist some coordinate subspace of rank r onto which the row space projects isomorphically. If we find one, we can pull back the standard basis for this coordinate subspace to a basis of the row space via this isomorphism. This provides at most " m choose r " distinguished bases for the row space of an n by m

matrix of rank r . The only one in echelon form is the one whose choice of r coordinates among the m possible coordinates, is “lexicographically minimal”. This provides a unique choice of ordered basis for the row space, hence the reduced echelon form of a given matrix is unique. I.e. its pivots correspond to the unique lexicographically minimal choice of coordinate subspace onto which the row space projects isomorphically. That is, for each $j \leq r$, the row space projects surjectively onto the coordinate j -plane spanned by the first j pivot columns, but does not surject onto any earlier coordinate j - plane. Then the actual rows of the reduced matrix are the unique vectors in the row space corresponding under this isomorphism to the standard basis of the coordinate r - plane spanned by all the pivot columns.

QED.

Corollary: Two matrices are row equivalent, i.e. differ by a sequence of elementary row operations, if and only if they have the same row space.

Proof: We have seen that two row equivalent matrices do have the same row space. Moreover the uniqueness proof above showed that the reduced echelon form of a matrix is entirely determined by its row space. Thus two matrices with the same row space have the same reduced echelon form, and hence differ by a sequence of elementary row operations. **QED.**

Remark: The step by step algorithm in this existence proof sketch will always work, but in practice you will usually see shortcuts that make the work go faster. Remember however to be *very* methodical or, if you are like me, you will make numerous arithmetical errors. In fact, because I am usually considered to be much more accurate at mental arithmetic than average, even among mathematicians, I was quite surprised that I make so many errors in these matrix operations.

Exercise: Reduce this one:

$$\begin{array}{cccc|c} 4 & 6 & 2 & -2 & \\ \hline \end{array}$$

$$\begin{array}{cccc|c} 2 & -1 & 5 & 3 & \\ \hline \end{array}$$

$$\begin{array}{cccc|c} -2 & 0 & -4 & -2 & \\ \hline \end{array}$$

$\begin{array}{cccc|c} 5 & 3 & 7 & 2 & \\ \hline \end{array}$ and see if you agree with me that the reduced form is this:

$$\begin{array}{cccc|c} 1 & 0 & 2 & 1 & \\ \hline \end{array}$$

$$\begin{array}{cccc|c} 0 & 1 & -1 & -1 & \\ \hline \end{array}$$

$$\begin{array}{cccc|c} 0 & 0 & 0 & 0 & \\ \hline \end{array}$$

$\begin{array}{cccc|c} 0 & 0 & 0 & 0 & \\ \hline \end{array}$. You should check it by seeing whether you do get valid linear relations among your original column vectors.

To make up your own examples, start with any two independent column vectors and then choose the third column vector to be some simple linear combination of the first two. Then you should end up with a reduced matrix that has exactly two non zero rows. To check your accuracy do it again by different steps and see if you get the same result. Or take a matrix with two or three rows, reduce it, then begin again, but interchange the rows before reducing it again. The result should be the same reduced matrix.

Excursion: the Grassmann variety $G(1,3)$ of projective lines in P^3

Consider “projective” 3 space P^3 defined as all one dimensional subspaces of 4 dimensional coordinate vector space R^4 . So a point of P^3 is represented by a non zero vector in R^4 , and proportional vectors represent the same point. Thus projective “lines” in P^3 are represented by 2-dimensional subspaces of R^4 , and projective planes correspond to 3 - dimensional subspaces of R^4 . We want to describe the set $G(1,3)$ of all projective lines in P^3 , i.e. of all 2 dimensional subspaces of R^4 , by means of reduced 2 by 4 echelon matrices of rank 2. I.e. by our uniqueness theorem, two such matrices have the same row space if and only if they have the same reduced echelon form, so assigning to each reduced 2 by 4 echelon matrix of rank 2, its row space, sets up a 1-1 correspondence between such reduced matrices and all 2 dimensional subspaces of R^4 .

We claim the space $G(1,3)$ has a natural geometric “stratification”, i.e. is a disjoint union of copies of affine spaces of various dimensions. In fact, $G(1,3)$ has a natural disjoint decomposition into 6 subsets isomorphic to R^4 , R^3 , R^2 , R^2 , R^1 , and R^0 (a point). Moreover, viewing the space $G(1,3) =$ all 2 dimensional subspaces of R^4 , as the space of all possible reduced 2 by 4 echelon matrices of rank 2, we claim the six disjoint subsets correspond to subdividing reduced echelon matrices according to the 6 possible ways of choosing the location of the two pivot columns, from among the 4 possible columns. We explain that next.

Since points of P^3 are one dimensional subspaces of R^4 , it follows that for two subsets of P^3 to meet in at least one point, the corresponding sets of vectors in R^4 must share at least one non zero vector. In particular, since every 3 dimensional subspace of R^4 must meet every 2 dimensional subspace non trivially, it follows that in P^3 every line in P^3 meets every plane. Since in general, two 2 dimensional subspaces of R^4 will meet only in $\{0\}$, it follows that two general lines in P^3 will not meet. Hence to require a line to meet a given line or a given point, or to lie in a given plane in P^3 , does impose a restriction. But any two lines that lie in the same plane in P^3 must meet, since a plane in P^3 is a 3 dimensional subspace of R^4 , and any two 2 dimensional subspaces of the same

3 dimensional subspace must have a common non zero vector. So this reveals the special feature of projective geometry, any two lines in a plane always meet, and any line in 3 space meets every plane. I.e. the fact that in \mathbb{R}^4 all subspaces contain $\{0\}$ is reflected in the absence of “parallelism” in \mathbb{P}^3 .

Now take as coordinates (X, Y, Z, W) in \mathbb{R}^4 , called “homogeneous” coordinates for \mathbb{P}^3 , (since the coordinate vectors (a, b, c, d) and (ta, tb, tc, td) represent the same point of \mathbb{P}^3). Then choose a nested family of sets:

the (projective) “plane” $\Pi: X = 0$,

the “line” $L: X = Y = 0$, and

the “point” $P: X=Y=Z=0$.

Note that P lies on L , which lies in Π . (The homogeneous coordinates of P may be taken as $(X, Y, Z, W) = (0, 0, 0, 1)$, since any other vector of form $(0, 0, 0, d)$ with $d \neq 0$ defines the same point).

This gives us several natural subsets of projective lines defined by “incidence” relations, i.e. by how they meet the given point P , the given line L , and the given plane Π . So consider the special subsets of:

- 1) lines in \mathbb{P}^3 that meet the line L ,
- 2) lines in \mathbb{P}^3 that meet the point P , and
- 3) lines in \mathbb{P}^3 that lie in the plane Π .

These sets are not disjoint, since all lines of types 2 or 3 are also of type 1, but they allow us to define six disjoint subsets as complements of them as follows:

- i) lines that do not meet L ,
- ii) lines that meet L but not P , and do not lie in Π ,
- iii) lines that meet P , but do not lie in Π ,
- iv) lines that lie in Π but do not meet P ,
- v) lines that lie in Π and meet P but do not equal L , and finally
- vi) the one line L .

We claim this stratification of $G(1, 3)$ corresponds to the stratification of rank 2 reduced echelon 2 by 4 matrices by the 6 ways to choose the two pivot columns.

Recall that the pivot variables are those in the earliest coordinate plane onto which the row space projects isomorphically. Since projection from the row space onto a given coordinate plane is isomorphic if and only if the “center” of projection, i.e.

the kernel of the linear projection map, does not meet the row space non trivially, this lets us interpret the location of pivot columns in terms of incidence relations.

E.g. the kernel of projection onto the first two coordinates X, Y is the space $X=Y=0$. Thus projection onto the first 2 coordinates is isomorphic if and only if the row space does not meet the subspace $X=Y=0$ non trivially, if and only if neither row of the reduced echelon form has zeroes in the first two columns, iff the first two columns are pivots.

Thus the sets of lines in P^3 which do not meet the line $L: X=Y=0$, corresponds exactly to the reduced echelon matrices of form:

$$\begin{array}{c|cc} 1 & 0 & a & b \\ 0 & 1 & c & d \end{array}$$

Thus the set i) is 4 dimensional, parametrized by (a,b,c,d) , i.e. by R^4 .

Now what about lines that do meet L , but do not meet P and do not lie in Π ? Since the row space does meet $X=Y=0$ non trivially, X and Y are not both pivots. But since the line does not lie in the plane $\Pi: X=0$, the first column, namely X , must be a pivot. So the other pivot is either the third or the 4th column. But the row space does not meet the point $P: X=Y=Z=0$, so the second row cannot have form $(0,0,0,1)$ so the 4th column, i.e. W , is not a pivot. Thus the two pivots are X and Z .

The set of such lines hence corresponds to reduced echelon matrices of form:

$$\begin{array}{c|ccc} 1 & a & 0 & b \\ 0 & 0 & 1 & c \end{array}$$

Thus the set ii) therefore has dimension 3, parametrized by (a,b,c) , i.e. by R^3 .

Lines that do not lie in $\Pi: X=0$ again correspond to matrices with X as a pivot variable, and if they do meet $P: X=Y=Z=0$ non trivially, then W must be a pivot, so the pivots are X and W . These matrices then have form:

$$\begin{array}{c|ccc} 1 & a & b & 0 \\ 0 & 0 & 0 & 1 \end{array}$$

and this set iii) is hence two dimensional, parametrized by (a,b) , i.e. by R^2 .

Lines that do lie in Π , but do not meet P , similarly must not have X or W as pivots, hence must have Y, Z as pivots. These matrices have form:

$$\begin{array}{c|cc} 0 & 1 & 0 & a \\ 0 & 0 & 1 & b \end{array}$$

and the set iv) thus also has dimension two, parametrized by (a,b) , i.e. by \mathbb{R}^2 .

Lines that lie in Π do not have X as pivot, and if they do meet P they must have W as pivot. But if they are not equal to the line $L: X=Y=0$, they must have some non zero entry in the first two columns, so the pivots are Y,W . These have form:

$$\begin{array}{c} |0 \ 1 \ a \ 0| \\ |0 \ 0 \ 0 \ 1| \end{array}$$

and hence set v) is one dimensional, parametrized by a in \mathbb{R}^1 .

Finally the one line vi), $L: X=Y=0$ is the unique matrix

$$\begin{array}{c} |0 \ 0 \ 1 \ 0| \\ |0 \ 0 \ 0 \ 1| \end{array}$$

so vi) is of course zero dimensional, the single line L .

In fact $G(1,3)$ can be described more symmetrically using reduced but not echelon matrices, i.e. matrices with the same row space but in which the choice of pivot columns is any pair of coordinate variables onto which the row space projects isomorphically, not necessarily the earliest pair. Then each matrix has more than one reduced form, and each choice of pivot variables will form a 4 dimensional set, like case i) above. E.g. those with pivots in columns 1 and 3, now look like:

$$\begin{array}{c} |1 \ a \ 0 \ b| \\ |0 \ c \ 1 \ d| \end{array}$$

This just means that in carrying out row reduction, after finding the first pivot, we chose as the next pivot, any non zero column entry occurring below the first row, not necessarily the first such column entry.

This shows that $G(1,3)$ can be covered by a collection of 6 copies of \mathbb{R}^4 , in particular $G(1,3)$ is a 4 dimensional “manifold”. In general, the space $G(r-1,n-1)$ of r dimensional subspaces of \mathbb{R}^n is a manifold of dimension $r(n-r)$, as you can check by using the same ideas as in this excursion, i.e. rank r , reduced r by n echelon matrices. (Hint: what does a reduced r by n echelon matrix look like whose pivots are the first r columns?) Such spaces are quite interesting in geometry and topology. (Be careful in reading further, since in books where the perspective of projective geometry is not used, this space is denoted $G(r,n)$.)

End of excursion.

Matrices and linear maps

The most common use for matrices is to represent linear maps. Just as every linear map $T: \mathbb{R} \rightarrow \mathbb{R}$ is multiplication by some number, namely $x \mapsto T(x) = ax$, where $a = T(1)$, so every linear map $T: \mathbb{R}^n \rightarrow \mathbb{R}^m$ can be represented as multiplication by a matrix A of numbers, where the columns of A are the values of T on the standard basis vectors of \mathbb{R}^n . I.e. the j th column of A is the vector $T(e_j)$ in \mathbb{R}^m .

First we show how to multiply vectors, and then we use this to define matrix multiplication, one row vector at a time.

Multiplying a vector by a vector

Definition: The “inner product” or “dot product” $v \cdot w$ of two vectors $v = (a_1, \dots, a_n)$ and $w = (b_1, \dots, b_n)$ in \mathbb{R}^n is defined as: $v \cdot w = (a_1, \dots, a_n) \cdot (b_1, \dots, b_n) = a_1 b_1 + \dots + a_n b_n$. It is a real number.

Ex. Dot product is commutative and behaves like a multiplication: i.e. it distributes over vector addition, and commutes with scalar multiplication in each variable.

More precisely, for all vectors u, v, w , and scalars t , we have

$$v \cdot w = w \cdot v,$$

$$(u+v) \cdot w = u \cdot w + v \cdot w, \text{ and}$$

$$(tu) \cdot v = t(u \cdot v) = u \cdot (tv).$$

Cor: If $v = (a_1, \dots, a_n)$ is any vector in \mathbb{R}^n , v defines a linear map $f: \mathbb{R}^n \rightarrow \mathbb{R}$ by $f(w) = v \cdot w$.

Ex. In fact, the map $\mathbb{R}^n \rightarrow (\mathbb{R}^n)^*$ taking v to $v \cdot ()$ is a linear isomorphism. In fact it takes the standard basis of \mathbb{R}^n to the dual basis for $(\mathbb{R}^n)^*$; i.e. if e_1, \dots, e_n is the standard basis of \mathbb{R}^n , then $e_j \cdot ()$ is the linear function with value 1 on e_j and value zero on the other e_i .

Since we know that \mathbb{R}^n and $(\mathbb{R}^n)^*$ have the same dimension, it suffices to check that the map in the previous exercise is a linear injection, but one can also give an inverse map as described next. The point is to show that every linear function on \mathbb{R}^n is defined by dotting with some vector.

Ex. A linear map $f: \mathbb{R}^n \rightarrow \mathbb{R}$, is equal to dotting with $v = (f(e_1), \dots, f(e_n))$. i.e. if $w = (b_1, \dots, b_n) = b_1 e_1 + \dots + b_n e_n$, then $f(w) = b_1 f(e_1) + \dots + b_n f(e_n)$.

Since we already know from chapter one that $\text{Hom}(\mathbb{R}^n, \mathbb{R}^m) \approx (\mathbb{R}^n)^* \times \dots \times (\mathbb{R}^n)^*$, m factors, it follows that every linear map $\mathbb{R}^n \rightarrow \mathbb{R}^m$ is given by a sequence of m dot products. This gives the representation of a map by “matrix multiplication”.

Multiplying a matrix by a vector

An “ m by n ” matrix A is an array of m row vectors v_1, \dots, v_m , each row vector belonging to \mathbb{R}^n . We will define a multiplication by A taking column vectors of length n to column vectors of length m . Thus given a column vector w of length n , we may define $A \cdot w$ to be the column vector of length m whose j th entry is $v_j \cdot w$. I.e. start with a vector w in \mathbb{R}^n , and dot w with every row of A , obtaining m numbers, which are the entries of the vector $A \cdot w$. Since each dot product by a row is a linear map $\mathbb{R}^n \rightarrow \mathbb{R}$, we thus obtain a linear map $\mathbb{R}^n \rightarrow \mathbb{R}^m$, taking w to $A \cdot w$. If $\text{Mat}(m, n) =$ the space of all m by n matrices, this defines a map $\text{Mat}(m, n) \rightarrow \text{Hom}(\mathbb{R}^n, \mathbb{R}^m)$, which we will see is a linear isomorphism.

Ex. $\text{Mat}(m, n)$ is a vector space of dimension $m \cdot n$, with basis all matrices having a ‘1’ in a single entry and all other entries zero.

Ex. $\text{Mat}(m, n) \approx \text{Hom}(\mathbb{R}^n, \mathbb{R}^m)$.

If A, B belong to $\text{Mat}(m, n)$ and w belongs to \mathbb{R}^n , then not only is $A \cdot (w_1 + w_2) = A \cdot w_1 + A \cdot w_2$, and $A \cdot (c w) = c \cdot (A \cdot w)$, but also $(A + B) \cdot w = A \cdot w + B \cdot w$, and $(cA) \cdot w = c \cdot (A \cdot w)$. Thus the map $\text{Mat}(m, n) \rightarrow \text{Hom}(\mathbb{R}^n, \mathbb{R}^m)$ is linear. Prove it is also injective, hence isomorphic, since both spaces have dimension mn .

It is useful to spell out the inverse of the map $\text{Mat}(m, n) \rightarrow \text{Hom}(\mathbb{R}^n, \mathbb{R}^m)$, thus giving another proof that $\text{Mat}(m, n) \approx \text{Hom}(\mathbb{R}^n, \mathbb{R}^m)$. We do this next.

How to represent any linear map $\mathbb{R}^n \rightarrow \mathbb{R}^m$ by matrix multiplication

Given a linear map $T: \mathbb{R}^n \rightarrow \mathbb{R}^m$, arrange the image vectors $T(e_1), \dots, T(e_n)$ as columns in a rectangular matrix A . Then there are m rows and n columns. If $v = (a_1, \dots, a_n)$ is any vector in \mathbb{R}^n , then $T(v) = a_1 T(e_1) + \dots + a_n T(e_n)$, is the linear combination of the columns of A having the coordinates of v as coefficients. Thus the i th entry of $T(v)$ is obtained by dotting v with the i th row of A .

Thus $T(v)$ can be computed by multiplying A by v as follows: write v as a length n column vector to the right of A . The product Av is a length m column vector, where the i th entry of Av is the dot product of the i th row of A with v . Thus each linear map from \mathbb{R}^n to \mathbb{R}^m is represented by multiplying by a (unique) m by n matrix.

Eg: The matrix of the map $\mathbb{R}^2 \rightarrow \mathbb{R}^2$ defined by $T(v) = 6v$, has rows (and columns): $(6,0)$ and $(0,6)$. The matrix of the rotation map of \mathbb{R}^2 counter clockwise through $\pi/2$ radians has columns $(0,1)$ and $(-1,0)$.

Ex: Find the matrix of the reflection map of \mathbb{R}^2 in the line spanned by $(1,0)$, and the matrix for counter clockwise rotation about $(0,0)$ through t radians.

Ex: 1) The space of all m by n matrices forms a vector space $\text{Mat}(m,n)$ where $A+B$ is the matrix whose (i,j) entry, i.e. the entry in the i^{th} row and j^{th} column, is the sum of the (i,j) entries of A and B , and where cA is the matrix whose (i,j) entry is c times the (i,j) entry of A .

2) The space $\text{Hom}(\mathbb{R}^n, \mathbb{R}^m)$ is isomorphic to the space $\text{Mat}(m,n)$, (note the indices n,m occur correctly in the reverse order here).

3) The dimension of $\text{Mat}(m,n)$, hence that of $\text{Hom}(\mathbb{R}^n, \mathbb{R}^m)$, is mn .

More fundamental subspaces associated to a matrix

Definition: Given an m by n matrix A , the “nullspace” of A , written $N(A)$, is the kernel of the associated map $A: \mathbb{R}^n \rightarrow \mathbb{R}^m$. It is thus a subspace of \mathbb{R}^n .

These are exactly the solution vectors of the system of homogeneous equations associated to the matrix A . I.e. w in \mathbb{R}^n belongs to the null space of A if and only if $v_j \cdot w = 0$ for every row vector v_j of A , if and only if $v \cdot w = 0$ for every vector v in the row space of A . This gives an important viewpoint relating the nullspace and row space of a matrix.

Definition: Two vectors v,w in \mathbb{R}^n are called “orthogonal” if and only if $v \cdot w = 0$. Given any subset S of \mathbb{R}^n , the “orthogonal complement” of S , called Sperp , is the set of all vectors w in \mathbb{R}^n which are orthogonal to all vectors in S .

Note: Under the isomorphism $\mathbb{R}^n \rightarrow (\mathbb{R}^n)^*$ noted above, taking a vector v to “dotting with v ” or $v \cdot ()$, this subspace Sperp of \mathbb{R}^n maps isomorphically to the subspace Sperp of $(\mathbb{R}^n)^*$ of linear functions vanishing on all elements of S , as discussed earlier at least for subspaces S .

Ex. For any subset S of \mathbb{R}^n , Sperp is a subspace of \mathbb{R}^n . For any subset S of \mathbb{R}^n ,

$\text{dimension}(\text{Sperp}) + \text{dimension}(\text{span}(S)) = n.$

Ex. For any matrix A , the nullspace and row space of A are orthogonal complements of each other, i.e. $N(A) = R(A)\text{perp}$, and $R(A) = N(A)\text{perp}$.

Remarks: The statement of the previous exercise reflects the basic fact that there are always two complementary ways to represent any subspace, namely implicitly and parametrically. I.e. one either gives a finite set that *spans* the space, so that elements of the space are obtained as linear combinations of the given spanning vectors, or else one gives *equations* for the space, so that elements of the subspace are those vectors that satisfy the equations. Thus the rows themselves give a spanning set for the row space, while any basis for the null space gives a finite set of equations for the row space.

These two ways of representing a subspace have complementary virtues. If you want to *produce* elements of a subspace, you take any linear combination of a spanning set. On the other hand if you want to *recognize* an element of your subspace, in the sense that someone presents you with a vector and asks whether it belongs to your subspace, then you want equations which are satisfied if and only if the answer is yes. Thus the actual rows give a spanning set for the row space, while the basis for the null space gives a set of equations for recognizing, or characterizing, elements of the row space.

Remark: Since when a matrix acts by multiplication on a column vector, the individual row vectors act as linear functions on that column vector, it is natural to think of the row vectors as elements of $(\mathbb{R}^n)^*$. I.e. a 1 by n matrix, a single row vector, defines a linear map $\mathbb{R}^n \rightarrow \mathbb{R}$. Similarly, since the column vectors of A are the values of the map $\mathbb{R}^n \rightarrow \mathbb{R}^m$ represented by A on the standard basis vectors, it is natural to think of column vectors of A as elements of \mathbb{R}^m . I.e. column vectors are “vectors”, while row vectors are “dual vectors” or linear functions.

Proposition: A matrix and its row reduced echelon form have the same null space.

Proof: We already know that performing a row operation leaves the row space unchanged, hence it also leaves the orthogonal complement, the null space, unchanged. **QED.**

Corollary: The homogeneous system of equations associated to a matrix or to its reduced echelon form, always have the same solutions.

Note: This was of course the reason for carrying out row reduction in the first place, i.e. the fact that it simplifies the equations without changing the solutions.

The null space as “relations” among the columns

If w is a vector in the null space of the m by n matrix A , then $w \cdot v_j = 0$ for every row v_j of A . But the j th row consists of the j th entries of every column of A . So all the entries of the columns satisfy the same relation. I.e. if $w = (a_1, \dots, a_n)$, and u_1, \dots, u_n are the columns of A , then $a_1 \cdot u_1 + \dots + a_n \cdot u_n = 0$, i.e. the scalars a_i are the coefficients of a relation among the columns $\{u_i\}$. Now since the null space stays the same under row operations, this means the columns of all row equivalent matrices satisfy the same relations. In particular a subset of the columns of A are independent, i.e. satisfy no non trivial relations, if and only if the corresponding columns of the reduced echelon form are independent. Thus since the pivot columns are the earliest subset of columns that form a basis of the column space of the reduced matrix, the columns of A that are in the same positions as the pivots of the reduced form are also independent. In fact, since the j th column depends on earlier columns if and only if there is a relation among the columns whose last non zero entry occurs in the j th position, we can characterize the pivot columns of A as those that do not depend on earlier columns.

Consequently, the process of reducing a sequence of vectors to an independent sequence, by eliminating those vectors that depend on earlier ones, can be carried out in \mathbb{R}^m as follows. Just place the given sequence of vectors as the columns of a matrix, and row reduced that matrix. After finishing, identify the pivot columns. Then go back and take as the independent subset with the same span, precisely those columns in the same positions as the pivot columns of the reduced matrix. These, the pivot columns of A , form a basis for the column space of A obtained by eliminating those columns that depend on previous ones.

Similarly one can extend an independent set of r vectors in \mathbb{R}^m to a basis by forming a matrix with the given vectors as the first r columns and completing the matrix by the standard basis as the next m columns. Then row reduce the matrix. The first r columns will be found to be among the pivots, and taking precisely all pivot columns of the original matrix will give a basis containing the original independent set. In this way the theoretical operations of reducing a spanning set to a basis or extending an independent set to a basis, can be carried out in practice in \mathbb{R}^m .

Remark:

These “relations” also give another proof of the uniqueness of the reduced echelon

form of A . Namely the location of the pivot columns is determined by the fact that they are exactly those columns of A that do not depend on previous columns, a property that is preserved under row operations. The pivot columns of A thus form a basis for the space spanned by all the columns of A . Each non pivot column is thus a unique linear combination of the pivot columns. After each row operation the space spanned by the columns may change, but the relations among the columns, which are exactly the elements of the null space, do not change. Hence with each row operation, the columns in the pivot positions are again a basis for the new column space, and each column in a non pivot position is still the same linear combination of the new pivot columns. Since the pivot columns of the reduced form are the first r standard basis vectors, where $r =$ the row rank of A , at that point the non pivot columns contain exactly the coefficients by which these columns depend on the pivot columns. Thus the unique relations by which the original non pivot columns depend on the original pivot columns, determine the reduced row echelon form.

This uniqueness argument using the null space also has a geometric version.

Conceptual proof of uniqueness of reduced echelon form:

It is fundamental that a matrix and its reduced echelon form have the same null space. Indeed that is the reason reduced echelon forms are useful for finding the null space of the original matrix. I claim that null space determines entirely the reduced echelon form. For simplicity take the case where the “pivot” columns all appear first, followed by the non pivot columns. (Note that a column is a pivot column if and only if it does not depend linearly on earlier columns. I.e. this is obvious for a reduced echelon matrix and hence also true for the original matrix, since the columns of both matrices satisfy exactly the same relations.)

If the matrix A is n by m with rank r , then the reduced echelon form has its last $n-r$ rows equal to zero and its upper left r by r block equal to an identity matrix. Thus it suffices to show that the null space characterizes the remaining upper right (r) by $(m-r)$ block. Such a block of course determines, and is determined by, a unique linear transformation from $m-r$ space to r space. Further, that linear transformation is determined by its graph, an $m-r$ dimensional linear subspace of m space.

That subspace, i.e. that graph, except for a minus sign, is precisely the null space. I.e. from looking at the reduced echelon form one can see that the negative of the upper right r by $(m-r)$ block is exactly the matrix of the linear map whose graph is the null space. I.e. the reason the reduced echelon form is useful for producing elements of the null space is its form lets you take any values at all for the non

pivot variables and solve for the unique corresponding values of the pivot variables that give an element of the null space. Thus the null space defines a linear function from the space of the non pivot variables to the space of the pivot variables. And in fact the non pivot part of the reduced echelon form of the matrix is just the negative of the matrix of that function.

Looked at another way, if A is the given matrix, the equation $AX=0$ determines implicitly a linear function from $m-r$ to r space whose matrix is the negative of the upper right r by $(m-r)$ block of the reduced echelon form of A .

I.e., the null space determines a linear map from the coordinate subspace spanned by the non pivot variables to that spanned by the pivot variables, whose matrix columns are (when augmented at the bottom by zeroes) exactly minus the sequence of non pivot columns of the reduced echelon form of A . Since both the location and the content of the pivot columns are known, the reduced form is determined by the null space.

Summary: Given a matrix A , the equation $AX=0$ determines implicitly a linear function from the space of non pivot variables to the space of pivot variables. (A pivot column is one that is not a linear combination of earlier ones, so A determines its pivot variables.) That function has two incarnations, its matrix and its graph. The (interesting part of the) reduced echelon form is (minus) that matrix, and the null space of A is the graph. Since A uniquely determines its null space, it also uniquely determines its reduced echelon form.

One can also give a geometric version of the uniqueness argument via the row space. I.e. since the row space projects isomorphically onto the coordinate subspace spanned by the pivot variables, it must be the graph of a linear map whose source space is the subspace of pivot variables and whose target is the subspace of non pivot variables.

Ex. Show that the non pivot part of the reduced row echelon form of a matrix A is the transpose of the matrix of the linear map from the subspace spanned by the pivot variables to the subspace spanned by the non pivot variables, and whose graph is the row space $R(A)$.

The span of the column vectors also has an intrinsic meaning in terms of the map defined by the matrix.

Definition: The “column space” $C(A)$ of an m by n matrix A , is the subspace of \mathbb{R}^m spanned by the columns of A .

Proposition:

The column space of a matrix is the image of the map defined by the matrix.

Proof: By definition, the columns of A are the images of the standard basis under the map defined by A . Since the standard basis spans the domain space, their images span the image of the domain space, i.e. the image of the map. **QED.**

Definition: The column rank of a matrix is the dimension of the span of its columns.

Proposition: The column rank and row rank of a matrix are equal.

Proof: One proof is to note that the row rank equals the number of non zero rows in the reduced echelon form, which equals the number of pivot columns. Since the non zero rows are a basis for the row space and the pivot columns are a basis for the column space this does it.

Another proof is to observe that the column space is the image space of the associated map, which is isomorphic to the quotient of the domain space by the null space. Since the rows are dual vectors that map the null space to zero, they span the dual space to the quotient of the domain by the null space, so the row space is dual to the column space, hence they have the same dimension. **QED.**

Remark: Since $\text{rank}R(A) = \text{rank}C(A)$, we call it simply the rank of A . In particular this equals the dimension of the image of the map defined by A , which thus agrees with the terminology of rank of a map.

Remark: A fourth subspace $N(A^*)$ is sometimes introduced, the null space of the “transpose” of A , the n by m matrix A^* obtained by interchanging the rows and columns of the m by n matrix A . This is the space of equations defining the image subspace $C(A)$. Hence $N(A^*) = C(A)^\perp$, and vice versa. In particular, as a corollary of the previous result, A and A^* have the same rank. When we see below how to represent any linear map by a matrix, using a basis, we will see that A^* represents the transpose of the map represented by A .

Two ways to represent a subspace, “parametrically” and “implicitly”

There are two complementary ways to represent a subspace, first by giving a basis or a spanning set for it, and second by giving equations for it. The first method parametrizes the subspace as the image of a map from some \mathbb{R}^n , and the second defines it implicitly as the solution space of some equations. A fundamental problem is to start from one representation and find the other. Thus a system of linear equations gives an implicit representation of the solution space, the null

space of that system. Solving the equations means finding a spanning set or basis for that null space. Thus given any matrix, the rows give a finite system of linear equations for the null space, and conversely a basis for the null space gives a finite system of linear equations for the row space. I.e. the rows parametrize the row space, while the null space represents the row space implicitly. Similarly, a basis for the null space of the transpose A^* gives linear equations for the column space $C(A)$, the subspace spanned by the columns of A .

Remarks on non linear geometry:

The study of the geometry of subsets of \mathbb{R}^n defined by higher degree equations, i.e. non linear ones, is much more complicated in this regard. E.g. a plane curve defined implicitly by a general equation of degree ≥ 3 , cannot be parametrized by any polynomial map from \mathbb{R}^1 to the curve, in fact all such maps are constant! It is feasible to find equations for the image of a polynomial map however and there even exist computer programs to do this. The point is that parametrizable subsets form a small subfamily of all implicitly definable “algebraic sets”. Indeed it turns out that if we consider also complex points of our sets, that every curve becomes a surface, and the parametrizable surfaces are all “spherical”. I.e. general algebraic curves correspond to surfaces that have the topology of a surface possibly with handles or holes, like a doughnut or multi - holed doughnut. Then considerations of topology show that one cannot parametrize a surface with holes by any surface with fewer holes. Parametrizing a real curve by \mathbb{R} corresponds to parametrizing the complex form of the curve by \mathbb{C} , and (after adding a point at infinity), \mathbb{C} corresponds to the sphere, i.e. the surface with no holes. One can map any surface onto one with fewer holes however and it is usual to study plane curves this way, e.g. by projecting them onto the X - axis and studying the fibers of the projection.

Summary: Any matrix can be reduced by elementary row operations to a unique matrix in reduced echelon form. Two matrices of the same size have the same reduced echelon form if and only if they have the same row space. After reduction, the rows of the reduced form will be a basis of the original row space.

The columns of the original matrix which are in the same positions as the “pivot columns” of the reduced form are a basis for the column space of the original matrix. In both the original matrix and the reduced form, a column is a pivot column if and only if it does not depend linearly on previous columns.

The null space of a matrix as well as the row space, are the same as the null space and row space of the reduced form. Row reduction however transforms the column space of a matrix into one of the standard coordinate subspaces, i.e. the

column space of the reduced form of a matrix usually shares only its dimension with the original column space.

An m by n matrix represents a linear map from \mathbb{R}^n to \mathbb{R}^m . The column space of a matrix is exactly the image space of its associated linear map. The null space of a matrix is exactly the kernel of the map it represents.

Given a matrix A , the subspaces $N(A)$ and $R(A)$ are dual in the sense that a basis for one gives linear equations for the other. Likewise, $N(A^*)$ and $C(A)$ are similarly dual subspaces.

Worked example:

Start from a system of 3 equations in 3 variables:

$$2X - Y + 2Z = 0$$

$$3X + 3Y + Z = 0$$

$X - 5Y + 3Z = 0$, and form the associated matrix A of coefficients

$$\begin{bmatrix} 2 & -1 & 2 \\ 3 & 3 & 1 \\ 1 & -5 & 3 \end{bmatrix}$$

Now row reduce it, (this took me a few failures).

$$\begin{bmatrix} 1 & 0 & 7/9 \\ 0 & 1 & -4/9 \\ 0 & 0 & 0 \end{bmatrix}$$

Now we have two spanning sets for $R(A)$, the original spanning set of rows:

$\{(2, -1, 2), (3, 3, 1), (1, -5, 3)\}$; and the rows of the reduced matrix

$\{(1, 0, 7/9), (0, 1, -4/9)\}$, which form a basis.

We also know $N(A)$ has dimension one, since there is only one non pivot variable, namely Z . We get a basis for $N(A)$ by starting from the standard basis of the non pivot space, and solving for X and Y using the reduced equations:

$X + (7/9)Z = 0$, so $X = (-7/9)Z$; and $Y - (4/9)Z = 0$, so $Y = (4/9)Z$. Since the non

pivot space is just the Z axis, the standard basis is $Z=1$, and we get $X = -7/9$, $Y = 4/9$. Thus the basis vector for $N(A)$ is $(-7/9, 4/9, 1)$. If we want integers, multiply by 9 and get $(-7, 4, 9)$.

This basis vector for $N(A)$ gives an equation for $R(A)$ when viewed as an equation, namely $-7X + 4Y + 9Z = 0$. You can check that all rows do satisfy this equation.

Of course the rows also give equations for $N(A)$, either the original equations, or more efficiently the (almost) reduced equations $9X+7Z = 0 = 9Y-4Z$.

Knowing where the pivot columns are also gives us a basis of the 2 dimensional column space, namely the two columns in the same positions as the pivot columns in the reduced matrix. Thus the first two columns (written horizontally here) $(2,3,1)$, $(-1,3,-5)$, form a basis of $C(A)$. If we want an equation for $C(A)$ we have to row reduce the transpose matrix, although now that we know the first two columns already span $C(A)$, we could just use them. I.e. we might as well just row reduce this matrix, the transpose of the pivot columns:

$$\begin{array}{ccc|c} 2 & 3 & 1 & \\ -1 & 3 & -5 & \end{array}. \text{ I claim the reduced form is this:}$$

$$\begin{array}{ccc|c} 1 & 0 & 2 & \\ 0 & 1 & -1 & \end{array}. \text{ So a basis of } N(A^*) \text{ is } \{(-2,1,1)\}, \text{ hence an equation for } C(A) \text{ is given by}$$

$-2U + V + W = 0$, where we have chosen to use different variables U, V, W for the column space.

Another example:

$$\begin{array}{cccc|c} -1 & 1 & 2 & 0 & = B \\ 2 & 2 & -1 & -1 & \\ -1 & 5 & 5 & -1 & \end{array}. \text{ This matrix } B \text{ reduces to the following matrix:}$$

$$\begin{array}{cccc|c} 1 & 0 & (-5/4) & (-1/4) & \\ 0 & 1 & (3/4) & (-1/4) & \\ 0 & 0 & 0 & 0 & \end{array}.$$

If the variables are X, Y, Z, W , this time there are 2 non pivot variables, Z, W so to get a basis of $N(B)$, we take the 2 standard basis vectors of that space namely $(Z, W) = (1, 0)$, and $(Z, W) = (0, 1)$, and for each of these we solve the reduced equations for X and Y . This gives the following two basis vectors for $N(B)$: $\{((5/4), (-3/4), 1, 0), ((1/4), (1/4), 0, 1)\}$. If we prefer integers, we get $\{(5, -3, 4, 0), (1, 1, 0, 4)\}$. As before these give equations for $R(B)$.

From the rows of the reduced matrix we find the basis $\{(4, 0, -5, -1), (0, 4, 3, -1)\}$ of $R(B)$. And since the first two columns are the pivots, we get the basis $\{(-1, 2, -1), (1, 2, 5)\}$ for $C(B)$. Thus to find equations for $C(B)$ we can reduce

$\begin{vmatrix} -1 & 2 & -1 \\ 1 & 2 & 5 \end{vmatrix}$, to find a basis for $N(B^*)$, since this matrix has the same row space and hence the same null space as the full B^* .

Ex. Find a basis for $N(B^*)$, hence find equations for $C(B) = R(B^*)$.

The matrix associated to a linear map $T:V \rightarrow W$ by bases of V, W .

Since any finite dimensional vector space is isomorphic to some \mathbb{R}^n by choosing a basis, any linear map between finite dimensional spaces can be represented by a matrix by choosing bases. I.e. if T is any linear map from one finite dimensional vector space V to another W , then by choosing bases for V and W we obtain isomorphisms between these abstract spaces and some coordinate spaces \mathbb{R}^n and \mathbb{R}^m . Hence, if $\dim(V) = n$ and $\dim(W) = m$, we obtain a resulting linear map from \mathbb{R}^n to \mathbb{R}^m which has a matrix A . This A is called the matrix of T associated to the given bases for V and W .

So given $T:V \rightarrow W$ and bases v_1, \dots, v_n , and w_1, \dots, w_m of V, W , then A is the matrix of the composition $\mathbb{R}^n \rightarrow V \rightarrow W \rightarrow \mathbb{R}^m$, where the first (left) map $\mathbb{R}^n \rightarrow V$ sends each e_j to v_j , the second map is $T:V \rightarrow W$, and the 3rd map $W \rightarrow \mathbb{R}^m$ sends w to the coefficient vector of its representation in the basis $\{w_j\}$.

Thus if $N: \mathbb{R}^n \rightarrow V$, and $M: \mathbb{R}^m \rightarrow W$, are the parametrizations determined by the given bases, then A is the matrix of the composition $M^{-1} \circ T \circ N$. Thus the j^{th} column of the matrix A for T , is the coefficient vector (c_1, \dots, c_m) of the image of the standard basis vector e_j under the triple composition. Since $M(e_j) = v_j$, this is the coefficient vector of $T(v_j) = c_1 w_1 + \dots + c_m w_m$, in the given basis for W .

A map from V to itself has a matrix associated to any basis of V . I.e. if v_1, \dots, v_n , is a basis for V , and $N: \mathbb{R}^n \rightarrow V$ is the associated parametrization, the matrix for a map $T:V \rightarrow V$ in this basis is the matrix of the composition $N^{-1} \circ T \circ N: \mathbb{R}^n \rightarrow \mathbb{R}^n$.

Thus if $T:\mathbb{R}^n \rightarrow \mathbb{R}^n$ is already given by a matrix B in the standard basis, and we want to express it in a new basis v_1, \dots, v_n , then N is the matrix with the v 's as columns, and the new matrix A for T in terms of this new basis is $A = N^{-1} \cdot B \cdot N$.

Eg: If $D:V \rightarrow V$ takes a polynomial of degree ≤ 2 to its derivative, the matrix of D

in the basis $\{1, X, X^2\}$ has columns $(0,0,0)$, $(1,0,0)$, $(0,2,0)$, since $D(1) = 0 = 0(1,0,0) + 0(0,1,0) + 0(0,0,1)$, and $D(X) = 1 = 1(1,0,0) + 0(0,1,0) + 0(0,0,1)$, and $D(X^2) = 2X = 0(1,0,0) + 2(0,1,0) + 0(0,0,1)$.

Map composition corresponds to matrix multiplication

We can even compute compositions of maps by multiplying matrices. Namely, if B, A are matrices where the number of columns of B equals the number of rows of A , so that the rows of B have the same length as the columns of A , we define the product BA to be the matrix whose (i,j) entry is the dot product of the i^{th} row of B with the j^{th} column of A . The product BA then has the same number of rows as B and the same number of columns as A .

Ex: If $T:V \rightarrow W$ and $S:W \rightarrow U$ are linear maps, and we choose bases for all three spaces, the matrix of the composition SoT has as entry in its i^{th} row and j^{th} column, the dot product of the i^{th} row of the matrix for S with the j^{th} column of T . (Sketch: Assigning to each vector in a space its i^{th} coefficient in a given basis, is a linear function. Thus if the coefficient of u_i in the expansion of $S(w_1)$ is a_1, \dots , and the coefficient of u_i in the expansion of $S(w_n)$ is a_n , and if $T(v_j) = b_1 w_1 + \dots + b_n w_n$, then the coefficient of u_i in the expansion of $S(T(v_j)) = b_1 S(w_1) + \dots + b_n S(w_n)$, is $b_1 a_1 + \dots + b_n a_n$. Since the i^{th} row of S consists of the coefficients (a_1, \dots, a_n) of u_i in the expansions of the images $S(w_1), \dots, S(w_n)$ of the basis vectors w_1, \dots, w_n , and the j^{th} column of T consists of the coefficients (b_1, \dots, b_n) of the vector $T(v_j)$ in the basis w_1, \dots, w_n , the result follows.)

Thus if A is the matrix of T , and B is the matrix for S , the matrix product $BA =$ the matrix for SoT , (in the same bases).

E.g. If $B = \begin{vmatrix} 2 & 3 \\ 1 & -4 \end{vmatrix}$, $A = \begin{vmatrix} 5 & 7 \\ 8 & 2 \end{vmatrix}$, then $BA = \begin{vmatrix} 34 & 20 \\ -27 & -1 \end{vmatrix}$

Rmk: Since map composition is associative, the previous exercise implies that matrix multiplication is also associative. Of course since if a matrix A has n rows, and B has m columns, and AB is defined, the nm entries in the product AB just consist of nm operations of taking dot products of rows of A by columns of B , so the easy fact that dot product of vectors is associative also implies the result for matrices.

Ex. Find examples of two 2×2 matrices A, B such that $A \cdot B \neq B \cdot A$, hence matrix multiplication is (usually) not commutative. Find some (non zero) 2×2 matrices A, B that do commute, i.e. such that $A \cdot B = B \cdot A$.

Identities, isomorphisms, and matrix inverses

An n by n matrix with all diagonal entries equal to 1, and all other entries equal to 0, represents the identity map $\mathbb{R}^n \rightarrow \mathbb{R}^n$ and thus is called an identity matrix.

If the matrix A represents an isomorphism, then since isomorphic spaces have the same dimension, A is square, hence n by n for some n , and the inverse isomorphism of A is represented by some n by n matrix B with $AB = BA = I$, where I is the n by n identity matrix. We call B the (matrix) inverse of A . Note that if A is a square matrix, then a matrix B is a left inverse for A if and only if it is a right inverse, since this holds for linear maps.

We know that an n by n matrix A represents an isomorphism $\mathbb{R}^n \rightarrow \mathbb{R}^n$ if and only if the map it represents takes the standard basis vectors to a basis of \mathbb{R}^n , i.e. if and only if the columns of A are a basis of \mathbb{R}^n . Then since the column rank equals the row rank, this is equivalent to the fact that the rows of A form a basis of \mathbb{R}^n . Thus an n by n matrix A is invertible if and only if the rows form a basis of \mathbb{R}^n if and only if the columns form a basis of \mathbb{R}^n , if and only if A has rank n . Of course conceptually, since the rows represent the coordinate functions of the map represented by A , i.e. linear functions from \mathbb{R}^n to \mathbb{R} , we might more properly view the rows as a basis for $(\mathbb{R}^n)^*$, and the columns as a basis for \mathbb{R}^n .

The matrix of the transpose is the transpose of the matrix

We know an m by n matrix A defines a linear map $A: \mathbb{R}^n \rightarrow \mathbb{R}^m$ and that there is an associated transpose map $A^*: (\mathbb{R}^m)^* \rightarrow (\mathbb{R}^n)^*$ defined by “preceding by A ”. We also know this transpose map has a matrix representation as an n by m matrix, using the bases of $(\mathbb{R}^m)^*$ and $(\mathbb{R}^n)^*$ dual to the standard bases of \mathbb{R}^m and \mathbb{R}^n . We have already defined the transpose of the matrix A to be the matrix A^* obtained by interchanging the rows and columns of A . Fortunately this transpose matrix does represent the matrix of the transpose map. Sometimes I denote the matrix of a map T by $[T]$. In that notation, then we claim $[T^*] = [T]^*$.

I still find this stuff confusing but here is my attempt to argue this.

Let $T: \mathbb{R}^n \rightarrow \mathbb{R}^m$ be the map defined by the m by n matrix A . Thus the 1st column of A are the coefficients of the image vector $T(e_1)$ in terms of the basis vectors e_1, \dots, e_n . Similarly the first row of A consists of just the first coefficient of

each of the image vectors $T(e_1), \dots, T(e_n)$, i.e. the first row of A is the sequence of numbers $e_1^*(T(e_1)), \dots, e_1^*(T(e_n))$.

Likewise, the matrix of $T^*: (R^m)^* \rightarrow (R^n)^*$, in the dual basis, has as 1st column, the coefficients of the function $T^*(e_1^*)$ in terms of the dual basis e_1^*, \dots, e_n^* . And just as the e_j coefficient of a vector v in R^n equals the value of e_j^* on v , dually the e_j^* coefficient of the function $T^*(e_1^*)$ in $(R^m)^*$, is the value of this function on e_j . Thus the entries of the first column of the matrix for T^* is the sequence of numbers $T^*(e_1^*)(e_1), \dots, T^*(e_1^*)(e_n)$. By definition of T^* as “preceding by T ”, this equals the sequence $e_1^*(T(e_1)), \dots, e_1^*(T(e_n))$. Thus the first row of the matrix for T , in the standard bases for R^n, R^m , equals precisely the first column of the matrix for T^* , in the dual bases for $(R^n)^*, (R^m)^*$. I.e. the matrix in the natural bases, for the transpose map $T^*: (R^m)^* \rightarrow (R^n)^*$, is just the transpose of the matrix for the map $T: R^n \rightarrow R^m$, obtained by interchanging rows and columns.

To argue it just in terms of matrices, note that the isomorphism $R^n \rightarrow (R^n)^*$ taking a vector v to “dotting with v ”, or $v(\cdot)$, just takes a column vector to its transpose, since a row vector acts on a column vector by dotting with it. Then given a matrix A defining a map $T: R^n \rightarrow R^m$, by definition the matrix of the map T^* is that of the composition $R^m \rightarrow (R^m)^* \rightarrow (R^n)^* \rightarrow R^n$, where the first map takes a column vector to a row vector, the last map does the opposite, and the middle map is T^* = preceding by T . Since T is represented by A , T^* can be represented by preceding by A . Thus this composition sends the column vector e_1 to the row vector e_1^* , whose matrix is the row vector $[1, 0 \dots 0]$, so preceding by A gives the product $[1 \ 0 \dots 0].A$ = the first row of A . The last map then sends this to a column vector, the first column vector of the matrix for T^* .

As a corollary, for maps S, T , we have $(SoT)^*(g) = go(SoT) = (goS)oT = T^*(goS) = T^*(S^*(g)) = (T^*oS^*)(g)$, hence $(SoT)^* = T^*oS^*$. Thus the same is true for transposes of matrices, i.e. as long as the product $A.B$ is defined, then $(AB)^* = B^*A^*$. If we had checked this directly, as is not too hard, we could have obtained the result just argued for transposes. I.e., using what we just claimed about transposing matrices, since $[e_1^*] = [e_1]^*$, we have: transpose of first row of $A = ([e_1^*].A)^* = ([e_1]^*.A)^* = A^*.[e_1] =$ first column of A^* .

If you are still puzzled by this, just take two column vectors v, w of the same length, and denote the corresponding row vectors by v^*, w^* , and check that $v^*.w = w^*.v$, as matrix products. That’s all that’s going on here. I.e. you just do this one row and column at a time to get the same result for matrices.

I.e. if $v =$

$$\begin{bmatrix} 2 \\ 4 \\ 1 \end{bmatrix}, \text{ so } v^* = [2 \ 4 \ 1]$$

and if $w =$

$$\begin{bmatrix} 5 \\ 2 \\ 7 \end{bmatrix}, \text{ so } w^* = [5 \ 2 \ 7], \text{ then,}$$

$$v^* \cdot w = 2 \cdot 5 + 4 \cdot 2 + 1 \cdot 7 = 5 \cdot 2 + 2 \cdot 4 + 7 \cdot 1 = w^* \cdot v = 25.$$

Well, this is even more enlightening, since it shows that the result just depends on commutativity of multiplication in the field of scalars. I.e. this whole thing is just a reflection of the fact that $v \cdot w = w \cdot v$ for dot products of vectors. That was not as obvious, at least to me, in those more abstract arguments.

We also get a nice corollary about dot products, since for two n dimensional column vectors v, w their dot product, which we could write as $\langle v, w \rangle$, equals the matrix product $v^* \cdot w$. Now if A is an n by n matrix then the dot product of Av with w equals $(Av)^* \cdot w$. Thus $\langle Av, w \rangle = (Av)^* \cdot w = v^* \cdot A^* \cdot w = \langle v, A^* w \rangle$. Thus you can move a matrix from one factor in a dot product to the other, if you change the matrix into its transpose.

Elementary row operations as multiplication by “elementary matrices”

Now that we know more about multiplying matrices, we can reinterpret the process of row reduction more conceptually in terms of linear isomorphisms. Recall that in an n by m matrix A representing a linear map $R^n \rightarrow R^m$, the columns are the coefficients of the values $A(e_j)$ of the map at the standard basis vectors. If we interchange the first two rows, we have changed the coefficients expressing these values, and thus we have changed the values of the map on the standard basis vectors. In this case we have changed the map A into the composition of the original map A , with the map on the target space that interchanges the first two basis vectors there. Since this operation is composing with a linear map on the target space R^m , it must be achieved by left - multiplying A by some square m by m matrix that represents interchanging the first two standard basis vectors of R^m . The matrix that does this is just the result of performing the same row operation on the identity m by m matrix. The same holds for all elementary row operations as is easy to check in examples.

Definition: An elementary matrix is a square matrix obtained from the identity matrix by performing one of the elementary row operations.

Then each elementary row operation on an m by n matrix A , can be performed by left multiplying A by the corresponding elementary m by m matrix E .

For example, here are some elementary matrices:

$$\begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix},$$

$$\begin{pmatrix} c & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix},$$

$$\begin{pmatrix} 1 & 0 & c \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

Do some experiments and see what these do to a 3- rowed matrix A , when used to multiply A from the left.

Of course we will not stop using our old row operations in favor of this new process, since that is less efficient, but just noticing this new interpretation of row operations gives us some new insight into them. E.g. since row operations are reversible by other row operations, it follows that every elementary matrix is invertible, hence so is every product of them. But more is true.

Proposition: Every invertible matrix is a product of elementary matrices.

Proof: It suffices to show that every invertible matrix can be obtained by applying elementary row operations to the identity matrix. Then since row operations are invertible, it also suffices to show that every invertible matrix can be transformed into the identity matrix by row operations. This is true however, since every matrix can be reduced by row operations to reduced echelon form, and the reduced echelon form of an n by n matrix of rank n , has all n columns as pivots, hence is the identity matrix. Thus every invertible matrix A has reduced echelon form equal to the identity matrix, hence there is a sequence of elementary row operations, hence a product of elementary matrices, that multiplies A into the identity matrix.

Cor: Row operations can be used to calculate the inverse of a matrix.

Proof: Take the n by n matrix A and add to it a copy of the n by n identity matrix, to form an n by $2n$ matrix. Then do row operations that reduce A to the identity matrix, while simultaneously carrying out the same operations on the identity matrix. At the end, A will be transformed into the identity by left multiplying by some product B of elementary matrices such that $BA = I$, so that $B = A^{-1}$. Thus the identity matrix will have been transformed into the product $BI = B = A^{-1}$.

QED.

Example: To find the inverse of $A =$

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

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

We row reduce the augmented matrix $[A \ I]$ as follows:

$$\begin{array}{l} | 1 \ 0 \ 0 | 1 \ 0 \ 0 | \\ | 1 \ 1 \ 0 | 0 \ 1 \ 0 | \approx \\ | 1 \ 1 \ 1 | 0 \ 0 \ 1 | \end{array}$$

$$\begin{array}{l} | 1 \ 0 \ 0 | 1 \ 0 \ 0 | \\ | 0 \ 1 \ 0 | -1 \ 1 \ 0 | \\ | 0 \ 1 \ 1 | -1 \ 0 \ 1 | \approx \end{array}$$

$$\begin{array}{l} | 1 \ 0 \ 0 | 1 \ 0 \ 0 | \\ | 0 \ 1 \ 0 | -1 \ 1 \ 0 | = [I \ A^{-1}] \\ | 0 \ 0 \ 1 | 0 \ -1 \ 1 | \end{array}$$

I.e. now the right half of this matrix is the inverse of A . You can check this by multiplying it by A :

$$\begin{array}{l} | 1 \ 0 \ 0 | | 1 \ 0 \ 0 | \quad | 1 \ 0 \ 0 | \\ | 1 \ 1 \ 0 | | -1 \ 1 \ 0 | = | 0 \ 1 \ 0 | = I. \text{ So it checks.} \\ | 1 \ 1 \ 1 | | 0 \ -1 \ 1 | \quad | 0 \ 0 \ 1 | \end{array}$$

Note: You do not need to know in advance that A is invertible, since that will be revealed by the fact that the reduced echelon form of A is I . In fact the same procedure finds a left - inverse of an m by n matrix A whose n columns are independent. In this case one takes only the first n rows of the resulting B to give the left inverse.

Classification of linear maps up to equivalences:

There are several natural equivalence relations on linear maps $V \rightarrow W$ defined in terms of isomorphisms. Everything we have done so far can be reinterpreted in terms of these equivalence relations. Two linear maps $S: V \rightarrow W$, and $T: V \rightarrow W$ between finite dimensional spaces V, W , are called *equivalent* if there are isomorphisms $P: V \rightarrow V$ and $Q: W \rightarrow W$ of the source and target spaces such that $QoS = T$. We saw at the end of chapter one that this is possible if and only if S, T have the same rank. We call S and T *left - equivalent* if P can be taken as the identity, i.e. if $T = QS$ for some isomorphism Q . Similarly, S and T are *right - equivalent* if Q can be taken as the identity and $T = SP$. The row operations studied in this chapter provide canonical representatives for these equivalence relations as follows.

Since a matrix M is left equivalent to exactly those products of form EM where E is invertible, it follows that M is left equivalent to precisely those matrices into which it can be transformed by elementary row operations. In particular M is left equivalent exactly to those matrices (of the same size) with the same row space, hence also with the same null space. Intrinsically then, two maps $V \rightarrow W$ are left equivalent if and only if they have the same kernel. This of course implies they have the same rank, but is a much stronger restriction, since their kernels must not only have the same dimension, but must be actually the same subspace.

If we define elementary *column* operations in the same way as elementary row operations, (or just as row operations applied to the transpose of the matrix), it follows that two matrices (of the same size) are right equivalent if and only if they can be transformed into one another by elementary column operations, if and only if they have the same column space. Since the columns span the image of the corresponding map, two abstract maps $S, T: V \rightarrow W$ are right equivalent if and only if they have the same image, if and only if their transposes $S^*, T^*: W^* \rightarrow V^*$ have the same kernels.

Thus elementary operations allow us to compute equivalence classes of maps, since then two maps $S, T: V \rightarrow W$ are left - equivalent iff their matrices (in the same bases) are equivalent by row operations; they are right - equivalent iff their matrices are equivalent under column operations, and they are just equivalent iff their matrices are equivalent using both row and column operations. In particular, given bases for V, W , we can compute whether or not two maps $S \rightarrow W$ have the same kernel and/or image using matrices, and in either case we can find bases for those subspaces, which are “canonical” in terms of the given bases for V, W .

What comes next?

In the next chapter we take up the more subtle question of how simple can we make the matrix of a map $T:V \rightarrow V$ from a space to itself, where naturally we now require using the same basis in source and target since they are the same space. As it turns out, this “similarity” problem has a nice theoretical answer, the rational canonical form, which can again be computed using row and column operations, this time applied to matrices with polynomial entries. There is also a more refined theoretical answer, the general Jordan form which, except in some special cases, is less computable due to the practical difficulty of factoring polynomials.

Chapter Three: Decomposing V into “ T -cyclic” subspaces

We want to classify linear maps $T:V \rightarrow V$ of a finite dimensional space V up to similarity; e.g. how simple can we make the matrix of T by a good choice of basis of V ? The essential observation here is the fact that since the source and target space are the same, we can now compose T with itself, hence we can form powers of T , and thus also polynomials in T . Note that if S is similar to T , i.e. if $S = U \circ T \circ U^{-1}$, then also S^n is similar to T^n , since then $S^n = (U \circ T \circ U^{-1}) \dots (U \circ T \circ U^{-1}) = U \circ (T^n) \circ U^{-1}$, due to canceling adjacent pairs of $U^{-1} \circ U$ in the middle. Consequently, if S is similar to T , then every polynomial P in S is similar to the same polynomial in T : i.e. if $S = U \circ T \circ U^{-1}$, then $P(S) = U \circ P(T) \circ U^{-1}$.

Ex. Check the claim we just made for similarity of polynomials in S, T .

Similar polynomials satisfy the same polynomials

The observation that forming polynomials preserves similarity leads in some cases to a complete understanding of the similarity class of an operator. In particular, since only the zero operator is similar to itself, if we find a polynomial P such that $P(T) = 0$, then every operator S similar to T must also satisfy $P(S) = 0$. For some operators this is also sufficient as we shall see. First we discuss the general idea of polynomials satisfied by an operator.

Terminology: A polynomial is called “monic” if the lead coefficient equals one.

Lemma: Every linear operator $T:V \rightarrow V$ on a finite dimensional k - vector space V , satisfies some monic (hence non zero) polynomial over k .

Proof: If $\dim(V) = n$, then $\dim(\text{Hom}(V, V)) = n^2$, but $k[X]$ is infinite dimensional, with basis all monomials $\{1, X, X^2, X^3, \dots\}$. Thus the map $k[X] \rightarrow \text{Hom}(V, V)$ has a non zero kernel, i.e. for some $f \neq 0$, $f(T) = 0$. Dividing through by the leading non-zero coefficient makes the polynomial monic and T still satisfies it. **QED.**

Lemma: If There is a unique monic polynomial of least degree satisfied by T .
Indeed this minimal polynomial divides all other polynomials satisfied by T .

Proof: If f, g are two (non zero) polynomials of least degree satisfied by f , and we divide g by f , we get an equation of form $g = qf + r$, where $\deg(r) < \deg(f)$. Since $r = g - qf$, and T satisfies both f and g , it also satisfies r . Since r has degree less than f , but f has least degree among non zero polynomials satisfied by T , so $r = 0$, i.e. f divides g . Since similarly g divides f , they must be scalar multiples of one another. In particular, if both are monic they are equal. **QED.**

We give a name to the unique monic polynomial of minimal degree satisfied by T .

Defn: If $T:V \rightarrow V$ is a linear map, and $\dim(V)$ is finite, the monic polynomial f of least degree with $f(T) = 0$, i.e. such that $(f(T))(v) = 0$ for all v in V , is called the **minimal polynomial** of T .

Note it follows from the proof of existence of the minimal polynomial that it always has degree $\leq n^2$ where $n = \dim(V)$. In fact the minimal polynomial has degree $\leq n = \dim(V)$, a fact whose proof will be crucial in studying similarity.

Definition: If $T:V \rightarrow V$ is a linear map, $\dim(V)$ is finite, and v is a vector in V , the unique monic polynomial f of least degree with $f(T)(v) = 0$, is called the minimal polynomial of T at v .

Remark: There is some such polynomial since the minimal polynomial for T on all of V works. The uniqueness proof for the one of least degree is also the same.

The next result is key to the ideas of this chapter.

Lemma: If $T:V \rightarrow V$ is a linear map, $\dim(V)$ is finite, and v, w are vectors whose minimal T -polynomials are f, g , then there is a vector u in V whose minimal T -polynomial is the least common multiple of f, g .

Proof: First we prove it in case f, g are relatively prime, in which case their lcm is the product $f.g$. Then we claim that $u = v+w$ works. Since $(f.g)(T) = f(T)og(T) = g(T)of(T)$ does annihilate both v and w it also annihilates their sum. Now let h be any polynomial such that $h(T)$ annihilates $v+w$. We claim $f.g$ divides h , for which it suffices to show that each of f and g do so. Since f annihilates v and h annihilates $v+w$, thus $f.h$ annihilates both v and $v+w$, hence also w . Hence the minimal T -polynomial for w divides $f.h$, so g divides $f.h$. Since f and g are relatively prime, then g divides h . A similar argument shows f also divides h . So indeed $f.g$ is the minimal T -polynomial at $u = v+w$.

Now let the T -minimal polynomials of v, w , namely f, g , be arbitrary. Consider all irreducible factors of f and of g , and let A be the product of those irreducible factors that occur more often in f than in g , and let B be the product of those irreducible factors that occur at least as often in g as in f . Then A and B are relatively prime, and their product $A \cdot B = \text{lcm}(f, g)$. Moreover $f = A \cdot p$, and $g = B \cdot q$, for some polynomials p, q . Then since f, g are the T -minimal polynomials of v, w , it follows that A, B are the T -minimal polynomials of $p(v)$ and $q(w)$. Hence $A \cdot B = \text{lcm}(f, g)$ is the T -minimal polynomial of $u = p(v) + q(w)$. **QED.**

Corollary: If m is the minimal polynomial of T on the space V , then there is a vector w in V such that the minimal polynomial of T at w is also m .

Proof: Since $m(T)$ annihilates every vector in V , it follows that for each vector w , the minimal polynomial of T at w has degree at most that of m . Choose w to be a vector whose minimal polynomial has maximal degree among all vectors in V . Then for any other vector v , we claim the minimal polynomial of T at v divides that at w . It will follow that the minimal polynomial of T at w also annihilates every other v , and hence is the minimal polynomial of T for the whole space.

If the divisibility does not hold, some irreducible factor of the T -minimal polynomial at v occurs to a higher power than in the T -minimal polynomial of w . Then the lcm of the two minimal T -polynomials at v and w has degree greater than either of them. Thus by the lemma there is a vector whose T -minimal polynomial has that greater degree, contradicting the choice of w . **QED.**

Cor: The minimal polynomial of an operator T on V has degree $\leq \dim V$.

Proof: Let v be a vector whose minimal T -polynomial equals that for T on all of V , and consider the evaluation map at v , namely the map $k[X] \rightarrow V$ taking $f(X)$ to $(f(T))(v)$. If $n = \dim(V)$, then the $n+1$ monomials $\{1, X, X^2, \dots, X^n\}$, must have dependent images in V , i.e. the vectors $\{v, T(v), T^2(v), \dots, T^n(v)\}$ are linearly dependent in V . Thus some polynomial of degree $\leq n$ in T vanishes at v , but by choice of v , this polynomial vanishes also on all of V . **QED.**

Revisiting a key example, the derivative operator on polynomials

Now we can completely understand the example of the derivative operator on spaces of polynomials of bounded degree; i.e. recall the derivative D acting on the space V of real polynomials of degree $< n$. The key point is that the n th derivative of any such polynomial is zero, but the $(n-1)$ st derivative of $X^{(n-1)}$ equals $(n-1)!$, which is not zero. Hence the minimal polynomial of D on V divides t^n , but does not divide $t^{(n-1)}$. Hence t^n is the minimal polynomial, and $X^{(n-1)}$ is an

element of V which has the same minimal D -polynomial as the whole space. It follows that this polynomial and its 1st $n-1$ derivatives are independent, hence they form a basis of the space, and one in which the matrix for D is especially simple.

To make it even simpler we choose instead the scalar multiple $X^{(n-1)}/(n-1)!$ as our first basis vector. This gives us the basis $\{X^r/r!\}$ with $0 \leq r \leq n-1$, and this has the nice property that D just shifts each of the first $n-1$ basis vectors to the next basis vector. We call such a basis “cyclic”. Moreover in this special case D annihilates the last basis vector, namely the constant $1 = D^{(n-1)}(X^{(n-1)}/(n-1)!)$, so the matrix in this basis becomes:

$$\begin{bmatrix} 0 & 0 & \dots & 0 \\ 1 & 0 & 0 & \dots & 0 \\ 0 & 1 & 0 & \dots & 0 \\ \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & 0 & \dots & 1 & 0 \end{bmatrix}, \text{ i.e. } 1\text{'s just below the main diagonal, and zeros elsewhere.}$$

We claim that D is similar precisely to those linear operators on V that also have the same minimal polynomial, namely t^n . It suffices to show that any such operator has this same matrix in some basis, since then both operators are similar to the operator defined by this matrix on \mathbb{R}^n . But if $T:V \rightarrow V$ has minimal polynomial t^n , where $n = \dim(V)$, then by our key lemma, there is some vector v in V whose minimal T -polynomial is also t^n . Then for this v , the vectors $\{v, Tv, T^2v, \dots, T^{n-1}(v)\}$ must be independent, while $T^n(v) = 0$. Thus in this basis the matrix for T is also the one above:

$$\begin{bmatrix} 0 & 0 & \dots & 0 \\ 1 & 0 & 0 & \dots & 0 \\ 0 & 1 & 0 & \dots & 0 \\ \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & 0 & \dots & 1 & 0 \end{bmatrix}.$$

Perhaps the simplest model for this operator is the quotient space of all polynomials in X , modulo the subspace of those divisible by X^n , i.e. $V = k[X]/(X^n)$, with operator $T =$ multiplication by X . I.e. then the basis $\{1, X, X^2, \dots, X^{(n-1)}\}$ is shifted to the right by T , until finally $T(X^{(n-1)}) = 0$, and we get the matrix above.

The fundamental cyclic example:

If we generalize the derivative example only slightly, we are led to the fundamental model for all operators on finite dimensional spaces. Certainly, one of the most fundamental and ubiquitous vector spaces over a field k is the space $k[X]$ of polynomials in a variable X , with coefficients in k . To get a finite dimensional such example, we mod out by a large subspace, like the subspace (f) all multiples of some polynomial f . Then look at $V = k[X]/(f)$, the space whose elements are represented by polynomials, but where two polynomials are considered equal if they differ by a multiple of f . Equivalently the polynomial f is set equal to zero. By the division algorithm, we can write any polynomial g as $g = f \cdot h + r$, for some polynomials h and r , where $\deg(r) < \deg(f)$, and then g is equivalent to r . Since two polynomials of degree $< \deg(f)$ cannot differ by a multiple of f unless they are equal, V is represented precisely by those polynomials of degree less than $\deg(f)$. Thus V has dimension $= \deg(f) = n$, with basis (the equivalence classes of) the monomials $\{1, X, X^2, \dots, X^{(n-1)}\}$.

Now one of the simplest linear maps on this space is just multiplication by X . This obviously takes the basis above to the sequence $\{X, X^2, X^3, \dots, X^n\}$. But since f has degree n , and $f = 0$ in this space, we can reduce X^n to some polynomial of degree $\leq n-1$. Thus each of the first $n-1$ basis vectors are taken to the next basis vector. Then if $f = a_0 + a_1X + \dots + a_{n-1} X^{(n-1)} + X^n$, and since $f = 0$ in this space, the last basis vector is taken to $X^n = -a_0 - a_1X - \dots - a_{n-1} X^{(n-1)}$. Thus the matrix of this map, in this basis is the following:

$$\begin{array}{c} |0 \ 0 \dots \dots \dots 0 \ -a_0 \ | \\ |1 \ 0 \ \dots \dots \dots 0 \ -a_1 \ | \\ |0 \ 1 \ \dots \dots \dots 0 \ -a_2 \ | \\ \\ | \dots \dots \dots \dots \dots \dots \dots \ | \\ |0 \ 0 \ \dots \dots \dots 1 \ -a_{n-1} \ | \end{array}$$

We call this matrix C_f = the “companion matrix” of the polynomial f . If you remember that all similar operators have the same minimal polynomial, it follows that this is about the simplest matrix possible for our map. I.e. applying a polynomial g to the map “multiplication by X ”, just gives us the map which is multiplication by $g(X)$. Hence the minimal polynomial is the monic polynomial of least degree such that multiplication by it sends every element of V to zero. This of course is just f . I.e. if T is multiplication by X on the space $k[X]/(f)$, then f is the minimal polynomial of T . In particular f is the also the minimal polynomial of this companion matrix. Then since all similar operators have the same minimal

polynomial, it follows that every matrix for an operator must contain at least the information of that minimal polynomial. Since, except for some 0's and 1's, this companion matrix consists of nothing except the coefficients of the minimal polynomial, it cannot be made much simpler. We shall see later, in the section on Jordan forms, that another option would be to have the matrix display the coefficients of the irreducible factors of the minimal polynomial, or its roots if those lie in the field k . For now we explore the use of the companion matrix.

The key fact about the example we just gave, is that the minimal polynomial of the operator T has maximal degree, i.e. equal to the dimension of the vector space on which T acts. In this case, we have the analogous result to the one proved above about the derivative operator, i.e. such a T is completely determined up to similarity by its minimal polynomial.

Theorem: If $T:V \rightarrow V$ is a linear operator on a finite dimensional space V over k , and if the minimal polynomial f of T has degree $= n = \dim(V)$, then T is similar to the operator in the fundamental example just discussed, multiplication by X on the quotient space $k[X]/(f)$. I.e. T has in some basis the companion matrix C_f of f .

Remark: We are allowing ourselves to call two maps $T:V \rightarrow V$ and $S:W \rightarrow W$ similar if there exists an isomorphism $U:V \rightarrow W$ such that $S = U \circ T \circ U^{-1}$.

Definition: A vector v is called a “cyclic vector” for an operator $T:V \rightarrow V$ on a finite dimensional space V if the minimal polynomial of T on V equals the minimal T -polynomial at v . (These may or may not exist.)

Summary: What we have just proved shows that an operator $T:V \rightarrow V$ on a finite dimensional space V is similar to the standard model, i.e. the fundamental example, if and only if it has a cyclic vector, if and only if the minimal polynomial of T has degree $= \dim(V)$, if and only if in some basis the matrix of T is the companion matrix of the minimal polynomial.

How to construct an operator with no cyclic vector

To see a non cyclic example, we only have to take appropriate products, i.e. we want to look at pairs of operators $S:V \rightarrow V$ and $T:W \rightarrow W$, and consider $(S \times T):V \times W \rightarrow V \times W$, where $(S \times T)(v, w) = (S(v), T(w))$. But we cannot use operators with relatively prime minimal polynomials, since our previous arguments imply that in that case the sum of a cyclic vector for S and a cyclic vector for T would be cyclic for $(S \times T)$.

So to get a non cyclic example, let f and g be two polynomials of positive degree such that f divides g , and take the product of the two standard models $k[X]/(f)$ and $k[X]/(g)$. Then on the product space $V = k[X]/(f) \times k[X]/(g)$, the map T defined as multiplication by X , has minimal polynomial g . Since the minimal polynomial has degree less than $\dim(V)$, there can be no cyclic vector, and the matrix cannot be a single companion matrix. Of course since the space is a product of two subspaces on each of which there is a cyclic vector, we can have a matrix which consists of two blocks, each a companion matrix, one for f and one for g . And this is the best we can do, since no cyclic subspace can have dimension greater than the degree of g . I.e. the subspace $k[X]/(g)$ is a maximal T -cyclic subspace. So the general result is that this is typical. I.e. we will prove the following: if $T:V \rightarrow V$ is a finite dimensional operator with minimal polynomial g , there will be a maximal cyclic subspace, of dimension equal to the degree of g , on which T acts with minimal polynomial g . Then we can decompose V not a product of this subspace and another subspace to which we can apply the same reasoning. I.e. the minimal polynomial of T on this complementary subspace will be some factor of g , and we can again find a maximal cyclic subspace where the minimal polynomial equals that factor, etc.... Then the matrix of T on V will consist of blocks each a companion matrix for one of these subspaces.

Thus in general, the similarity class of T will be determined not just by the minimal polynomial, but by the minimal polynomial and a sequence of factors of that polynomial, each one dividing the next. These polynomials are called “invariant factors” of T and together they determine the similarity class of T . Here is the statement we will prove.

Thm. (Invariant factor theorem): If V is a vector space of finite dimension n over the field of scalars k , and $T:V \rightarrow V$ is a k -linear operator, then there is a unique finite sequence of non constant monic polynomials, g_1, g_2, \dots, g_t in $k[X]$, such that $\deg(g_1) + \dots + \deg(g_t) = n$, with the following properties:

- i) each g_j divides the next one, i.e. $g_1 | g_2, g_2 | g_3, \dots, g_{t-1} | g_t$;
- ii) g_t is the minimal polynomial of T ;
- iii) in some basis for V , the matrix for T consists of t blocks along the diagonal, in which the blocks are the companion matrices for the polynomials g_1, \dots, g_t .

Definition: The product of the polynomials g_j in the theorem is an invariant of the operator T , a monic polynomial of degree equal to $\dim(V)$, the “characteristic polynomial” of T . We will see later it can be computed as a determinant.

T-cyclic decomposition by “invariant factors”

If there is no cyclic vector, we want to decompose our space V into a product of cyclic subspaces. For any decomposition, the subspaces must be preserved by the action of T , i.e. they must be T -invariant, i.e. T must map the subspace into itself. Now if W is a subspace such that $T(W)$ is contained in W , then also $T^2(W) = T(T(W))$ is contained in $T(W)$ hence in W also. Thus W must be closed under the action of every power T^k of T , and since W is a subspace, it must contain every linear combination of every power of T as well. So for every vector w in W , and every polynomial P , W must contain $P(T)(w)$. Thus if we define a multiplication of $k[X]$ on W by $P(X)(w) = P(T)(w)$, W must be closed under multiplication by $k[X]$.

It follows that if W is T invariant and contains a vector w , then W also contains $P(T)(w)$, for every polynomial $P(X)$. Conversely, if w is any vector and W is the set of all vectors of form $P(T)(w)$, we claim W is a T invariant subspace. To see this note that if $P(X)$ is any polynomial, then $X.P(X)$ is also a polynomial, so for any polynomial P , $T(P(T)(w))$ is also a polynomial in T applied to w . Thus if w is any vector in V , and if $k[X]$ is the ring of polynomials in X with coefficients in the scalar field k , then $k[X].w$ (i.e. the set of all vectors of form $P(T)(w)$ with $P(X)$ a polynomial in X), is the smallest T -invariant subspace of V containing w .

So the simplest T invariant subspaces have form $k[X].w$ for some vector w . There is a natural basis for such a subspace. Namely, consider the sequence $\{w, Tw, T^2w, T^3w, \dots\}$. This is a spanning set since all elements of $k[X].w$ are linear combinations of powers of T applied to w . Moreover since we are assuming finite dimensions, some smallest one of these powers, say $T^n(w)$, (and hence also all larger powers), is dependent on earlier ones. Thus some finite sequence $\{w, Tw, T^2w, \dots, T^{n-1}w\}$ is a basis. From the appearance of this basis, we call such a subspace “ T -cyclic”. Since $T^n(w)$ is dependent on the previous powers of $T(w)$, we have $T^n(w) = a_0w + a_1T(w) + a_2T^2(w) + \dots + a_{n-1}T^{(n-1)}(w)$, for some scalars a_0, \dots, a_{n-1} . By definition then, the matrix of T in this basis is the following:

$$\begin{array}{l}
 |0 \ 0 \ 0 \ \dots \ -a_0| \\
 |1 \ 0 \ \dots \ -a_1| \\
 |0 \ 1 \ 0 \ \dots \ \dots \ \dots| \\
 |0 \ 0 \ 1 \ \dots \ \dots \ \dots| \\
 | \ \dots \ \dots \ \dots \ \dots \ \dots \ \dots| \\
 | \ \dots \ \dots \ \dots \ \dots \ \dots \ \dots| \\
 | \ \dots \ \dots \ \dots \ \dots \ \dots \ \dots| \\
 |0 \ \dots \ \dots \ 1 \ -a_{n-1}|
 \end{array}$$

We have called this matrix the “companion matrix” for the polynomial :
 $a_0 + a_1X + a_2X^2 + \dots + a_{n-1}X^{n-1}$.

A proof of the invariant factor theorem by the “splitting” technique

Let $T:V \rightarrow V$ be a linear operator on a finite dimensional vector space over the scalar field k , with minimal polynomial m , and define a multiplication of the polynomial ring $k[X]$ on V by setting $f(X).v = f(T)(v)$. Thus a subspace of V is T invariant if and only if it is closed under multiplication by the ring $k[X]$.

We want to decompose V as a product of T invariant subspaces, each spanned by a single vector under multiplication by $k[X]$, i.e. under multiplication by $k[T]$, hence into a product of T -cyclic subspaces. This is analogous to decomposing V by means of a k basis, since if v_1, \dots, v_n is a k basis then V is isomorphic to the product of the subspaces $k.v_1 \times k.v_2 \times \dots \times k.v_n$. In that earlier case however it was easier to do this, since the fact that k is a field makes two different notions of “dependent” become equivalent.

I.e. if v_1, \dots, v_m is any k spanning set for V , we can throw out any vector that does not depend k - linearly on earlier ones, and we will get a k spanning set, say v_1, \dots, v_n such that no non zero scalar multiple of any v_j is a k linear combination of the others. I.e. if we had $a_1v_1 + \dots + a_nv_n = 0$, with some term, say $a_nv_n \neq 0$, then $a_nv_n = -a_1v_1 - \dots - a_{n-1}v_{n-1}$, and we can divide by a_n to express v_n as a linear combination of the earlier v_j , hence v_n would have already been thrown out.

But if we use polynomial coefficients we can have $f_1v_1 + \dots + f_nv_n = 0$, and even if say $f_nv_n \neq 0$, so that $f_nv_n = -f_1v_1 - \dots - f_{n-1}v_{n-1}$, we still cannot necessarily divide by f_n to express v_n as a $k[X]$ linear expression in the other vectors. We could do so if and only if f_n actually divides the other coefficients f_j with $j \geq 2$.

In order to express V as a product of subspaces of form $k[X].v$, we need a set of $k[X]$ - generators v_1, \dots, v_r for V such that whenever $f_1v_1 + \dots + f_rv_r = 0$, then in fact all terms $f_jv_j = 0$. The problem is that although when dealing with scalar multipliers from a field, two subspaces like $k.v$ and $k.w$ are either the same or meet only in the vector zero. Since some polynomials have non trivial common factors, subspace like $f.v$ and $g.w$ can overlap non trivially without being equal. So we will have to deal with divisibility issues. Fortunately this can be done, since in $k[X]$ any set of polynomials have a greatest common divisor. Still it takes a little work.

We will proceed by induction of the k dimension of V . Since there is no need to decompose a one dimensional space, we will assume V has larger dimension and

that the decomposability into T cyclic subspaces is true for all spaces of smaller dimension. So we want to see how to split off a T -cyclic subspace of V , with the complementary factor still T -invariant.

So, if $T:V \rightarrow V$ is a linear operator on a finite dimensional vector space V over the scalar field k , with minimal polynomial m , we have a multiplication of the polynomial ring $k[X]$ on V by setting $f(X).v = f(T)(v)$. Then a subspace of V is T invariant if and only if it is closed under multiplication by the ring $k[X]$.

Now let v be a vector such that the minimal polynomial of T at v equals the minimal polynomial m of T on all of V . Then it follows that the minimal polynomial of T on every other vector in V is a factor of m . This is the divisibility property we need to make our argument work. We will show that V is isomorphic to a product of the T -cyclic subspace $k[X].v$ and another T -invariant subspace. Then we can use the inductive hypothesis to decompose that other T invariant subspace into T -cyclic subspaces.

Splitting Lemma: If the minimal polynomial of T at v equals the minimal polynomial m of T on all of V , then there is a T - invariant subspace W of V , such that V is isomorphic to $k[X].v \times W$

Proof:

First we mod out V by the T invariant subspace generated by v , i.e. by $k[X].v$, getting a space $V/(k[X].v)$ on which T induces a k - linear operator S , hence also a multiplication by the polynomial ring $k[X]$. Since the map S on the quotient is induced by T , the quotient map $V \rightarrow V/(k[X].v)$ is not only k linear but also $k[X]$ linear, i.e. it commutes with multiplication by polynomials as well as scalars. Note that the kernel of this map is the T -cyclic subspace $k[X].v$.

Now we know that abstractly V is isomorphic as k - vector space to the product of the quotient and the kernel, i.e. that $V \approx (k[X].v) \times V/(k[X].v)$ as k - vector spaces, but we want to show this also holds as $k[X]$ - subspaces of V . I.e. we want the isomorphism of $V/(k[X].v)$ with a k -subspace of V to also respect the action of the operator T . So we need to find a T -invariant subspace W of V which meets $k[X].v$ only in $\{0\}$, i.e. a $k[X]$ - subspace subspace W that is complementary to $k[X].v$.

To do this it suffices to find a $k[X]$ - linear map which is right inverse to the quotient map $V \rightarrow V/(k[X].v)$. I.e. we want a map $V/(k[X].v) \rightarrow V$ such the composition $V/(k[X].v) \rightarrow V \rightarrow V/(k[X].v)$ is the identity, and in order for the image of the map $V/(k[X].v) \rightarrow V$ to be T -invariant, we want the map to be $k[X]$ -linear. Then its image in V will be W , the complementary subspace to $k[X].v$.

(Note that the fact the map $V/(k[X].v) \rightarrow V \rightarrow V/(k[X].v)$ is the identity, hence injective, implies $k[X].v \rightarrow V$ is injective and meets the kernel of the quotient map, namely $k[X].v$, only in $\{0\}$.)

We know if T acts on a space U with minimal polynomial f at a vector u , then the T -cyclic subspace $k[X].u = k[T].u$ is isomorphic to $k[X]/(f)$, where the action of T on the subspace $k[T].u$ corresponds to multiplication by X on the quotient space $k[X]/(f)$.

By induction we can decompose the quotient space $V/(k[X].v)$ into a product of T -cyclic subspaces, hence the map $V \rightarrow V/(k[X].v)$ becomes a $k[X]$ linear map $V \rightarrow k[X]/(f_1) \times \dots \times k[X]/(f_s)$, with kernel $k[X].v$. Moreover, the minimal polynomial of T on V and the minimal polynomial of T at v , are both equal m , while each minimal polynomial f_j of X acting on $k[X]/(f_j)$ is a factor of m .

For each j , let u_j represent the equivalence class in $k[X]/(f_j)$ of the cyclic vector 1 , so that each $k[X]/(f_j) \approx k[X].u_j$. Then to define a $k[X]$ - linear, right - inverse of the map $Q: V \rightarrow k[X].u_1 \times \dots \times k[X].u_s$, it suffices to define it on each factor separately. So for each j , we first seek a vector w_j such that $Q(w_j) = u_j$. If also $f_j.w_j = 0$, where f_j is the minimal polynomial of S acting on u_j , then we can define the right inverse map $k[X].u_j \rightarrow V$ by sending u_j to w_j , and $g(X).u_j$ to $g(T)(w_j)$. I.e. that will define a map $k[X] \rightarrow V$ that sends f to zero, hence induces a $k[X]$ - linear map from the quotient $k[X]/(f)$ to V .

Since Q is surjective, it is always possible to find some vector z_j that maps to u_j , i.e. with $Q(z_j) = u_j$, and the challenge is to find a preimage that is annihilated by f_j . This is where the divisibility property mentioned above will come to our rescue, i.e. the fact that the minimal polynomial m of T at v , is a multiple of the minimal polynomial f_j of S at u_j .

If z_j is any vector in V with $Q(z_j) = u_j$, then $f_j.z_j$ will map by Q to $f_j.u_j = 0$, so at least $f_j.z_j$ lies in the kernel $k[X].v$ of Q . But we want it to be zero. Now if we change z_j by anything in the kernel of Q , it will still map to u_j . So we want to find something, say y_j in the kernel of Q , such that $f_j.(z_j - y_j) = 0$. This of course would require that $f_j.z_j = f_j.y_j$, with y_j in the kernel of Q . Next we see how to do this.

Since $f_j.z_j$ lies in the kernel $k[X].v$ of Q , thus $f_j.z_j = g.v$ for some polynomial g . And f_j divides m , so $m = f_j.A$ for some polynomial A . Since m annihilates all of V , thus $A.f_j.z_j = m.z_j = 0$, and since $f_j.z_j = g.v$, hence also $A.g.v = 0$. Now the minimal

polynomial of v is m , so m must divide $A \cdot g$, i.e. $m = A \cdot f_j$ must divide $A \cdot g$. But this implies that f_j must divide g , since if $A \cdot f_j \cdot B = A \cdot g$, then canceling A gives $f_j \cdot B = g$.

Now if $g = B \cdot f_j$, then $g \cdot v = B \cdot f_j \cdot v$, so we can choose $y_j = Bv$, and then setting $w_j = z_j - y_j$, gives us a vector w_j such that $Q(w_j) = Q(z_j) = u_j$, and also $f_j \cdot w_j = f_j \cdot z_j - f_j \cdot y_j = f_j \cdot z_j - f_j \cdot B \cdot w_j = g \cdot v - g \cdot v = 0$.

Then mapping each u_j back to the corresponding w_j , and mapping $g \cdot u_j$ to $g \cdot w_j$, defines a $k[X]$ -linear map R , right inverse to $Q: V \rightarrow k[X] \cdot u_1 \times \dots \times k[X] \cdot u_s$, such that the image W of R in V , is a T invariant subspace complementary to $k[X] \cdot v$, as desired. **QED.**

Cor: (invariant factor theorem):

We have split $V \approx W \times k[X] \cdot v \approx W \times k[X]/(f_1)$, as a product of T -invariant subspaces where $k[X] \cdot v \approx k[X]/(f_1)$ is T -cyclic. By induction on dimension, W is also a product of T -cyclic subspaces $W \approx k[X]/(f_2) \times \dots \times k[X]/(f_s)$. This splits V as a product of T -cyclic subspaces $V \approx k[X]/(f_1) \times \dots \times k[X]/(f_s)$.

Applying the appropriate number of powers of T to the cyclic vectors $v = v_1, \dots, v_s$ gives us a k -basis of V in which the matrix of T has the desired form. In particular, note that by induction the minimal polynomials f_j successively divide each other, as well as dividing m . **QED.**

Next we want to give a matrix algorithm to compute not only the minimal polynomial of a given matrix, but the full sequence of invariant factors that determine its similarity class.

Computing the invariant factor decomposition from a “ $k[X]$ - presentation”

The lesson from the decomposition theorem we have proved is that, as a $k[X]$ space, the vector space V with operator T , is isomorphic to a product of cyclic $k[X]$ spaces, of form $k[X]/(f)$ where f is some polynomial. Thus, as a $k[X]$ space, $V \approx k[X]/(f_1) \times \dots \times k[X]/(f_s)$, is a quotient of $k[X]$ spaces, namely $V \approx (k[X] \times \dots \times k[X]) / ((f_1) \times \dots \times (f_s))$. Equivalently there is a surjective map of $k[X]$ spaces $k[X] \times \dots \times k[X] \rightarrow V$ whose kernel is $(f_1) \times \dots \times (f_s)$. Our goal is to find such a representation starting just from a matrix for T , i.e. we want a surjection and then we want to calculate the kernel.

First we want a $k[X]$ -spanning set for V , and the simplest choice is a k -basis, $\{v_1, \dots, v_n\}$, in which the map T is represented by an n by n matrix A over k . Now

this k -basis is necessarily a $k[X]$ spanning set, since $k[X]$ contains k . This then defines a surjective $k[X]$ linear map $(k[X])^n \rightarrow V$, taking the standard basis vector $e_j = (0, \dots, 0, 1, 0, \dots, 0)$, where the 1 is in the j th position, to the basis vector v_j . Since the map is extended to be $k[X]$ linear, the vector $(0, \dots, 0, X, 0, \dots, 0) = X \cdot e_j$, must go to $X \cdot v_j = T(v_j)$. Now if A is the matrix for T in the given basis v_1, \dots, v_n , then the j th column of A also describes a vector that equals $T(v_j)$, i.e. the entries (b_1, \dots, b_n) in that column are the coefficients of the expansion of $T(e_j)$ in the basis: $T(e_j) = b_1 v_1 + \dots + b_n v_n$. Hence under the surjective map $k[X]^n \rightarrow V$, both vectors $X \cdot e_j$ and (b_1, \dots, b_n) map to $T(v_j)$, and thus their difference $X e_j - b_1 e_1 - \dots - b_n e_n$ maps to zero.

Hence the columns of the matrix $X \cdot I - A$ belong to the kernel of the surjective map $k[X]^n \rightarrow V$. We will show below that those columns span that kernel.

Consequently the matrix $X \cdot I - A$ defines a $k[X]$ -linear map $k[X]^n \rightarrow k[X]^n$ whose image, the span of its columns, equals the kernel of the surjection $k[X]^n \rightarrow V$. Exactly as in the case of k linear maps, the induced map $k[X]^n / \text{Im}(X \cdot I - A) \rightarrow V$, is $k[X]$ linear and bijective, hence defines an isomorphism as $k[X]$ spaces. Thus the T - cyclic structure, or $k[X]$ structure of V , is displayed by the matrix $X \cdot I - A$.

This computation is analogous to expressing the k vector space structure of V as a product of copies of k . This is contained at least implicitly in the matrix computations from chapter 2. I.e. choosing a k -spanning set for V defines a surjective k linear map $k^m \rightarrow V$ with a kernel U , so that V is isomorphic to the quotient space k^m / U . The problem is to compute this quotient space. In order to take advantage of matrix operations we proceed as follows. Choose a k -spanning set also for U , which gives another surjective map $k^s \rightarrow U$. Then by composition with $U \rightarrow k^m$, we get a map $A: k^s \rightarrow k^m$ which is defined by a matrix A . Then V is isomorphic to the quotient space $k^m / \text{Im}(A)$, the “cokernel” of the map A .

Definition: A sequence of linear maps $k^n \rightarrow k^m \rightarrow V$ where $k^m \rightarrow V$ is a surjection whose kernel is the image of the map $A: k^n \rightarrow k^m$, is called a “presentation” of V . Such a presentation induces an isomorphism of V with the quotient space $k^m / \text{Im}(A) =$ the “cokernel” of the matrix A .

Representing a space as the cokernel of a matrix A has the advantage that such cokernels can be computed by diagonalizing the matrix A using elementary matrix operations. This was accomplished for k linear maps in chapter one when we computed the equivalence class of a matrix. I.e. using both row and column operations on the m by n matrix $A: k^n \rightarrow k^m$, we get a diagonal matrix B where say the first r columns are standard basis vectors, and all other entries are zeroes. Thus the image of this map in k^m is just the span of the first r coordinate axes, so

the quotient space $k^m/\text{Im}(A)$ is isomorphic to the quotient space $k^m/\text{Im}(B) = (k \times \dots \times k)/(k \times \dots \times k \times \{0\} \times \dots \times \{0\})$, where there are m copies of k in the top, and r copies of k in the bottom. Hence the quotient is visibly isomorphic to $(k/k) \times \dots \times (k/k) \times (k/\{0\}) \times \dots \times (k/\{0\}) \approx \{0\} \times \dots \times \{0\} \times k \times \dots \times k \approx k^{(m-r)}$.

Next we want to imitate this procedure to compute the $k[X]$ structure of a space V and a map $T:V \rightarrow V$.

Definition: We call $[X.I-A]$, the “characteristic matrix” of A . Its determinant, $\det(X.I-A) = \text{ch}(A)(X)$ is the “characteristic polynomial” of A .

Remark: If A is a matrix for an operator T , defined by a basis, then the characteristic polynomial is an invariant of T as well as of A since the characteristic polynomials of any two matrices for a given operator T are the same. I.e. if A, B are two matrices for T , then they are similar, so there is some invertible matrix U such that $B = U^{-1}AU$. In particular, $\det(B) = \det(U^{-1}AU) = \det(U^{-1}) \cdot \det(A) \cdot \det(U) = \det(U^{-1}) \cdot \det(U) \cdot \det(A) = \det(U^{-1} \cdot U) \cdot \det(A) = \det(I) \cdot \det(A) = 1 \cdot \det(A) = \det(A)$. But if A and B are similar, then their characteristic matrices are also similar, since then $U^{-1}(X.I-A)U = U^{-1}(X.I)U - U^{-1}AU = X.I.U^{-1} \cdot U - U^{-1}A \cdot U = X.I - U^{-1}A \cdot U = X.I - B$. Thus if A, B are two matrices for T , then $\text{ch}(A)(X) = \text{ch}(B)(X)$. Thus we may call this common polynomial the characteristic polynomial of the operator T .

The characteristic presentation:

Proposition: In the case of the $k[X]$ structure defined on V by the n by n matrix A , we have the presentation $k[X]^n \rightarrow k[X]^n \rightarrow V$, where the map $(X.I-A):k[X]^n \rightarrow k[X]^n$ is given by the “characteristic matrix” $X.I-A$. Thus as $k[X]$ spaces, $V \approx k[X]^n/\text{Im}(X.I-A)$.

Proof: To prove this it remains to show that the columns of $[X.I-A]$ do span the kernel of the natural surjection $k[X]^n \rightarrow V$ defined by a basis of V . Since we know those columns do belong to the kernel of the surjection, it follows that the induced map $k[X]^n/\text{Im}(X.I-A) \rightarrow V$ is surjective. Consequently to deduce that it is an isomorphism, it would suffice to show the k -dimension of the quotient $k[X]^n/\text{Im}(X.I-A)$ is the same as the k -dimension of V , namely n . One way to show this is by the theory of determinants. The fact that the entries on the main diagonal are all of degree one in X , and all other entries are constants, implies the degree of the determinant of $X.I-A$ is n . Then we claim the k -dimension of the cokernel of $X.I-A$ equals the degree of this determinant. This will follow from the fact proved below that this matrix can be diagonalized. I.e. the determinant does not change during diagonalization, and after diagonalization the determinant equals

the product of the diagonal entries. Since the degree of that product also equals the dimension of the cokernel, this proves our result. **QED.**

Remark: There is also a nice direct argument for the previous proposition, in Jacobson's beautiful book Basic Algebra I, p.190.

Note there is no $k[X]$ linear combination of the columns of $X.I - A$ that equals zero, since that would imply the determinant is zero, whereas it is a monic (hence non zero) polynomial of degree n . Hence the characteristic matrix defines an injective map, and so the kernel subspace K is isomorphic to $k[X]^n$ itself. Of course although isomorphic, these two spaces are not equal, since their quotient space $k[X]^n/K \approx V$ is not zero.

Diagonalizing the characteristic matrix

The admissible row and column operations consist of,

- 1) interchanging two rows or two columns;
- 2) multiplying through a row or a column by an invertible polynomial, i.e. a non zero scalar;
- 3) adding to any row (or to any column), a polynomial multiple of any other row (or of any other column).

These are the exact analogs of the elementary operations involving scalar matrices, but now applied to polynomial matrices. In particular, in order for the operations to be invertible, i.e. reversible, we must use only invertible multipliers in step 2. The fact that some non zero polynomials are not invertible means we will not be able to reduce the diagonal elements to constants.

Proposition: The characteristic matrix $X.I - A$ associated to an n by n scalar matrix A , can be diagonalized by row and column operations. This can be done so that the diagonal entries successively divide each other. When completed, the diagonal entries will be a sequence of monic polynomials, which divide each other successively, and whose product equals the characteristic polynomial.

Remark: The sequence of diagonal entries will usually begin with a certain number of copies of the constant monic polynomial $= 1$.

Proof: The elementary operations allow us to form linear combinations of elements belonging to the same row or column, and hence, by the Euclidean algorithm in $k[X]$, to replace any element by the gcd of all elements in its same row or same column. Hence we can make the entry in position $(1,1)$ equal to the

gcd of all entries in the first row, using only column operations. (Note that since the determinant of $[X.I-A]$ is not zero, the rows, and columns, are not $k[X]$ dependent, and in particular there is a non zero entry in every row and every column.)

Now that we have the $(1,1)$ entry dividing every other entry in the first row, we can use entry $(1,1)$ to replace every other entry in the first row by zero, again by column operations. Now we can use the same procedure on the elements of the first column, using row operations, until the only non zero entry in the first row and first column is the $(1,1)$ entry. Now ignoring the first row and column we can proceed to the second row, and by induction we can diagonalize the rest of the matrix, thus obtaining a matrix that is completely diagonal. Since the determinant is non zero, all diagonal entries will be non zero.

Now that the matrix is diagonal, its determinant equals the product of the polynomials along the diagonal, so the sum of the degrees of these polynomials must equal $n = \dim(V)$. (If we make all the polynomials monic, the determinant actually equals the original determinant, not only in degree.) Also since the matrix is diagonal, the quotient space of $k[X]^n$ by the $k[X]$ -span of the columns is $k[X]$ -isomorphic to the product of the quotients $k[X]/(f_1) \times \dots \times k[X]/(f_n)$, where f_1, \dots, f_n is the sequence of polynomials along the diagonal. Consequently the k -dimension of this space is also equal to the sum of the degrees of the polynomials f_j , namely it equals $n = \dim(V)$. Thus the quotient space not only maps onto V , but isomorphically to it since the two spaces have the same k -dimension. We thus have a decomposition of V as a product of T -cyclic subspaces corresponding under the isomorphism to the factors $k[X]/(f_j)$.

We can improve this decomposition so that the polynomials f_j successively divide each other. Namely after the matrix is diagonal, we can add every later column to the first column so that all diagonal entries are in the first column. Now repeating the earlier procedure using the Euclidean algorithm, we can replace the entry $(1,1)$ by the gcd of all the entries in the first column, which element is then the gcd of every entry in the matrix. Then we repeat our earlier work to make all other entries zero in the first row and column, leaving in position $(1,1)$ an element that still divides every entry in the matrix. Now ignoring the first row and column, proceed to the rest of the matrix, and by induction we can obtain a diagonal matrix in which each diagonal entry divides the next one. It may be true now that some of the earlier diagonal entries equal 1. For instance if the minimal polynomial of T has degree $n = \dim(V)$, then all diagonal entries will be 1 except the last entry, which will be the minimal polynomial.

After the refinement just described, it follows in general that the last diagonal entry is a multiple of all the others, hence equals the minimal polynomial of A , hence of T . After diagonalizing $X.I-A$, we have remarked the product of all the (monic) diagonal entries equals the original determinant of $X.I-A$, the characteristic polynomial.

Corollary (Cayley - Hamilton theorem): The minimal polynomial divides the characteristic polynomial. All irreducible factors of the characteristic polynomial occur in the minimal polynomial.

Proof: This follows from the refined diagonalization discussed above. **QED.**

Corollary: If the characteristic polynomial is a product of distinct irreducible factors, then the minimal and characteristic polynomial are equal.

Remark: One can describe the diagonalization procedure above theoretically as follows: begin by making the $(1,1)$ entry as small as possible, using the Euclidean algorithm, i.e. row and column operations. It then divides every element in the matrix. Use it to make all other entries in 1st row and 1st column equal to zero. Now proceed to the rest of the matrix, and by induction, one obtains the final diagonal form described above.

In practice this suggests that one should begin by transferring the smallest visible non zero element in the matrix to the $(1,1)$ entry, by interchanging rows and columns, before starting the Euclidean algorithm. This minimal $(1,1)$ entry will in general be a unit, hence in particular will already be the gcd of the whole matrix. The advantage of this procedure is that it can be carried out in practice using nothing more than the Euclidean algorithm in $k[X]$, which is reasonably efficient. Hence at least in theory it should be applicable to actual examples, not only “cooked” ones, but as we know, real life is messy.

Terminology: This normal form may be called the “rational canonical form”, or decomposition by “invariant factors”.

Remark: In case the minimal polynomial is irreducible, then all invariant factors are equal to it, since they are non constant factors of the minimal polynomial. Hence the rational canonical matrix is in block form with r copies of the companion matrix of the minimal polynomial along the diagonal, where $\dim(V)$ equals r times the degree of the minimal polynomial.

Examples of diagonalization:

I have a really hard time not making mistakes in these calculations so let's keep these examples simple. Take this matrix over the rational field, $A =$

$$\begin{vmatrix} 1 & 0 & 1 \\ 1 & 1 & 0 \\ 0 & 1 & 1 \end{vmatrix}, \text{ hence the characteristic matrix is } [X.I-A] =$$

$$\begin{vmatrix} X-1 & 0 & -1 \\ -1 & X-1 & 0 \\ 0 & -1 & X-1 \end{vmatrix}, \text{ now interchange 1st and 3rd columns to get a unit at } (1,1):$$

$$\begin{vmatrix} -1 & 0 & X-1 \\ 0 & X-1 & -1 \\ X-1 & -1 & 0 \end{vmatrix}, \text{ now add } (X-1) \text{ times the first row to the 3rd row:}$$

$$\begin{vmatrix} -1 & 0 & X-1 \\ 0 & X-1 & -1 \\ 0 & -1 & (X-1)^2 \end{vmatrix}, \text{ multiply 1st row by } -1, \text{ and add } (X-1) \text{ times 1st column to 3rd:}$$

$$\begin{vmatrix} 1 & 0 & 0 \\ 0 & X-1 & -1 \\ 0 & -1 & (X-1)^2 \end{vmatrix}, \text{ focus on lower right 2 by 2 submatrix, interchange rows:}$$

$$\begin{vmatrix} 1 & 0 & 0 \\ 0 & -1 & (X-1)^2 \\ 0 & (X-1) & -1 \end{vmatrix}, \text{ now add } (X-1) \text{ times second row to third row:}$$

$$\begin{vmatrix} 1 & 0 & 0 \\ 0 & -1 & (X-1)^2 \\ 0 & 0 & (X-1)^3 - 1 \end{vmatrix}, \text{ multiply 2nd row by } -1, \text{ then can kill off the } (X-1)^2:$$

$$\begin{vmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & (X-1)^3 - 1 \end{vmatrix}, \text{ ok, we got it! "simplify":}$$

$$\begin{vmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & X^3 - 3X^2 + 3X - 2 \end{vmatrix}, \text{ and I can even factor this over } \mathbb{Q}:$$

$$\begin{vmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & (X-2)(X^2-X+1) \end{vmatrix}.$$

Well that's it. And this seems to be what should usually happen randomly, i.e. there is only one diagonal entry that is not a '1', so the minimal polynomial equals the characteristic polynomial, and the whole space has a cyclic vector. But of course the cooked examples you find in books will not usually do this. Anyway we have from the unfactored form, the coefficients of the minimal polynomial that give us the companion matrix representing A in some appropriate basis, as follows:

$$\begin{vmatrix} 0 & 0 & 2 \\ 1 & 0 & -3 \\ 0 & 1 & 3 \end{vmatrix}.$$

If you want to know what basis gives this matrix, you need a cyclic vector for A. There is an algorithm for computing it by tracking all the steps of the diagonalization, given in the very nice book by Dummit and Foote. In this case however it seems most vectors should be cyclic, and if we just try e1, it works. I.e. {e1, A(e1), A^2(e1)} is a basis, and a cyclic basis, which thus gives this matrix. In particular it seems there are lots of cyclic bases and any one gives this matrix.

I also used this to check myself, by computing {e1, A(e1), A^2(e1), A^3(e1)}, and finding a relation between them, i.e. thus computing again the minimal polynomial. Remember, since {e1, A(e1), A^2(e1)} is a basis, e1 is a cyclic vector, and thus its minimal polynomial is minimal for the whole space.

Example:

Here is another easy one; let A be this matrix:

$$\begin{vmatrix} 1 & 1 & -1 \\ 1 & -1 & 1 \\ -1 & 1 & 1 \end{vmatrix}, \text{ then } X.I-A \text{ is the following:}$$

$$\begin{vmatrix} (X-1) & -1 & 1 \\ -1 & (X+1) & -1 \\ 1 & -1 & (X-1) \end{vmatrix}; \text{ interchange 1st and 3rd columns:}$$

$$\begin{vmatrix} 1 & -1 & (X-1) \end{vmatrix}$$

$$\begin{array}{ccc|c} -1 & (X+1) & -1 & \\ \hline (X-1) & -1 & 1 & \end{array}; \text{ add } (X-1) \text{ times 2nd row to 3rd row:}$$

$$\begin{array}{ccc|c} 1 & -1 & (X-1) & \\ \hline -1 & (X+1) & -1 & \\ 0 & (X^2-2) & (2-X) & \end{array}; \text{ add 1st row to 2nd row:}$$

$$\begin{array}{ccc|c} 1 & -1 & (X-1) & \\ \hline 0 & X & (X-2) & \\ 0 & (X^2-2) & (2-X) & \end{array}; \text{ use 1st column to zero out 1st row:}$$

$$\begin{array}{ccc|c} 1 & 0 & 0 & \\ \hline 0 & X & (X-2) & \\ 0 & (X^2-2) & (2-X) & \end{array}; \text{ negate 3rd column, then add to 2nd column:}$$

$$\begin{array}{ccc|c} 1 & 0 & 0 & \\ \hline 0 & 2 & (2-X) & \\ 0 & (X^2+X-4) & (X-2) & \end{array}; \text{ multiply 3rd row by } -2, \text{ then add } (X^2+X-4) \text{ times 2nd row to 3rd row:}$$

$$\begin{array}{ccc|c} 1 & 0 & 0 & \\ \hline 0 & 2 & (2-X) & \\ 0 & 0 & (2-X)(X^2+X-2) & \end{array}; \text{ divide 2nd column by 2, then add } (X-2) \text{ times 2nd column to 3rd column:}$$

$$\begin{array}{ccc|c} 1 & 0 & 0 & \\ \hline 0 & 1 & 0 & \\ 0 & 0 & (X^3 - X^2 - 4X+4) & \end{array}; \text{ which gives the following companion matrix similar to } A:$$

$$\begin{array}{ccc|c} 0 & 0 & -4 & \\ \hline 1 & 0 & 4 & \\ 0 & 1 & 1 & \end{array} \approx A.$$

In particular, again the characteristic and minimal polynomials are equal, both to $(X^3 - X^2 - 4X+4)$. To check this then we could use determinants to compute the characteristic polynomial of A , as follows: i.e. $\det(X.I-A) =$

$$\det \begin{array}{ccc|c} (X-1) & -1 & 1 & \end{array}$$

$$\begin{vmatrix} -1 & (X+1) & -1 \\ 1 & -1 & (X-1) \end{vmatrix}, \text{ expanding by LaGrange formula in 1st row:}$$

$$= (X-1)(X^2-2) + (2-X) - X = X^3 - X^2 - 4X + 4. \text{ check!}$$

Refinements: In both these examples, the characteristic polynomial is easy to factor into irreducible factors over \mathbb{Q} . In the 1st example above we have two distinct irreducible factors $(X-2)(X^2-X+1)$. In the next chapter we will learn how to refine the companion matrix to display these irreducible factors, i.e. the general “Jordan form”, giving us a block matrix over \mathbb{Q} , consisting of one companion matrix for each factor, i.e. this matrix:

$$\begin{vmatrix} 2 & 0 & 0 \\ 0 & 0 & -1 \\ 0 & 1 & 1 \end{vmatrix}.$$

We will also learn how to further refine this matrix over the complex field, by factoring the polynomial fully, yielding a diagonal matrix with the roots of this polynomial, namely $\{2, (1+i\sqrt{3})/2, (1-i\sqrt{3})/2\}$, on the diagonal.

In the second example just above, the polynomial $(X^3 - X^2 - 4X + 4)$ factors completely over \mathbb{Q} as $(X-1)(X-2)(X+2)$. This will allow us to display the matrix in Jordan form as fully diagonal over \mathbb{Q} , with the roots 1,2,-2 on the diagonal.

Here is another example I did, that you can do on your own. Let $A =$

$$\begin{vmatrix} 1 & 0 & -1 \\ 0 & -1 & 1 \\ -1 & 1 & 0 \end{vmatrix}.$$

My answer was $X(X^2-3) =$ both the characteristic and minimal polynomial. This means we get a single companion matrix as rational canonical form as in this chapter, and in the next chapter, a Jordan matrix with two blocks of companion matrices over \mathbb{Q} , one for each factor. By passing to the reals, we will get a diagonal Jordan matrix with some irrational numbers on the diagonal.

Remark: A key point is that the theory in this chapter yields a “rational canonical” matrix with the minimum number of companion matrices, one for each “invariant factor” of the characteristic polynomial, i.e. one for each non constant diagonal entry in the diagonalized characteristic matrix. In the next chapter we will refine

this to a “Jordan form” which displays a companion matrix for each irreducible factor of the characteristic polynomial. The thing to remember is that the cruder rational canonical form is effectively computable by the techniques in this chapter. The more refined Jordan form discussed in the next chapter does reveal more information, but only in case the characteristic polynomial can be factored completely into irreducible factors over the desired field. Unfortunately, there is no practical computational technique for obtaining such a factorization even over \mathbb{Q} , except for low degrees, such as 3 by 3 matrices or smaller. Thus although we will prove the existence of a Jordan form, we will not be able to give computational methods for actually finding it in full generality, unless the irreducible factors of the characteristic polynomial are known.

Chapter Four: General and classical Jordan form

In the previous chapter, given a linear map $T:V \rightarrow V$ with $\dim(V)$ finite, we proved the existence of a decomposition of V into a product of subspaces, each of which has a T -cyclic vector, and hence on each of which the corresponding matrix of the restricted map is a companion matrix of an “invariant factor” of the characteristic polynomial. The decomposition we found was as crude as possible, in the sense that the cyclic factor spaces were as large as possible. In particular if the entire space V has a T -cyclic vector, then the invariant factor decomposition has only one piece and the matrix of T on the full space V is the companion matrix of the characteristic polynomial of T .

In general this is the best we can do, because it is so difficult in practice to factor polynomials, but in cases where we can factor the characteristic polynomial into irreducible factors, e.g. for small matrices, or any triangular matrix, then we can further decompose the space V into cyclic factors which are as small as possible.

Indecomposable T - invariant subspaces

A space with an operator T , is “decomposable” (with respect to the action of T), if it splits into a product of two non zero T - invariant subspaces, and is “indecomposable” otherwise. We claim that any space with operator whose minimal polynomial has more than one irreducible factor is decomposable, and that a space with operator is indecomposable if and only if it is both cyclic, and has minimal polynomial equal to a power of a single irreducible polynomial. We will use the following relatively prime decomposition lemma.

Decomposition Lemma: If $T:V \rightarrow V$ satisfies the polynomial $P =$

$(f_1(X))^{r_1}(f_2(X))^{r_2}\dots(f_t(X))^{r_t}$ with the f_i distinct, irreducible polynomials, then V is isomorphic to the product of the subspaces $V_i = \ker(f_i^{r_i}(T))$.

Proof: As always, to show two spaces are isomorphic we look for a natural map from one to the other. Since the subspaces V_i each map by inclusion to V , there is an induced linear map from the product of the V_i to V . We claim this map, taking (w_1, \dots, w_t) to $w_1 + \dots + w_t$, is injective and surjective. The trick is to use the fact that products involving distinct factors are relatively prime, and then apply the Euclidean algorithm. First define polynomials P_1, \dots, P_t , where each P_i is the product of all factors of the minimal polynomial except the i th one. Thus

$$P_1 = (f_2(X))^{r_2}\dots(f_t(X))^{r_t},$$

$$P_2 = (f_1(X))^{r_1}(f_3(X))^{r_3}\dots(f_t(X))^{r_t},$$

$$\dots, P_t = (f_1(X))^{r_1}(f_2(X))^{r_2}\dots(f_{t-1}(X))^{r_{t-1}}.$$

Since there is no irreducible factor common to all of these polynomials, the Euclidean algorithm gives polynomials Q_1, \dots, Q_t such that $P_1Q_1 + \dots + P_tQ_t = 1$. (I.e. recall the Euclidean algorithm says that the greatest common divisor of a set of polynomials over a field can be written as a finite linear combination, with polynomial coefficients, of the elements of that set, and the gcd of a set with no common irreducible factor is one.)

Hence for any vector w in V , we have $w = 1 \cdot w =$

$P_1(T) \circ Q_1(T)(w) + \dots + P_t(T) \circ Q_t(T)(w)$ is a sum of images of the polynomials $P_i(T)$.

Since (by definition of the $P_i(T)$) composing $(f_i(T))^{r_i}$ with $P_i(T)$ is zero, $\text{Im}(P_i(T))$ is in $\ker(f_i(T))^{r_i} = V_i$, so we have proved every vector in V is a sum of vectors in the V_i . This proves surjectivity of the map $V_1 \times \dots \times V_t \rightarrow V$.

For injectivity, assume $(w_1, \dots, w_t) \rightarrow (w_1 + \dots + w_t) = 0$. Then $w_i =$

$-w_1 - \dots - w_{i-1} - w_{i+1} - \dots - w_t$ is a linear combination of the other w_j , hence each w_i lies in the kernel of P_i , since all the other w_j do. But, for the same reason, each w_i also lies in the kernel of every P_j with $j \neq i$, hence each w_i is in the kernel of $P_1(T) \circ Q_1(T) + \dots + P_t(T) \circ Q_t(T) = \text{Id}$, so every $w_i = 0$. This proves injectivity.

QED.

Remark: The decomposition lemma is true for any polynomial satisfied by the operator T , but if the polynomial is not minimal, some of the subspaces V_i may be $\{0\}$. In any case, we claim the operator T maps each subspace V_i into itself, i.e.

the subspaces V_i are “invariant” under T . This is true because T commutes with any polynomial in T . Thus if w lies in V_i , i.e. if $P_i(T)(w) = 0$, then also $T((P_i(T))(w)) = 0 = (T \circ P_i(T))(w) = P_i(T)(T(w))$, so $T(w)$ is also in V_i .

Thus the restriction of T to each V_i is an operator on V_i , and hence has itself a minimal polynomial, which must be a factor of the minimal polynomial of T . Not surprisingly, if $P = (f_1(X))^{r_1}(f_2(X))^{r_2}\dots(f_t(X))^{r_t}$ is the minimal polynomial of T on V , then we claim $f_i^{r_i}(X)$ is the minimal polynomial of the restriction of T to $\ker(f_i^{r_i}(T))$. I.e. if the restriction of T to V_i satisfies some polynomial, then as argued above the product of these polynomials annihilates T on V . (Just apply the product to a basis composed of bases of the subspaces V_i .) Hence if any one restriction satisfied a polynomial of lower degree than r_i , substituting this polynomial for the factor $f_i^{r_i}(X)$ of P , gives a polynomial of lower degree than P satisfied by T , a contradiction since P is the minimal polynomial for T .

Corollary: Given $T:V \rightarrow V$, a T -invariant subspace U of V is indecomposable if and only if U is T -cyclic and has restricted minimal polynomial equal to a power of an irreducible polynomial.

Proof: The previous lemma shows that every T -invariant subspace on which the minimal polynomial has more than one irreducible factor is decomposable. Hence an indecomposable subspace has minimal polynomial a power of an irreducible polynomial. The splitting lemma in the previous chapter implies that any non cyclic T -invariant subspace is also decomposable. **QED.**

Corollary: Given an operator $T:V \rightarrow V$ with $\dim(V)$ finite, there is a product decomposition of V into indecomposable T -invariant subspaces on each of which T is cyclic with minimal polynomial a power of an irreducible polynomial.

Proof: Just keep decomposing the space until it is no longer decomposable; then the factors are of the form stated in the previous corollary.

More explicitly, first decompose the space into T -cyclic factors as in the previous chapter. Then use the previous relatively prime decomposition lemma to further decompose each cyclic factor according to the factorization of each restricted minimal polynomial into powers of irreducible factors. Then we claim each such factor is still T -cyclic. To see this, we argue as follows:

Given a T -cyclic factor of form $k[X].v$ in the invariant factor decomposition, let the restricted minimal polynomial equal $f.g$ where f and g are relatively prime.

Then the vector $f.v = f(T)(v)$ has minimal polynomial g and the vector $g.v = g(T)(v)$ has minimal polynomial f , so we can split the T -cyclic subspace $k[X].v$ further into the product of the T -cyclic subspaces $k[X].(g.v)$ and $k[X].(f.v)$. In this way we can split off one T -cyclic factor for each power of an irreducible polynomial occurring in the minimal polynomial of $k[X].v$, until we have split $k[X].v$ into a product of T -cyclic subspaces each of whose minimal polynomials is a power of a different irreducible factor of the minimal polynomial of $k[X].v$. Thus we eventually have a T -cyclic, indecomposable, prime power decomposition of V . **QED.**

Corollary: We can always find a matrix for T composed of blocks along the diagonal, with each block a companion matrix for a power of an irreducible polynomial. The product of all the corresponding polynomials is the characteristic polynomial of T . For each irreducible factor of the characteristic polynomial, we may consider the highest power of that factor which occurs as a companion matrix; the product of these highest occurring powers, is the minimal polynomial of T .

Remark: I.e. if the minimal polynomial for T in the previous corollary has form $f_1^{r_1}. f_2^{r_2}.....f_s^{r_s}$, then there will be at least one companion matrix for each power $f_i^{r_i}$, but there may also be other companion matrices associated to equal or lower powers of each f_i . Next we show how to obtain a matrix in terms of the companion matrices of just the irreducible factors f_i .

General Jordan form: Given a T cyclic subspace, i.e. one with a T -cyclic basis of form $\{v, T(v), T^2(v), \dots, T^{(n-1)}(v)\}$, the matrix associated to this basis is the companion matrix of the restricted minimal polynomial of T on this subspace. If that minimal polynomial is a power f^r of some polynomial f , we can tweak the basis slightly to obtain a matrix containing r copies of the companion matrix of f , plus some 1 's. This displays more precise information than the companion matrix of f^r . I.e. it is more informative to know the coefficients of f than the coefficients of f^r , just as it is more informative to represent a number as 3^6 than as 729 .

To see how to do this, let $f(X) = a_0 + a_1X + a_2X^2 + \dots + a_{n-1}X^{n-1} + X^n$, and assume the minimal polynomial of T on a given T -cyclic subspace is f^r , with T -cyclic vector v . This means the T -cyclic subspace has dimension $n.r$ with basis $\{v, Tv, \dots, T^{(nr-1)}v\}$. In this basis, the matrix of T will be the companion matrix of the polynomial f^r , all 1 's below the main diagonal, and (minus) the first $n.r$ coefficients of f^r in the last column. Instead of this basis, the Jordan form uses instead the sequence $\{v, Tv, \dots, T^{n-1}(v), f(T)(v), f(T)(T(v)), \dots, f(T)(T^{n-1}(v)), \dots, f^{(r-1)}(v), \dots, f^{(r-1)}(T^{n-1}(v))\}$.

Try this for $r=2$ to see what the matrix looks like. Each of the first $n-1$ basis vectors is taken by T to the next basis vector. Hence the first $n-1$ columns are $[0, 1, \dots, 0]$, $[0, 0, 1, \dots, 0]$, \dots , $[0, 0, \dots, 0, 1, 0, \dots, 0]$.

But the n th basis vector, $T^{n-1}(v)$, is taken by T to

$T^n(v) = -a_0v - a_1Tv - \dots - a_{n-1}T^{n-1}(v) + f(T)(v)$, whose coefficient vector is $[-a_0, -a_1, \dots, -a_{n-1}; 1, 0, \dots, 0]$. I.e. we have a copy of the companion matrix for f followed by a single '1' in the n th column. Then the next $n-1$ basis vectors are each again taken to the next basis vector, while the last one, $T^{n-1}(f(T)(v))$, is taken to $T^n(f(T)(v)) =$

$$\begin{aligned} & -a_0.f(T)(v) - a_1.f(T)(T(v)) - \dots - a_{n-1}.T^{n-1}(f(T)(v)) + f(T)(f(T)(v)) \\ & = -a_0.f(T)(v) - a_1.f(T)(T(v)) - \dots - a_{n-1}.T^{n-1}(f(T)(v)), \text{ since } f(T)(f(T)(v)) \\ & = f^2(T)(v) = 0, \text{ by hypothesis.} \end{aligned}$$

Thus the lower right hand block of the matrix is the companion matrix of f .

Thus the matrix of T in this basis is this:

$$\begin{array}{l} | 0 \ 0 \dots \dots \ 0 \ -a_0 \ | \ 0 \dots \dots \dots \ 0 | \\ | 1 \ 0 \dots \dots \ 0 \ -a_1 \ | \ 0 \dots \dots \dots \ 0 | \\ | 0 \ 1 \dots \dots \ 0 \ -a_2 \ | \ 0 \dots \dots \dots \ 0 | \\ | 0 \ 0 \dots \dots \dots \ | \dots \dots \dots \ | \\ | \dots \dots \dots \dots \dots \ | \dots \dots \dots \ | \\ | 0 \ 0 \dots \dots \ 1 \ -a_{n-1} \ | \ 0 \dots \dots \dots \ 0 | \\ | 0 \ 0 \dots \dots \ 0 \ 1 \ | \ 0 \ 0 \dots \dots \dots \ -a_0 | \\ | 0 \dots \dots \dots \ 0 \ | \ 1 \ 0 \dots \dots \dots \ -a_1 | \\ | 0 \dots \dots \dots \ 0 \ | \dots \dots \dots \ | \\ | \dots \dots \dots \dots \dots \ | \dots \dots \dots \ | \\ | 0 \dots \dots \dots \ 0 \ | \ 0 \dots \dots \ 1 \ -a_{n-1} \ | \end{array}$$

Thus the upper left and lower right blocks are copies of the companion matrix of f , the upper right block is all zeroes, while the lower left block has a single '1' in its upper right corner, and zeroes elsewhere. This is the elementary Jordan block associated to f^2 .

In general, for f^r , f of degree n , we get an $n.r$ by $n.r$ matrix with r copies of the companion matrix for f along the diagonal, but also with a single '1' below the last entry of each companion matrix except the lower right one. Thus if the n by n companion matrix of f is as follows:

$$\begin{vmatrix} 0 & 0 & \dots & 0 & -a_0 \\ 1 & 0 & \dots & 0 & -a_1 \\ 0 & 1 & \dots & 0 & -a_2 \end{vmatrix} = Cf$$

$$\begin{vmatrix} \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & \dots & 1 & -a_{n-1} \end{vmatrix}$$

and we define a special n by n matrix with a single '1' in it as follows:

$$\begin{vmatrix} 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{vmatrix} = E.$$

Then the elementary Jordan block associated to f^r is the following $r.n$ by $r.n$ matrix, containing r copies of Cf , and $r-1$ copies of E :

$$\begin{vmatrix} Cf & 0 & 0 & 0 & 0 & 0 \\ E & Cf & 0 & 0 & 0 & 0 \\ 0 & E & Cf & 0 & 0 & 0 \\ \dots & \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & E & Cf & 0 & 0 \\ 0 & 0 & 0 & E & Cf & 0 \\ 0 & 0 & 0 & 0 & E & Cf \end{vmatrix} = J(f),r$$

Remark: Thus every k -linear operator has a matrix representation consisting of blocks of this type where all polynomials f are irreducible over k , although finding a corresponding basis in practice assumes the ability to factor the minimal polynomial into irreducible factors. Moreover, the number and sizes of the blocks corresponding to each irreducible factor are uniquely determined by the operator. Indeed they are determined by the dimensions of the subspaces $\ker(f^s(T))$, for all irreducible factors f of the minimal polynomial and all powers f^s that divide that polynomial. Precisely, the dimension of the quotient $\ker(f^s(T))/\ker(f^{s-1}(T))$, as a vector space over the field $k[X]/(f)$, equals the number of blocks $J(f),r$ that occur for $r \geq s$.

It is usual to rearrange the full block matrix so that all occurrences of the same

irreducible factor are adjacent. If we choose an ordering for these irreducible factors, and if within the part of the matrix devoted to a single factor we order the blocks by size, then the matrix is uniquely determined.

Since the general Jordan form is a refinement of the rational canonical form, we can work backwards and recover the rational canonical form from the general Jordan form. E.g. take the largest Jordan block corresponding to each irreducible factor of the minimal polynomial of T ; the product of the minimal polynomials of these blocks is the largest invariant factor, i.e. the minimal polynomial of T . Then take the second largest Jordan block corresponding to each irreducible factor; the product of their minimal polynomials equals the second largest invariant factor of T ; etc... In this way the general Jordan form determines all the invariant factors of T . The uniqueness of the general Jordan form thus implies the uniqueness of the rational canonical form as well.

One can describe how to reconstruct the invariant factor decomposition from the Jordan one as follows. We know that if two T -cyclic subspaces have relatively prime restricted minimal polynomials, then their product is still T -cyclic. Indeed from our lemmas about minimal polynomials, we can take the sum of the two cyclic vectors of the subspaces as a cyclic vector for their product. Hence in the Jordan decomposition we choose for each irreducible factor f_j of the full minimal polynomial, a T -cyclic subspace corresponding to a power of f_j , and we choose these subspaces as large as possible. Then the product of these subspaces is a maximal T -cyclic subspace. Then we repeat this process, choosing from among the remaining T -cyclic subspaces, again choosing as large a one as possible for each irreducible factor f_j of the minimal polynomial. Continuing, we eventually have the T -cyclic invariant factor decomposition, in which the T -cyclic subspaces are as large as possible.

Examples: Recall the earlier example $A =$

$$\begin{vmatrix} 1 & 0 & 1 \\ 1 & 1 & 0 \\ 0 & 1 & 1 \end{vmatrix}$$

$$\begin{vmatrix} 1 & 1 & 0 \\ 0 & 1 & 1 \end{vmatrix}$$

$\begin{vmatrix} 0 & 1 & 1 \end{vmatrix}$. In this case we found minimal polynomial $X^3 - 3X^2 + 3X - 2$, hence the rational canonical form is just the companion matrix:

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

$$\begin{vmatrix} 1 & 0 & -3 \\ 0 & 1 & 3 \end{vmatrix}$$

$\begin{vmatrix} 0 & 1 & 3 \end{vmatrix}$. But we can factor this polynomial as $(X-2)(X^2-X+1)$, so the

Jordan form consists of two companion matrix blocks, one associated to each irreducible factor:

$$\begin{vmatrix} 2 & 0 & 0 \\ 0 & 0 & -1 \\ 0 & 1 & 1 \end{vmatrix}.$$

The upper left block is the companion matrix of $(X-2)$, i.e. just the one by one matrix $[2]$, and the lower right block is the 2 by 2 companion matrix of X^2-X+1 . In both cases the coefficients have their signs reversed, and begin with the constant coefficient.

The second example from above was $A =$

$$\begin{vmatrix} 1 & 1 & -1 \\ 1 & -1 & 1 \\ -1 & 1 & 1 \end{vmatrix},$$

with minimal polynomial $(X^3 - X^2 - 4X + 4)$, which factors

completely over \mathbb{Q} as $(X-1)(X-2)(X+2)$. Since the irreducible factors are all linear and occur only to the first power, the Jordan matrix is composed entirely of companion matrices of the linear factors, hence is “diagonal”, as follows:

$$\begin{vmatrix} -2 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 2 \end{vmatrix},$$

where we have chosen to order the characteristic roots by size.

Of course these computations were done over the rational field \mathbb{Q} . If we extend our field to the complex field, we could factor the first example also into distinct linear factors, $(X-2)(X^2-X+1) = (X-2)(X - (1+i\sqrt{3})/2)(X - (1-i\sqrt{3})/2)$. As noted there, this yields a diagonal Jordan matrix, with the roots of this polynomial, namely $\{2, (1+i\sqrt{3})/2, (1-i\sqrt{3})/2\}$, on the diagonal.

Similarly, the 3rd example mentioned above, namely $A =$

$$\begin{vmatrix} 1 & 0 & -1 \\ 0 & -1 & 1 \\ -1 & 1 & 0 \end{vmatrix},$$

with minimal polynomial $X(X^2-3)$, has the following Jordan form

$$\begin{vmatrix} 0 & 0 & 0 \\ 0 & 0 & 3 \\ 0 & 1 & 0 \end{vmatrix},$$

over the rational field, and the following diagonal Jordan form over \mathbb{R} :

$$\begin{vmatrix} 0 & & 0 \\ 0 & -\sqrt{3} & 0 \end{vmatrix}$$

$$\begin{vmatrix} 0 & 0 & \sqrt{3} \end{vmatrix}$$

Remark: Existence of the Jordan form is often asserted only over the complex field, or with the added hypothesis that the minimal polynomial does factor into linear factors (not necessarily distinct) over the given field. This special case, the “classical” Jordan form, is the most important one, so we discuss it separately in the next section. For now, observe that the theory of this section implies the Jordan form is diagonal, with the roots of the minimal polynomial along the diagonal, if and only if the minimal polynomial factors into *distinct* linear factors over k .

Eigenvectors, diagonalizable matrices, and classical Jordan form

The classical Jordan form is the special case of the general Jordan form when the minimal polynomial of the operator T factors into linear factors, not necessarily distinct, over the field k . This is always true over the complex numbers or any algebraically closed field, and can be forced in any case by passing from the given field to an extension of it which contains the roots of the minimal polynomial.

This provides us with a canonical form for our matrix which is as simple as possible, frequently diagonal, and always “almost diagonal”. In this case all irreducible factors of the minimal polynomial are of form $(X-c)$, so the minimal polynomial for a T -cyclic subspace is of form $(X-c)^r$, for $r \geq 1$. From our earlier theory, an elementary Jordan block, for $(X-c)^r$, looks like this r by r matrix:

$$\begin{vmatrix} c & 0 & 0 & 0 & 0 & 0 \\ 1 & c & 0 & 0 & 0 & 0 \\ 0 & 1 & c & 0 & 0 & 0 \\ \dots & \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & 1 & c & 0 & 0 \\ 0 & 0 & 0 & 1 & c & 0 \\ 0 & 0 & 0 & 0 & 1 & c \end{vmatrix} = J(c),r$$

Thus when the minimal polynomial of a matrix has all its roots in the field k , the Jordan form will be composed of blocks like the one above, of various sizes, and for various constants c . These constants c are called “generalized eigenvalues”. Notice that on a T -cyclic subspace where T has the elementary Jordan block matrix above, the operator $(T-c)^r$ is identically zero. Thus analyzing T in this setting involves studying the null spaces of operators of form $(T-c)^r$. We want to introduce the usual terminology for elements of such null spaces.

Defn: An **eigenvector** of a k - linear map T , is a non zero vector v such that $T(v)$ is a scalar multiple of v , i.e. such that $T(v) = cv$ for some scalar c in the field k .

The scalar c is the **eigenvalue** associated to v . (Other terms used in place of “eigen”, rough translations from German, are “proper” and “characteristic”.)

Rmk: Any non zero vector in the kernel of T is an eigenvector with eigenvalue zero. Since an eigenvector of T with eigenvalue c is a non zero vector in the kernel of $(T-cI)$, a non zero vector in the kernel of $(T-cI)^r$ is called a **generalized eigenvector** of T .

Defn: Given a linear map $T:V \rightarrow V$, and a scalar c , the eigenspace of T corresponding to c , is the kernel of $(T-c)$; the kernel of a power $(T-c)^r$ is called a generalized eigenspace of T .

Remark: Thus the eigenspaces and generalized eigenspaces consist of the corresponding eigenvectors, or generalized eigenvectors, plus the zero vector. Since this is a rather fine distinction we may sometimes make the careless error of referring to an eigenspace as a subspace of eigenvectors.

Eg. If c is constant, the functions $a.e^{ct}$ with $a \neq 0$, i.e. the non zero elements of the kernel of $(D-c)$, are eigenvectors of the operator D acting on smooth functions. The function $a.t.e^{ct}$ is in the kernel of $(D-c)^2$, and the functions $a.t^{(r-1)}.e^{ct}$ are elements of the kernel of $(D-c)^r$, hence generalized eigenvectors of D . In this setting these eigenvectors are often called eigenfunctions.

E.g. A 90 degree counter-clockwise rotation T of the real plane \mathbb{R}^2 about the origin, is a linear map with no eigenvectors. Since this operator satisfies $T^2 = -I$, the minimal polynomial is X^2+1 . The matrix is

$$\begin{vmatrix} 0 & -1 \\ 1 & 0 \end{vmatrix}.$$

If we consider the action of this operator on the complex “plane” \mathbb{C}^2 , defined by multiplying complex coordinate vectors by the same matrix, the extended operator will then have complex eigenvalues $\pm i$, and corresponding complex eigenvectors.

Geometry of eigenvectors: Recall that a vector v in real Euclidean space has both a length and (if $v \neq 0$) a direction. An eigenvector is a non zero vector v such that either $T(v) = 0$, or v and $T(v)$ have the same (or opposite) direction. Hence v spans a line that is mapped by T into itself.

As illustrated by the rotation operator above, the concept of eigenvalue is dependent on the given field.

Proposition: The eigenvalues of the k -linear map T are exactly those roots of the minimal polynomial of T which lie in the field k .

Proof: If c is a root in k of the minimal polynomial f , then $f(X)$ factors over k as $(X-c)g(X)$, for some g of lower degree. Since f is minimal there is some vector v that is annihilated by f but not by g . Then $f(T)(v) = 0 = (T-c)(g(T)(v))$. Since $g(T)(v) \neq 0$, $g(T)(v)$ is an eigenvector of T with eigenvalue c .

Conversely, if c is an eigenvalue of T in k , then there is some vector $v \neq 0$ with $(T-c)(v) = 0$, i.e. T satisfies $(X-c)$ at v . Since the polynomial $(X-c)$ is linear, it is the minimal polynomial of T at v . Since the minimal polynomial of T on all of V also vanishes at v , it must be a multiple of $(X-c)$, i.e. c is a root in k of that minimal polynomial. **QED.**

Corollary: If the minimal polynomial of T factors completely into linear factors over the field k , then the eigenvalues of T are exactly the roots of the minimal polynomial, hence also exactly the roots of the characteristic polynomial.

Proof: Since the characteristic and minimal polynomials have the same irreducible factors, they also have the same roots, so this follows from the previous proposition. **QED.**

The fundamental result, as before, is that we can reduce the study of an operator to subspaces on which the minimal polynomial is as simple as possible.

Decomposition Lemma: If a linear map $T:V \rightarrow V$ has minimal polynomial $(X-c_1)^{r_1}(X-c_2)^{r_2}\dots(X-c_t)^{r_t}$ with all c_i distinct, i.e. if all the roots c_i are in the field of scalars, then V is isomorphic to the product of the subspaces $V_i = \ker(T-c_i)^{r_i}$. The restriction of T to the subspace V_i has minimal polynomial $(X-c_i)^{r_i}$.

Proof: This is just the special case of the general decomposition lemma proved earlier, when the roots of the minimal polynomial all lie in the field of scalars. **QED.**

Corollary: Whenever the minimal polynomial of $T:V \rightarrow V$ has all its roots in the field of scalars, there is always a basis for V composed of generalized eigenvectors.

Proof: Just choose a basis for each subspace $\ker(T-c_i)^{r_i}$ and combine them to get a basis for V . **QED.**

The simplest case thus occurs when there is a basis of (ordinary) eigenvectors, the “diagonalizable” case.

Definition: An operator $T:V \rightarrow V$ is “diagonalizable” if and only if there is a basis

in which the matrix of T is diagonal, i.e. the matrix consists of all zeroes except possibly along the main diagonal.

Theorem: A linear map $T:V \rightarrow V$ is diagonalizable iff there is a basis for V composed of eigenvectors or “eigenbasis”, iff the minimal polynomial has form $(X-c_1)(X-c_2)\dots(X-c_t)$ where all c_i are distinct.

Proof: If V has dimension n , the rules for forming a matrix from a basis $\{v_1, \dots, v_n\}$, imply the matrix of $T:V \rightarrow V$ will be diagonal if and only if for every basis vector v_j , we have $T(v_j) = c_j \cdot v_j$, for some scalar c_j . This says exactly that a matrix is diagonal if and only if it is associated to a basis of eigenvectors, hence an operator is diagonalizable if and only if there exists a basis of eigenvectors.

If V has an eigenbasis and c_1, \dots, c_t , is a maximal sequence of distinct scalars among its sequence of eigenvalues, then the polynomial $(X-c_1)(X-c_2)\dots(X-c_t)$ annihilates every vector in the eigenbasis when T is substituted for X . I.e. the factors $(T-c_1)(T-c_2)\dots(T-c_t)$ commute with one another, so since $(T-c_i)$ annihilates all the basis eigenvectors associated to the eigenvalue c_i , this product annihilates the entire basis, hence the whole space. Thus the minimal polynomial must divide this polynomial, hence must also have distinct linear factors. (In fact no proper factor can annihilate the whole basis, since e.g. the map $(T-c_2)\dots(T-c_t)$ takes v_1 to $(c_1-c_2)\dots(c_1-c_t) \cdot v_1 \neq 0$. So $(X-c_1)(X-c_2)\dots(X-c_t)$ is the minimal polynomial.)

Conversely, if the minimal polynomial factors into distinct linear factors, then the decomposition lemma above implies that the space is a product of subspaces of form $\ker(T-c_i)$. Hence choosing a basis for each of these and combining them into a basis for V gives a basis of eigenvectors. **QED.**

Corollary: If the characteristic polynomial of T factors into distinct linear factors, then it equals the minimal polynomial, and T is diagonalizable.

Remark: The last part of the proof of the previous Theorem shows that if an operator T is diagonalizable on V , and if we can actually compute the roots of its minimal polynomial, then it is straight forward to compute a basis that diagonalizes the operator. Namely for each eigenvalue c , compute a basis for the kernel of $(T-c)$. Combining these bases, for all eigenvalues c , gives an eigenbasis of V .

Remark: Over the complex numbers, most polynomials factor into distinct linear factors, so statistically speaking, “most” complex matrices are diagonalizable, i.e. diagonalizable matrices form a dense open subset of all complex $n \times n$ matrices. For

example a quadratic polynomial X^2+bX+c has a repeated root if and only if $b^2 = 4c$, so in the (b,c) plane of all monic quadratic polynomials, those with a repeated root lie on the parabola $b^2 - 4c = 0$. Thus, the characteristic polynomial of a 2 by 2 matrix with rows $[a \ b]$, $[c \ d]$ has a repeated root if and only if $(a+d)^2 - 4(ad-bc) = 0$. Hence such matrices lie on a 3 dimensional quadratic “hypersurface” in the 4 dimensional space of all 2 by 2 matrices.

Remark: If $T:R^n \rightarrow R^n$ is given by a matrix B in the standard basis, and if $\{v_1, \dots, v_n\}$ is an eigenbasis, and if N is the matrix with the eigenvectors v_j as columns, then the matrix $N^{-1}.B.N$ will be diagonal. I.e. N takes the standard basis vectors e_j to the eigenvectors v_j , then B stretches those eigenvectors to $c_j.v_j$, so N^{-1} takes those stretched eigenvectors back to the stretched standard basis vectors $c_j.e_j$, and the resulting matrix, for the composition taking each e_j to $c_j.e_j$, has the c_j 's on the diagonal and zeroes elsewhere.

Eg: The map $T:R^2 \rightarrow R^2$ with $T(1,0) = (0,1)$, and $T(0,1) = (0,0)$, satisfies $X^2 = 0$, i.e. $T(T(v)) = 0$ for all v , but $T(1,0) \neq (0,0)$, so T has minimal polynomial X^2 . By the theorem, T is not diagonalizable.

Eg: Prove directly that in the previous example, all eigenvectors have form $(0,y)$.

Eg: The map $c:V \rightarrow V$ multiplication by the scalar c , has diagonal matrix $c.I$ in every basis. In particular, every basis is an eigenbasis for the identity map.

Eg. The map $T:R^2 \rightarrow R^2$ sending $(1,0)$ to $(1,0)$ and $(0,1)$ to $(0,2)$ has diagonal matrix with columns $(1,0)$ and $(0,2)$ in the standard basis, i.e. the standard basis is an eigenbasis. But $(1,1)$ is not an eigenvector, so in the basis $\{(1,0), (1,1)\}$ the matrix of f is not diagonal, but has columns $(1,0)$, $(-1,2)$.

Ex: The derivative map D acting on polynomials of degree ≤ 2 , has no eigenbasis, but any non zero constant polynomial is an eigenvector with eigenvalue zero.

Ex: Let V be the solution space of the differential equation $f'' - f = 0$, on the space of real-valued differentiable functions of a real variable. This is the kernel of D^2-1 , where D is the derivative operator, with minimal polynomial X^2-1 on V . It follows (from linear algebra and the mean value theorem) that V has dimension 2 over R , hence $\{e^t, e^{-t}\}$ is an eigenbasis for the operator D on V .

[Sketch of proof that $f'' - f = 0$ has 2 dimensional solution space: Let W be any

finite dimensional subspace of solutions containing e^t and e^{-t} . Since $D^2 - 1 = (D+1)(D-1)$ is identically zero on W , $(D+1)$ must map W into the kernel of $(D-1)$. Thus $W/\ker(D+1)$ injects into $\ker(D-1)$, so $\dim(W) \leq \dim(\ker(D-1)) + \dim(\ker(D+1))$. Thus it suffices to show $\dim(\ker(D \pm 1)) = 1$. But if f solves $f' - af = 0$ for any constant a , then the derivative of (f/e^{at}) equals $e^{at} \cdot (f' - af)/e^{2at} = 0$, hence by the mean value theorem $f/e^{at} = c$ is constant, so $f = c \cdot e^{at}$. Thus the kernel of $(D-a)$ is one dimensional with basis e^{at} .]

Ex: Let V be the solution space of the differential equation $f'' + f = 0$, acting on the space of real-valued differentiable functions of a real variable. Again V has dimension 2 over \mathbb{R} , with basis $\{\cos(t), \sin(t)\}$. The minimal polynomial of D on V is $X^2 + 1$, so there are no real eigenvalues or eigenvectors. (Note the matrix for D in this basis is the same as the matrix for the 90 degree cc rotation of the real plane). If we consider the solutions as a subspace of complex - valued functions of a real variable, then an eigenbasis exists over \mathbb{C} : $\{\cos(t) + i \cdot \sin(t), \cos(t) - i \cdot \sin(t)\} = \{e^{it}, e^{-it}\}$.

Ex: Let A be the following matrix over \mathbb{Q} :

$$\begin{vmatrix} 1 & 1 & 1 \\ 0 & 2 & 2 \\ 0 & 0 & 3 \end{vmatrix}$$
 Then the characteristic polynomial is the determinant of $[X-A] =$

$$\begin{vmatrix} (X-1) & -1 & -1 \\ 0 & (X-2) & -2 \\ 0 & 0 & (X-3) \end{vmatrix},$$

which equals $(X-1)(X-2)(X-3)$, with roots $\{1,2,3\}$, all in \mathbb{Q} and distinct. Hence A is diagonalizable with eigenvalues 1,2,3. We claim an eigenbasis is given by $\{(1,0,0), (1,1,0), (3,4,2)\}$.

Check this by finding bases for each of the kernels of the operators $(A-I)$, $(A-2.I)$, $(A-3.I)$.

Check also that if N is the matrix with these eigenvectors as columns, then $N^{-1} \cdot A \cdot N$ is diagonal.

Remark: We chose this matrix to be “triangular”, i.e. all zeroes below the main diagonal, so we could be sure the eigenvalues were exactly what we wanted them to be. This is the only general class of matrices I know whose characteristic polynomials are easy to factor into linear factors, and even to arrange to have any desired linear factors.

How to determine diagonalizability over \mathbb{Q} in practice

Notice that although it is not so practical to find all irreducible factors of a monic polynomial over \mathbb{Q} , it is always feasible to find all *linear* factors, since by the rational root theorem they must have form $(X-c)$ where c is an integer factor of the constant term. I.e. after factoring out any powers of X from our polynomial, we get a non-zero constant term with a finite number of integer factors, hence a finite list of other possible linear factors of our polynomial. Thus given an integer matrix, we can always decide whether or not it is diagonalizable over \mathbb{Q} , and if so, we can find an eigenbasis that diagonalizes it. Namely we can compute the characteristic polynomial either by diagonalizing the characteristic matrix, or by taking its determinant, and we can find all linear factors. If those suffice to factor it completely, then we can see whether they are all distinct or not. If the polynomial does factor completely into distinct linear factors, then we can diagonalize it over \mathbb{Q} by finding bases for the eigenspaces. If there are not enough integral linear factors to factor it completely, or if there exist repeated factors, then the matrix is not diagonalizable over \mathbb{Q} . As always you must take this algorithm with a few grains of salt, since if the constant term of the characteristic polynomial turns out to be say 5040, it might take you awhile to find the rational roots by hand.

Ex: Decide whether the following matrix is diagonalizable over \mathbb{Q} , and if it is, find a basis of eigenvectors. $A =$

$$\begin{vmatrix} 5 & -4 & 2 \\ 12 & -10 & 6 \\ 12 & -10 & 7 \end{vmatrix}.$$

Hint: Start by computing the characteristic polynomial of A , then use the rational root theorem to find its rational, i.e. integral, roots.

You should find that it is indeed diagonalizable over \mathbb{Q} . I made up this example by starting from a diagonal integral matrix D and an integral matrix N whose determinant I cooked up to equal one. Try it, it's not too hard if you keep the numbers small and throw in some zeroes. Then the inverse matrix N^{-1} is also integral, and can be found by row reduction as discussed earlier. Then I multiplied $N \cdot D \cdot N^{-1}$ to get this matrix A . So if you find a matrix N whose columns are a basis of eigenvectors for A , then $N^{-1} \cdot A \cdot N$ should be the diagonal matrix D . Please verify that your eigenvectors are correct by multiplying them by A . If they don't work, try recomputing the characteristic polynomial. (I got $X^3 - 2X^2 - X + 2$).

How to compute a Jordan basis

If the characteristic polynomial does factor completely into linear factors over a field k , even if there are repeated linear factors, then the matrix has a (classical) Jordan form over k . Given this factorization, we can compute a basis of generalized eigenvectors that puts the matrix in Jordan form. As stated above, the factorization exists theoretically if the field is algebraically closed, such as the complex numbers. In practice we usually use integer matrices, and then it is unusual for the characteristic polynomial to factor completely into linear factors over \mathbb{Q} . Still one can always decide whether or not this is true, by computing the characteristic polynomial, and then using the rational roots theorem to find all linear factors. If the characteristic polynomial does factor completely into linear factors over \mathbb{Q} , then as stated above one can compute the Jordan form as well as a basis of generalized eigenvectors that put the matrix into that form. Thus over \mathbb{Q} , although one cannot feasibly calculate the general Jordan form of an arbitrary matrix, due to the difficulty of finding non linear irreducible factors of the characteristic polynomial, one can determine whether a classical Jordan form exists for the matrix over \mathbb{Q} . If it does, one can calculate both the classical Jordan form and, with considerably more effort, a basis that achieves this form.

Notice that it is immediate from the relatively prime decomposition theorem that a basis of generalized eigenvectors exists whenever the characteristic polynomial factors into linear factors; what is more difficult is that such a basis can be chosen to be “cyclic” and hence give a Jordan matrix. This is a corollary of the general theory we gave for existence of rational canonical form, but we will deduce it again without assuming that theory. The benefit is that one can obtain the classical Jordan form without the abstraction of the theory of $k[X]$ spaces and their presentations, but the price is that the direct argument we will give is more complicated. Recall also that we did not explain how to choose the cyclic basis that gives the rational canonical form associated to the invariant factors of the characteristic polynomial, but referred to the algorithm in Dummitt and Foote. So another plus is that the more complicated argument that follows will show how to actually choose a cyclic basis that puts the matrix in Jordan form.

Jordan canonical form - “almost diagonalizable” maps

We have shown that an operator is diagonalizable if its minimal polynomial factors into distinct linear factors. Over an algebraically closed field like the complex numbers, the minimal polynomial of an operator will always factor into linear factors, but not necessarily distinct ones. If not, the map will not be diagonalizable, but it is almost diagonalizable, as we show next. The essential new wrinkle is the explicit consideration of “nilpotent” maps. Indeed, it will turn out that in this setting every operator is uniquely a sum of a diagonalizable part and a

nilpotent part.

Jordan canonical form

An operator is called “nilpotent”, if some power of it is identically zero. E.g. the differentiation operator D acting on the space of polynomials of degree $\leq n$, is nilpotent of order $n+1$, since the $n+1$ st derivative of a polynomial of degree $\leq n$ is zero. In the decomposition lemma above, transformations T whose minimal polynomials have repeated linear factors $(X-c)^r$, with $r \geq 2$, give a decomposition of the space into a product of subspaces on which the map $(T-c)$ is not identically zero but is nilpotent. I.e. $(T-c)^r$ is identically zero on $\ker(T-c)^r$ by definition of kernel. Since $T = c.Id + (T-c)$, on this subspace T is thus the sum of the diagonal operator $c.Id$ and the nilpotent operator $(T-c)$. This leads to what is called “Jordan form”, or a “Jordan matrix” for T .

The Jordan matrix of a nilpotent operator

We want to show how to choose a basis that puts a matrix in Jordan form. The simplest case is for a nilpotent n by n matrix with minimal polynomial X^n . This case is not too bad. I.e. if V is n dimensional and $T:V \rightarrow V$ satisfies the minimal polynomial $X^n = 0$, we want to find a cyclic vector v such that $\{v, T(v), T^2(v), \dots, T^{n-1}(v)\}$ is a basis. In this basis the matrix will be the following Jordan form, with 1's just below the main diagonal:

$$\begin{array}{c} | 0 \dots\dots\dots 0 | \\ | 1 0 \dots\dots\dots 0 | \\ | 0 1 \dots\dots\dots 0 | \\ \dots\dots\dots \\ | 0 \dots\dots\dots 1 0 | \end{array}$$

Hence we need a v that does not lie in the kernel of $T^{(n-1)}$, or equivalently we need v to be a basis for the quotient space $V/\ker(T^{(n-1)})$. In chapter one we observed that to find a basis for a quotient space we first find a basis for the subspace in the bottom, then extend that to a basis for V , and the extra vectors added to extend the basis are a basis for the quotient space.

Of course here $\ker T^{(n-1)}$ is $n-1$ dimensional and we just need to choose one vector not in this kernel. If we have computed the matrix for $T^{(n-1)}$ that is easy. I.e. by hypothesis this matrix is not zero and the j th column is the image of e_j under $T^{(n-1)}$. Hence if the j th column is non zero, then e_j is a cyclic vector. We need to prove if $T^{(n-1)}(v) \neq 0$, then $\{v, T(v), T^2(v), \dots, T^{n-1}(v)\}$ is a basis.

Since each vector is a basis of one of the successive quotients $\ker T^s / \ker T^{(s-1)}$, it follows from our earlier theory that their union is a basis for V .

To check it directly, assume we have a linear relation $a_0.v + a_1.T(v) + a_2.T^2(v) + \dots + a_{n-1}.T^{n-1}(v) = 0$. Apply $T^{(n-1)}$ to this equation getting $T^{(n-1)}(a_0.v) = a_0.T^{(n-1)}(v) = 0$, since $T^{(n-1)}$ annihilates all other vectors in the proposed basis. But since $T^{(n-1)}(v) \neq 0$, this implies that $a_0 = 0$. Since $a_0 = 0$, we have $a_1.T(v) + a_2.T^2(v) + \dots + a_{n-1}.T^{n-1}(v) = 0$, to which we can apply $T^{(n-2)}$ and conclude $a_1 = 0$. Continuing, all coefficients $a_j = 0$, so the sequence $\{v, T(v), T^2(v), \dots, T^{n-1}(v)\}$ is independent, hence a basis of our n dimensional space.

The case we just argued corresponds to a matrix with only one elementary Jordan block. Then next simplest case is where there is more than one block, but all blocks are the same size. This happens if the minimal polynomial of T on V is X^r , and the dimension of $\ker(T^r) / \ker(T^{(r-1)}) = s$, where $\dim(V) = r.s$. Then there are s elementary blocks, each of them of dimension r by r . In this case we choose a basis of the quotient space $\ker(T^r) / \ker(T^{(r-1)})$, say x_1, \dots, x_s , and then we claim the union of the cyclic sequences $\{\dots; x_j, T(x_j), \dots, T^{(r-1)}(x_j); \dots\}$ form a basis for V . I.e. $\{x_1, T(x_1), \dots, T^{(r-1)}(x_1); x_2, T(x_2), \dots, T^{(r-1)}(x_2); \dots; x_s, T(x_s), \dots, T^{(r-1)}(x_s)\}$ is a cyclic basis of V . The same arguments for independence work again. E.g. applying $T^{(r-1)}$ to a linear relation among these vectors gives a linear relation among the x 's, which are independent; continuing as before, the sequence is independent. By hypothesis $\dim(V) = r.s$, so these independent vectors form a basis, and are visibly cyclic.

The easiest way to choose a basis for $\ker(T^r) / \ker(T^{(r-1)}) \approx V / \ker(T^{(r-1)}) \approx \text{Im}(T^{(r-1)})$, may be to row reduce the matrix for $T^{(r-1)}$ to identify the pivot columns, and choose the standard basis vectors corresponding to the pivot columns of $T^{(r-1)}$ as basis. I.e. we want a sequence of s vectors such that only the trivial linear combination of them lies in $\ker(T^{(r-1)})$; equivalently, only the trivial linear combination of their images under $T^{(r-1)}$ is zero. This is given by standard basis vectors whose images under $T^{(r-1)}$ are independent, e.g. the standard basis vectors corresponding to pivot columns of $T^{(r-1)}$. So this is not too bad either; we only need to compute the matrix for $T^{(r-1)}$ and then row reduce it.

The next simplest case seems to be two blocks of different sizes, one smaller than the other. Here we are looking for two cyclic vectors v, w , with basis $\{w, T(w), \dots, T^{(s-1)}(w); v, T(v), \dots, T^{(r-1)}(v)\}$, where $r < s$. We need w to be a basis for the one dimensional space $\ker T^s / \ker T^{(s-1)} \approx V / \ker T^{(s-1)} \approx \text{Im} T^{(s-1)}$. We can proceed as above and choose w as a standard basis vector corresponding to a non zero column of the matrix for $T^{(s-1)}$.

Then we need to choose v so that $\{T^{(s-r)}(w), v\}$ is a basis of $\ker T^r / \ker T^{(r-1)}$.

This seems more complicated; we could proceed as follows. First choose a basis for $\ker T^{(r-1)}$, and add to this sequence the vector $T^{(s-r)}(w)$, and then add on also a basis for $\ker T^r$. Now reduce that sequence to an independent set, which will consist of first the basis for $T^{(r-1)}$, then the vector $T^{(s-r)}(w)$, and finally one more vector v . Then the pair $\{T^{(s-r)}(w), v\}$ is the basis we want of $\ker T^r / \ker T^{(r-1)}$, and hence $\{w, T(w), \dots, T^{(s-1)}(w); v, T(v), \dots, T^{(r-1)}(v)\}$ is our cyclic basis for V . At each stage we have bases for the successive quotient spaces, so their union is indeed a basis for V .

That was a little complicated but we can prove the following theorem by these arguments, without assuming our earlier results on existence of cyclic bases.

Theorem: If a k -linear operator $T: V \rightarrow V$ has minimal polynomial

$(X-c_1)^{r_1}(X-c_2)^{r_2}\dots(X-c_t)^{r_t}$ with all c_i distinct in k , then in some basis T has a matrix in Jordan form over k . For each factor $(X-c)^r$, there is a sequence of elementary Jordan blocks of size $\leq (r \text{ by } r)$, with c along the diagonal. The number of such blocks equals $\dim \ker(T-c)$. For each $s \leq r$, the number of such blocks which are at least s by s in size equals $\dim(\ker((T-c)^s) / \ker((T-c)^{(s-1)})$). At least one such block has size $(r \text{ by } r)$; the sum of the sizes of all blocks with c on the diagonal equals the largest power of $(X-c)$ which divides the characteristic polynomial of T .

Rmk: In particular, if T is diagonalizable, then the Jordan form is diagonal, and if T is not diagonalizable, the Jordan form is the closest thing to being diagonal. Moreover, the description at the end of the statement of the theorem shows that the number and size of elementary Jordan blocks is determined by the operator T , so the Jordan form is uniquely determined by T , up to ordering of the various blocks.

Proof: By the decomposition lemma V is the product of the subspaces $\ker(T-c_i)^{r_i}$, and since all polynomials in T commute with each other, each such subspace is T -invariant. So it suffices to show each subspace $W = \ker((T-c)^r)$ has a basis putting the matrix in Jordan form. Since on this subspace $(T-c)^r = 0$, but $(T-c)^{r-1} \neq 0$, the operator $N = (T-c)$ is nilpotent on W of order r , so $W = \ker(N^r)$. Consider the quotient space $W / \ker(N^{r-1})$, and choose a basis $[x_1], \dots, [x_n]$ for it. Since N induces an injection of $W / \ker(N^{r-1})$ to the quotient $\ker(N^{r-1}) / \ker(N^{r-2})$, we may extend the independent set $\{[N(x_1)], \dots, [N(x_n)]\}$ to a basis $\{[N(x_1)], \dots, [N(x_n)], [y_1], \dots, [y_m]\}$ for $\ker(N^{r-1}) / \ker(N^{r-2})$. Next we extend the independent set $\{[N^2(x_1)], \dots, [N^2(x_n)], [N(y_1)], \dots, [N(y_m)]\}$ to a basis for $\ker(N^{r-2}) / \ker(N^{r-3})$ Continuing, we obtain a basis

$\{[N^{r-1}(x_1)], \dots, [N^{r-1}(x_n)], [N^{r-2}(y_1)], \dots, [N^{r-2}(y_m)], \dots, [z_1], \dots, [z_q]\}$ for $\ker(N)$.

Then the union of the vectors representing these bases is a basis for V :

x_1, \dots, x_n ;

$N(x_1), \dots, N(x_n), y_1, \dots, y_m$;

$N^2(x_1), \dots, N^2(x_n), N(y_1), \dots, N(y_m), \dots$;

.....

$N^{r-1}(x_1), \dots, N^{r-1}(x_n), N^{r-2}(y_1), \dots, N^{r-2}(y_m), \dots, z_1, \dots, z_q$.

Re-ordering the vectors along the columns in the list above gives the cyclic basis we want, namely:

$\{x_1, N(x_1), \dots, N^{r-1}(x_1); x_2, N(x_2), \dots, N^{r-1}(x_2); \dots; x_n, N(x_n), \dots, N^{r-1}(x_n)$;

$y_1, \dots, N^{r-2}(y_1); y_2, \dots, N^{r-2}(y_2); y_m, \dots, N^{r-2}(y_m)$;

.....;

$z_1, \dots, z_q\}$.

Then there are $n = \dim(V/\ker(N^{r-1}))$, elementary Jordan blocks of size r by r corresponding to the x 's; there are m blocks of size $(r-1)$ by $(r-1)$, where $n+m = \dim(\ker(N^{r-1})/\ker(N^{r-2}))$, corresponding to the y 's;..... Finally, in the matrix for $N = (T-c)$, there are q blocks of size 1 by 1 , i.e. one q by q block of all zeroes, corresponding to the remaining eigenvectors $\{z_1, \dots, z_q\}$, where $n+m+\dots+q = \dim(\ker(N^2)/\ker(N)) + q = \dim(\ker(N))$.

Finally the matrix for $T = N+cI$, in this basis, is obtained from the matrix for N , by adding c 's everywhere along the diagonal. **QED.**

Rmk: Actually computing these Jordan bases is tedious, but one might proceed like this. For an integral matrix, compute the minimal polynomial by diagonalizing the characteristic matrix, and find its linear factors over Z by the rational root theorem. If they suffice to factor it completely you may continue to find the Jordan form. If $(X-c)^r$ is the factor of the minimal polynomial corresponding to the root c , set $N = (T-c)$ and compute bases of the subspaces $\ker(N), \ker(N^2), \dots, \ker(N^r)$. Write the basis for $\ker(N^r)$ after the basis for $\ker(N^{r-1})$, and then reduce to a new basis for $\ker(N^r)$ that contains the basis for $\ker(N^{r-1})$. The additional vectors give a basis for $\ker(N^r)/\ker(N^{r-1})$. These are the x 's in the proof above. Applying N to this basis injects it into $\ker(N^{r-1})/\ker(N^{r-2})$. Write this injected set after the basis for $\ker(N^{r-2})$, then add on the basis for

$\ker(N^{r-1})$, and reduce to a new basis for $\ker(N^{r-1})$ that contains, first the basis for $\ker(N^{r-2})$, and then a basis for $\ker(N^{r-1})$ that contains the injected basis from $\ker(N^r)/\ker(N^{r-1})$. Thus we have extended the image of that basis to a basis for $\ker(N^{r-1})/\ker(N^{r-2})$. The extra vectors added are the y 's from above. Continue....

For $r=3$, this looks as follows. Let $\{x_1, \dots, x_n\}$ be a basis of $\ker(N^3)$, let $\{y_1, \dots, y_m\}$ be a basis for $\ker(N^2)$, and let $\{z_1, \dots, z_p\}$ be a basis for $\ker(N)$. Then reduce the sequence $\{y_1, \dots, y_m, x_1, \dots, x_n\}$ to a basis for $\ker(N^3)$, eliminating m of the x 's, getting say $\{y_1, \dots, y_m, x_1, \dots, x_r\}$, for simplicity, where $r+m = n$. Thus $\{x_1, \dots, x_r\}$ is our basis for $\ker(N^3)/\ker(N^2)$.

Then reduce $\{z_1, \dots, z_p, N(x_1), \dots, N(x_r), y_1, \dots, y_m\}$ to a basis for $\ker(N^2)$, eliminating $p + r$ of the y 's, getting say $\{z_1, \dots, z_p, N(x_1), \dots, N(x_r), y_1, \dots, y_s\}$, for simplicity, where $p+r+s = m$. Thus $\{N(x_1), \dots, N(x_r), y_1, \dots, y_s\}$, is our basis for $\ker(N^2)/\ker(N)$.

Finally reduce $\{N^2(x_1), \dots, N^2(x_r), N(y_1), \dots, N(y_s), z_1, \dots, z_p\}$, to a basis for $\ker(N)$, eliminating $r+s$ of the z 's, getting say $\{N^2(x_1), \dots, N^2(x_r), N(y_1), \dots, N(y_s), z_1, \dots, z_t\}$, where $r+s+t = p$. This last sequence is our basis for $\ker(N)$.

Putting them all together gives our basis for $\ker(N^3)$:

$$\{x_1, \dots, x_r; N(x_1), \dots, N(x_r), y_1, \dots, y_s; N^2(x_1), \dots, N^2(x_r), N(y_1), \dots, N(y_s), z_1, \dots, z_t\}.$$

Reordering the basis gives the one we want for our Jordan form matrix:

$$\{x_1, N(x_1), N^2(x_1); \dots; x_r, N(x_r), N^2(x_r); y_1, N(y_1); \dots; y_s, N(y_s); z_1, \dots, z_t\}.$$

The Jordan matrix will thus consist of r elementary blocks of size 3×3 , s blocks of size 2×2 , and t blocks of size 1×1 . In particular there is a total of $r+s+t = p = \dim \ker(N)$ blocks of all sizes, since only the last column in each block corresponds to a basis vector for $\ker(N)$.

E.g. A Jordan matrix for an operator T acting on a space $V = \ker(T-c)^3$ of dimension 9. I.e. the minimal polynomial is $(X-c)^3$ and the characteristic polynomial is $(X-c)^9$.

$$\begin{matrix}
c & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
1 & c & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 1 & c & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & c & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 1 & c & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & c & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 1 & c & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & c & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & c
\end{matrix}$$

This matrix consists of 4 Jordan blocks, one of them 3x3 and three of them 2x2. As before, set $T-c = N$. Each of the 4 columns having only one 'c' in it and no '1', i.e. the right most column of each block, represents a basis vector for $\ker(N)$, which thus has dimension 4. The 4 columns just to the left of each of these columns, each having a 1 to the left of the previously mentioned c, represent a basis for $\ker(N^2)/\ker(N)$, which thus also has dimension 4. The first column, the only column having a 1 in it and also lying 2 columns to the left of a column with only a c, represents a basis for $\ker(N^3)/\ker(N^2)$, which thus has dimension 1.

Thus the fact there are 4 blocks means $\ker(N)$ has dimension 4. The fact there are 4 blocks at least 2x2 in size, means $\ker(N^2)/\ker(N)$ has dimension 4. The fact there is one 3x3 block means $\ker(N^3)/\ker(N^2)$ has dimension one. The minimal polynomial of N is X^3 , so the minimal polynomial of T is $(X-c)^3$.

Here the Jordan basis cannot all come from a basis of $\ker(N^3)/\ker(N^2)$. A basis vector v_1 of that space only represents the leftmost column of the matrix, and then $N(v_1)$ and $N^2(v_1)$ only give you the next 2 columns. To get the rest of the Jordan basis we need to extend $N(v_1)$ to a basis of $\ker(N^2)/\ker(N)$, getting three more vectors v_2, v_3, v_4 , representing columns 4, 6 and 8 of the matrix above. Then applying N to these gives the rest of the basis of $\ker(N)$, representing columns 5,7,9. The Jordan basis is $\{v_1, N(v_1), N^2(v_1), v_2, N(v_2), v_3, N(v_3), v_4, N(v_4)\}$.

E.g. If $\dim(V) = 3$, and the minimal polynomial of T is $(X-c)^3$, let $N = T-c$. Then $\ker(N)$ is one dimensional, $\ker(N)^2$ is two dimensional and $\ker(N)^3 = V$ is three dimensional. Just pick any basis $\{w_2, w_3\}$ of $\ker(N)^2$, and extend it to a basis $\{v_1, w_2, w_3\}$, adding in one new vector v_1 . Throw away $\{w_2, w_3\}$ and apply N twice

to v_1 ; then $\{v_1, N(v_1), N^2(v_1)\}$ is the Jordan basis. In this basis the 3×3 matrix for T has c 's along the main diagonal, and a '1' under each of the first two occurrences of ' c '. I.e. the three columns are the transposes of $[c, 1, 0]$, $[0, c, 1]$, $[0, 0, c]$, so the matrix forms a single 3×3 Jordan block.

If the minimal polynomial is $(X-c)$ then any basis is an eigenbasis and diagonalizes the matrix, since then $T = cI$. If $\dim(V) = 3$, the 3×3 matrix for T has c 's along the main diagonal and no 1's. The matrix is composed of three 1×1 Jordan blocks.

If $\dim(V) = 3$, and T has minimal polynomial $(X-c)^2$, again let $N = (T-c)$. Then $\ker(N)$ is two dimensional, so pick a basis $\{w_2, w_3\}$ of $\ker(N)$ and extend to a basis $\{w_2, w_3, v_1\}$ of V . Throwing away $\{w_2, w_3\}$ gives us a basis $\{v_1\}$ of $\ker(N)^2/\ker(N)$. Then $N(v_1) \neq 0$ belongs to $\ker(N)$, so we can reduce the sequence $\{N(v_1), w_2, w_3\}$ to a basis $\{N(v_1), w\}$ of $\ker(N)$, where w is either w_2 or w_3 . Then the Jordan basis for V is $\{v_1, N(v_1), w\}$. In this basis the 3 by 3 matrix for T has c 's along the diagonal, and a single "1" occurring just below the first c . I.e. the columns of the matrix are (the transposes of) $[c, 1, 0]$, $[0, c, 0]$, $[0, 0, c]$. Thus the matrix has one 2×2 Jordan block followed by a single 1×1 Jordan block.

Eg: The Jordan form of a linear map $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ with minimal polynomial $(X-c)^2$ has columns $(c, 1)$, $(0, c)$. The derivative $D: V \rightarrow V$ acting on polynomials of degree ≤ 2 , has X^3 as minimal polynomial, Jordan basis $\{X^2, 2X, 2\}$ and Jordan matrix with columns (the transposes of) $[0, 1, 0]$, $[0, 0, 1]$, $[0, 0, 0]$.

Ex. What is the (only possible) Jordan form of a 4×4 matrix with minimal polynomial $(X-c)^3$? What are all possible Jordan forms of a 5×5 matrix with minimal polynomial $(X-c)^3$, or $X^2(X-2)^3$?

Rmk: The map S whose matrix in the Jordan basis is just the diagonal entries is of course diagonalizable, and the map N whose matrix is just the off diagonal entries is nilpotent. Moreover these matrices, hence these maps, commute. Thus the original map $T = S+N$, is a sum of a diagonalizable map S and a nilpotent map N , such that S and N commute. This sums up the theoretical content of this theorem, i.e. a linear operator T whose minimal polynomial factors completely into linear factors, can be written (uniquely) as the sum of a diagonalizable ("semi simple") operator S and a nilpotent operator N which commute; i.e. $T = S+N$, and $SN = NS$.

Rmk: The uniqueness of the Jordan decomposition can be sketched as follows. If

$T = S + N$ where S is diagonalizable, N is nilpotent, and $SN = NS$, then S and N also commute with T and hence also with $T - cI$ for any constant c . Thus if c is an eigenvalue of T , both N and S map the kernel of $(T - cI)^k$ to itself. But $(T - cI)$ and N are both nilpotent on this subspace, and they commute, so their difference $(T - N - cI) = (S - cI)$ is also nilpotent on this subspace. Arguing backwards, $T - cI$ is nilpotent on any subspace on which $S - cI$ is nilpotent. I.e. the same subspace, corresponding to the eigenvalue c , occurs in the decomposition lemma for both T and S .

Thus the Jordan form for S on this subspace is a block in the general Jordan form for S . Since S is diagonalizable, that Jordan form on this subspace is diagonal with the constant c along the diagonal. I.e. on this subspace, the operator S is the diagonalizable summand we found above for T . Since this holds on every subspace in the decomposition, this is true on all of V , and S is thus unique.

In summary, the decomposition lemma yields a unique decomposition of V into subspaces of form $\ker(T - cI)^k$. Then for any decomposition of T as $S + N$ with S diagonalizable, N nilpotent and $SN = NS$, the restriction of S to each subspace $\ker(T - cI)^k$ must equal cI . Thus S is uniquely determined on each subspace $\ker(T - cI)^k$, and hence also on their sum, i.e. on the whole space. Since $N = T - S$, N is also unique.

Note: We have seen that a linear map between two different spaces is determined up to isomorphisms of source and target just by its rank, or equivalently the dimension of its kernel. Although the situation is more complicated for a map from a space V to itself, the theory of Jordan form shows that if the minimal polynomial is a product of powers of linear factors like $(X - c)^r$, then the map T is determined up to an isomorphism of V just by the dimensions of the kernels of the powers $(T - cI)^j$, for $1 \leq j \leq r$, and for all roots c of the minimal polynomial.

Example: Consider the matrix $A =$

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

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

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

$$\begin{vmatrix} 1 & -2 & -1 & 0 \\ 1 & 0 & -3 & 0 \\ -1 & -2 & 1 & 0 \\ 1 & 2 & 1 & 2 \end{vmatrix} \text{ The characteristic matrix is thus } X.I - A =$$

$$\begin{vmatrix} (X-1) & 2 & 1 & 0 \\ -1 & X & 3 & 0 \\ 1 & 2 & (X-1) & 0 \\ -1 & -2 & -1 & (X-2) \end{vmatrix}$$

$$\begin{vmatrix} -1 & X & 3 & 0 \\ 1 & 2 & (X-1) & 0 \\ -1 & -2 & -1 & (X-2) \end{vmatrix}$$

$$\begin{vmatrix} 1 & 2 & (X-1) & 0 \\ -1 & -2 & -1 & (X-2) \end{vmatrix}$$

I actually diagonalized this little guy and got the two invariant

factors $(X-2)$, $(X-2)^2(X+2)$. Hence the minimal polynomial is $(X-2)^2(X+2)$, and the characteristic polynomial is $(X-2)^3(X+2)$. This already tells us exactly what the Jordan form J of this matrix is: there will be two blocks corresponding to characteristic value 2, one of size 2 and one of size 1, and there will be a single block of size 1 corresponding to characteristic value -2. Namely $J =$

$$\begin{array}{l} |-2 \ 0 \ 0 \ 0| \\ |0 \ 2 \ 0 \ 0| \\ |0 \ 1 \ 2 \ 0| \end{array}$$

$|0 \ 0 \ 0 \ 2|$. It is more work to find a basis that puts it in this form, but it is doable if tedious arithmetic.

Since there are two powers of $(X-2)$ in the invariant factor sequence, we know that $\ker.(A-2)$ will be two dimensional, and since the number of exponents ≤ 2 add to 3, we know $\ker.(A-2)^2$ will be three dimensional. Also $\ker.(A+2)$ will be one dimensional.

I got this basis for $\ker.(A-2)$: $\{(0,0,0,1), (1,-1,1,0)\}$, and extended it to this basis of $\ker.(A-2)^2$: $\{(0,0,0,1), (1,-1,1,0), (0,-1,1,0)\}$. Thus the last vector $(0,-1,1,0)$ is a basis for the quotient space $\ker.(A-2)^2/\ker.(A-2)$. Since we want an $(A-2)$ - cyclic basis, we want to use $(A-2)(0,-1,1,0) = (1,-1,1,-1)$ as part of our basis for $\ker.(A-2)$. Thus we reduce the sequence $\{(1,-1,1,-1), (0,0,0,1), (1,-1,1,0)\}$ to a new basis for $\ker.(A-2)$, getting $\{(1,-1,1,-1), (0,0,0,1)\}$. Thus our new basis of $\ker.(A-2)^2$ is $\{(0,-1,1,0), (1,-1,1,-1), (0,0,0,1)\}$, which was denoted $\{x, (A-2)(x), y\}$ in our abstract notation. Then we compute as basis for $\ker.(A+2)$: $\{(-1,-1,-1,1)\} = \{z\}$. In the resulting Jordan basis $\{(-1,-1,-1,1); (0,-1,1,0), (1,-1,1,-1), (0,0,0,1)\}$, the matrix of A should be the Jordan matrix J above. I.e. if Q is the matrix whose columns are the vectors of this Jordan basis, we should have $Q^{-1}AQ = J$. To check this without computing Q^{-1} we could check equivalently that $AQ = QJ$, and I checked it that way. But I made a lot of typos in copying it here, so please verify everything. [I am also wondering why I ordered my Jordan blocks from largest to smallest, since the invariant factors go in the opposite order.)

Example: Consider the matrix $B =$

$$\begin{array}{l} |5 \ -1 \ -3 \ 2 \ -5| \\ |0 \ 2 \ 0 \ 0 \ 0| \\ |1 \ 0 \ 1 \ 1 \ -2| \\ |0 \ -1 \ 0 \ 3 \ 1| \end{array}$$

$|1 \ -1 \ -1 \ 1 \ 1|$, and diagonalize the characteristic matrix, getting the invariant factor sequence: $(X-2); (X-2)^2(X-3)^2$. (For me, this diagonalization was the

hard part.) Thus there are two Jordan blocks for the characteristic value 2, one of size 2 by 2 and one of size 1 by 1, and one 2 by 2 block for the characteristic value 3. I.e. the Jordan matrix J is this:

$$\begin{array}{c|cccc} 2 & 0 & 0 & 0 & 0 \\ 1 & 2 & 0 & 0 & 0 \\ \hline 0 & 0 & 2 & 0 & 0 \\ 0 & 0 & 0 & 3 & 0 \\ \hline 0 & 0 & 0 & 1 & 3 \end{array}$$

Again it takes more arithmetic to find the Jordan basis.

By row reduction I got this basis for $\ker(B-2)$: $\{(2,1,0,0,1), (1,0,1,0,0)\}$. Then I got this additional vector in a basis for $\ker(B-2)^2$: $(1,1,1,1,0)$, and hence as a basis for $\ker(B-2)^2/\ker(B-2)$. Fortunately, $(B-2)$ of this last vector equals $(1,0,1,0,0)$, so our new basis of $\ker(B-2)^2$ is just: $\{(1,1,1,1,0), (1,0,1,0,0), (2,1,0,0,1)\}$, which has the form $\{x, (B-2)(x), y\}$.

Computing gives this basis for $\ker(B-3)$: $\{(-1,0,0,1,0)\}$, and a basis for $(B-3)^2$ is: $\{(2,0,0,0,1), (-1,0,0,1,0)\}$, so $(2,0,0,0,1)$ is a basis for $\ker(B-3)^2/\ker(B-3)$. Again fortunately, $(B-3)(2,0,0,0,1) = (-1,0,0,1,0)$, so our Jordan basis of $\ker(B-3)^2$ is $\{(2,0,0,0,1), (-1,0,0,1,0)\}$. Ok, so our combined Jordan basis for the whole space is $\{(1,1,1,1,0), (1,0,1,0,0), (2,1,0,0,1), (2,0,0,0,1), (-1,0,0,1,0)\}$. The matrix with these as columns is the change of basis matrix Q. Thus we should have $Q^{-1}BQ = J$, which I verified, lazily as before, by checking that $BQ = QJ$.

Example: Just for fun, and as easy calculation, let's compare the difference between the companion matrix and the Jordan matrix for the same operator. Consider the polynomial $(X-2)^4 = x^4 - 8X^3 + 24X^2 - 32X + 16$, and the associated companion matrix A =

$$\begin{array}{c|cccc} 0 & 0 & 0 & -16 \\ 1 & 0 & 0 & 32 \\ \hline 0 & 1 & 0 & -24 \\ \hline 0 & 0 & 1 & 8 \end{array}$$

Then $(X-2)^4$ is the characteristic and minimal polynomial for A. We can easily confirm this by computation as follows. From the columns of A we see that the standard basis is A-cyclic; i.e. $Ae_1 = e_2$, $Ae_2 = e_3 = A^2e_1$, $Ae_3 = e_4 = A^3e_1$, while $Ae_4 = -16e_1 + 32e_2 - 24e_3 + 8e_4 = A^4e_1$. Expressing everything in terms of e_1 , gives $(-16 + 32A - 24A^2 + 8A^3)(e_1) = A^4e_1$, so $(A^4 - 8A^3 + 24A^2 - 32A + 16)(e_1) = 0$. Since A commutes with every polynomial in A, this polynomial in A also annihilates $A(e_1) = e_2$, $A^2(e_1) = e_3$, and $A^3(e_1) = e_4$. Hence this polynomial in A does annihilate every vector in a basis, hence equals zero. Since

$e_1, A(e_1) = e_2, A^2(e_1) = e_3, \text{ and } A^3(e_1) = e_4,$ are independent, no polynomial of degree less than 4 can annihilate e_1 , so the previous polynomial is indeed the minimal polynomial of A , both at e_1 and on the whole space.

Now to get a Jordan basis, since $(A-2)$ satisfies the polynomial t^4 , hence is nilpotent of order 4, we want an $(A-2)$ - cyclic basis. This is easy to arrange from what we know about A . Namely take $u_1 = e_1$, and then we want $u_2 = (A-2)(u_1) = A(u_1) - 2u_1 = A(e_1) - 2e_1 = e_2 - 2e_1$. Then we want $u_3 = (A-2)(u_2) = A(e_2) - 2e_2 - 2A(e_1) + 4e_1$. Finally we want $u_4 = (A-2)(u_3) = e_4 - 6e_3 + 12e_2 - 8e_1$.

Then we can check that $(A-2)(u_4) = A(e_4) - 8e_4 + 24e_3 - 32e_2 + 16e_1 = 0$. So in the basis $\{u_1, u_2, u_3, u_4\}$, $(A-2)$ has this matrix:

$$\begin{vmatrix} 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{vmatrix}$$

so in that same basis, the operator A has this Jordan matrix:

$$\begin{vmatrix} 2 & 0 & 0 & 0 \\ 1 & 2 & 0 & 0 \\ 0 & 1 & 2 & 0 \\ 0 & 0 & 1 & 2 \end{vmatrix}$$

As you can see, the difference between these matrices is that the first one, the companion matrix, displayed the coefficients of the minimal polynomial, while the second, the Jordan matrix, displayed its roots. You will also recall that the coefficients of the minimal polynomial are always computable by diagonalization, while the roots require factorization also. Hence the Jordan form, although simpler, may be impractical to compute for more difficult examples.

For more details and worked examples see my web notes, pp.11-16, at:

<http://alpha.math.uga.edu/%7Eroy/845-2.pdf>

Eg: If $f:V \rightarrow V$ is a linear map on a real vector space, with minimal polynomial X^2+1 , V decomposes as a sum of subspaces isomorphic to $\mathbb{R}[X]/(X^2+1)$, i.e. of 2 dimensional real subspaces, each of them one - dimensional over the field $\mathbb{R}[X]/(X^2+1)$. This quotient field is isomorphic to the complex numbers \mathbb{C} , where X corresponds to $i = \sqrt{-1}$, so the operator f on this space corresponds to multiplication by i . If V has dimension s over \mathbb{C} , the real rational canonical matrix of f on V , consists of exactly s blocks, each 2 by 2, with columns $(0,1), (-1,0)$.

Ex: A scalar c is an eigenvalue of f if and only if $\det(f-c.I) = 0$, iff c is a root of the characteristic polynomial of f , if and only if c is a root of the minimal polynomial of f .

Ex: Find all Jordan and rational canonical forms of linear maps of \mathbb{R}^3 with minimal polynomials $(X-2)^3$, $(X-2)^2$, and $(X-2)$.

Chapter Five: Spectral theorems, detecting “orthogonal” diagonalizability

In this chapter we consider vector spaces over the fields of real and complex numbers. With a view to applications to mechanics and rigid motions, we will introduce notions of length and angle for vector spaces over the real numbers, and learn which maps are diagonalizable for some choice of mutually perpendicular axes. These criteria are called “spectral theorems” and can be expressed in terms of symmetries. We also apply our theory to solving linear differential equations.

Recall a diagonalizable operator over \mathbb{Q} , whose eigenvalues are thus rational numbers, can be detected using the rational root theorem. But a diagonalizable operator over \mathbb{R} whose eigenvalues are irrational can be hard to recognize, since we often cannot factor the minimal polynomial. Thus it is useful to have a sufficient criterion for diagonalizability that does not require finding the irreducible factors of the minimal polynomial. Such criteria exist and we discuss these next.

Notice that a diagonalizable operator T has the property that every subset of eigenvectors spans a T -invariant subspace. Thus the first step in an inductive proof of diagonalizability of an operator T on V , is to find one eigenvalue, and hence one eigenvector. The next step would be to find a complementary T invariant subspace on which one can then find another eigenvalue and eigenvector. This would allow the proof to proceed by induction on dimension. Now if we work over say the complex numbers then there is no problem of existence of an eigenvalue on any invariant subspace, since the minimal polynomial always has a complex root and hence there is at least one eigenvector. So if we have a complex linear T , such that for the span of each eigenvector we can always find a complementary T -invariant subspace then we would be done. Over the reals we have to work a little harder, but there is a similar theorem.

One idea for constructing an operator T that has this nice splitting process, that every invariant subspace has an invariant complement, is to introduce a notion of

angles and perpendicularity, and then choose an operator T that preserves angles, or at least preserves perpendicularity. Then since T preserves the line spanned by an eigenvector, it would also preserve the subspace perpendicular to that line, and this would give us a complementary T -invariant subspace.

The simplest operator that preserves perpendicularity would be a “rigid motion”, or length preserving operator. This property can be expressed in terms of dot products. Recall this notion.

Definition: The dot product $\langle v, w \rangle = v \cdot w$ of two vectors $v = (a_1, \dots, a_n)$, $w = (b_1, \dots, b_n)$ in \mathbb{R}^n is defined as $a_1 b_1 + \dots + a_n b_n$.

By the n dimensional Pythagorean theorem, it follows that the squared length of a vector equals its dot product with itself. We make this a definition.

Definition: If $v = (a_1, \dots, a_n)$ is a vector in \mathbb{R}^n , then $v \cdot v = a_1^2 + \dots + a_n^2 \geq 0$, so we can define its “length” to be $\sqrt{v \cdot v} = |v| \geq 0$.

It is easy to check that the dot product is a “product” in the sense that it is distributive over addition and also commutes with scalar multiplication in each variable separately. That is,

Ex. Given vectors u, v, w in \mathbb{R}^n , and a scalar c , we have $u \cdot (v+w) = u \cdot v + u \cdot w$, and $u \cdot (cv) = (cu) \cdot v = c(u \cdot v)$.

This lets us make a useful calculation. Recall that if two vectors v, w starting from the origin are completed into a triangle by joining their “heads”, the vector joining the heads is parallel to and has the same length as the vector $w-v$. Then we can relate the length of this third side of the triangle to the lengths of the first two sides by expanding a dot product. I.e. then $|w-v|^2 = (w-v) \cdot (w-v) = w \cdot w - 2v \cdot w + v \cdot v = |v|^2 + |w|^2 - 2v \cdot w$. Now by Pythagoras’ theorem $|v|^2 + |w|^2 = |w-v|^2$ if and only if v and w are perpendicular, and by our calculation, this is equivalent to $v \cdot w = 0$. Thus we make another definition:

Definition: Two vectors v, w in \mathbb{R}^n are called “perpendicular”, also “orthogonal”, if and only if $v \cdot w = 0$.

Definition: If U is a subspace of \mathbb{R}^n , define $U_{\text{perp}} = \{\text{all vectors } w \text{ in } \mathbb{R}^n \text{ such that } v \cdot w = 0 \text{ for all } v \text{ in } U\}$.

Note: We defined length in terms of dot product, but the previous calculation also shows we could define dot product in terms of length. I.e. we computed that $(v \cdot w) = (1/2)(|v|^2 + |w|^2 - |v-w|^2)$.

Remark: Back in chapter 2, we defined a dot product on k^n for any field k , and noted that this gave an isomorphism from k^n to k^{n*} , sending each vector v to “dotting with v ”, i.e. to the linear function $v(\cdot)$, or sending each column vector to its transpose, the corresponding row vector. Under this isomorphism, for each subspace U , the orthogonal complement U^{\perp} in $(k^n)^*$, of functions in k^n that vanish identically on U , corresponds to the subspace U^{\perp} of k^n of vectors that dot to zero with every vector in U . But this copy of U^{\perp} inside of k^n may not be perpendicular to U in the intuitive sense, i.e. it may overlap U in a non zero subspace, or equivalently some non zero vectors in k^n may be perpendicular to themselves! E.g. if we consider the complex field $k = \mathbb{C}$, then the vector $(1, i)$ in \mathbb{C}^2 has dot product $1^2 + i^2 = 1 + (-1) = 0$, hence is perpendicular to itself in this sense. Over the reals however, since \mathbb{R} is ordered, we have $v \cdot v > 0$ if $v \neq 0$. This lets us use orthogonal complements to decompose \mathbb{R}^n .

Ex. If U is a subspace of \mathbb{R}^n , and U^{\perp} is its orthogonal complement in \mathbb{R}^n with respect to the dot product, show that the map $U \times U^{\perp} \rightarrow \mathbb{R}^n$, taking (v, w) to $v+w$, is injective, and hence an isomorphism.

We often want to “normalize” our vectors by taking them to be length one, which we can do just by dividing a non zero vector by its length. This gives another term,

Definition: A set of vectors is called “orthonormal” if they are all of length one and mutually perpendicular.

Ex. A set of orthonormal vectors in \mathbb{R}^n must be independent.

Now we can define a length preserving operator in terms of dot products as well.

Definition: An operator $T: \mathbb{R}^n \rightarrow \mathbb{R}^n$ is called length preserving, or an isometry, if and only if for every vector v , $|v| = |Tv|$.

Ex. Show that T is an isometry if and only if $v \cdot w = (Tv) \cdot (Tw)$ for all vectors v, w . (Hint: Use our observation above that length determines dot product.) In particular an isometry preserves perpendicularity.

There is another equivalent description of an isometry in terms of transposes.

Lemma: $T: \mathbb{R}^n \rightarrow \mathbb{R}^n$ is an isometry if and only if its matrix is the inverse of its transpose.

proof: If $TT^* = \text{Id}$, then for all v, w , $(Tv) \cdot (Tw) = v \cdot (T^*Tw) = v \cdot w$, recalling that we can move a matrix across a dot product by changing it into its transpose. In the other direction, If $v \cdot w = (Tv) \cdot (Tw) = v \cdot (T^*Tw)$ for all v, w , then $v \cdot (w - T^*Tw) = 0$ for all v, w . But if $u = w - T^*Tw$, then $u \cdot u = |u|^2 = 0$ implies $u = 0$. Thus $w - T^*Tw = 0$ for all w , i.e. $T^*Tw = w$ for all w , so $T^*T = \text{Id}$ and $T^* = T^{-1}$. **QED.**

Definition: A real matrix A on \mathbb{R}^n , is called orthogonal if $A^* = A^{-1}$, i.e. orthogonal matrices are exactly those that define isometries.

Cor: A matrix defines an isometry if and only if its columns are an orthonormal basis for \mathbb{R}^n , if and only if its rows are as well.

Now we know that an isometry preserves perpendicularity, hence sends the subspace orthogonal to any eigenvector into itself. This does not suffice to let us argue diagonalizability however, since an isometry may not have eigenvectors. For example, a rotation of the plane about the origin does not have any eigenvector. Similarly a rotation of three space about an axis has one eigenvector, namely a vector spanning the axis, but usually has no eigenvectors in the plane perpendicular to the axis. So a rotation of \mathbb{R}^3 has one eigenvector but usually no eigenbasis. Thus it is quite interesting and beautiful that a slight tweak of the property of isometry does give us a class of diagonalizable operators, namely if we swap the property $T^* = T^{-1}$ for the even simpler property $T = T^*$, we hit pay dirt.

Definition: A matrix A is **symmetric** if and only if $A = A^*$, if and only if, for every i and j , the entry in the i^{th} column and j^{th} row equals the entry in the j^{th} column and i^{th} row.

Happily, you can check that symmetric matrices still preserve orthogonality, even though they usually do not preserve length. This provides the inductive step in the argument for diagonalizability.

Ex. If A is a symmetric matrix and v is an eigenvector for A , and w is a vector perpendicular to v , then Aw is also perpendicular to v .

The harder step of starting the induction by finding an eigenvector, is argued in the proof of the next theorem.

Thm: If A is a symmetric n by n real matrix, then \mathbb{R}^n has a basis of mutually orthogonal unit length eigenvectors for A , hence an “orthonormal eigenbasis”.

Pf: (thanks to Ed Azoff) The continuous real valued function $f(x) = Ax \cdot x$ has a minimum c on the unit sphere in \mathbb{R}^n , at some point y . Thus $Ax \cdot x \geq Ay \cdot y = c$ for all x in the unit sphere. Since $x \cdot x = 1$ for all such x , thus $Ax \cdot x \geq cx \cdot x$ for all such x , hence $Ax \cdot x - cx \cdot x = (A - cI)x \cdot x \geq 0 = (A - cI)y \cdot y$, for all such x .

In particular, if $B = A - cI$, then for all real t , and all x in the unit sphere,
 $0 \leq B(y + tx) \cdot (y + tx)$.

Since $By \cdot y = 0$, expanding the right hand side, and using the symmetry of A , hence also of B , gives $0 \leq 2t (By \cdot x) + t^2 (Bx \cdot x)$. Hence the discriminant $(By \cdot x)^2$ of this quadratic is ≤ 0 . Since it is a square it is also ≥ 0 , and thus $By \cdot x = 0$ for every x in the unit sphere. Thus $By = 0 = (A - c)y$, so $Ay = cy$, i.e. y is a (unit length) eigenvector for A .

Now restrict A to the subspace V of vectors orthogonal to y . If $v \cdot y = 0$, then $Av \cdot y = v \cdot Ay = v \cdot cy = c(v \cdot y) = 0$. Hence A preserves V . A still has the property $Av \cdot x = v \cdot Ax$ on V , so the restriction of A to V has an eigenvector in V . (Although V has no natural representation as \mathbb{R}^{n-1} , the argument for producing an eigenvector depended only the symmetry property $Av \cdot x = v \cdot Ax$.) Repeating the argument, A has an eigenbasis of mutually orthogonal unit vectors. **QED.**

Remark: There is a slightly more sophisticated version of this argument for existence of an eigenvalue that may seem simpler to some people. Namely consider again the function $f(x) = Ax \cdot x$ on the unit sphere, and compute that its gradient vector, or derivative, is $2Ax$, just as in the one variable case. Then at an extremum y , this derivative must be zero on the tangent space to the sphere. Thus the gradient vector $2Ay$ is perpendicular to the tangent space of the sphere, hence either zero or parallel to the radius vector y . In both cases we have an eigenvector. Notice this finds an eigenvector with the largest eigenvalue at a maximum of f , and an eigenvector with smallest eigenvalue at a minimum of f . Note too that Dr. Azoff's argument in the proof above also works at a maximum point y .

Cor: If A is a symmetric $n \times n$ matrix, there is an orthogonal $n \times n$ matrix E such that E^*AE is a diagonal matrix.

Pf: If E has columns equal to the eigenbasis for A , the matrix associated to the map A by that eigenbasis was essentially defined as $E^{-1}AE$. But since we chose the eigenbasis vectors of unit length, the fact that they are also mutually perpendicular implies $E^{-1} = E^*$. **QED.**

Ex. Prove the converse of the previous corollary, i.e. that if $A = EDE^*$ where D is diagonal and E is orthogonal, then A is symmetric.

Understanding isometries of \mathbb{R}^n

We have completely described real symmetric operators, i.e. those which equal their transpose. Next we will describe operators which are inverse to their transpose, i.e. “orthogonal” operators. Such an operator preserves distances and angles. I.e. by definition, the columns of the matrix are all of length one and mutually orthogonal, so A takes the standard basis to a sequence of n mutually orthogonal unit vectors, i.e. to an “orthonormal basis”.

We don't need to be actually in \mathbb{R}^n for this discussion, but we do need an inner product. This is worth discussing since one may encounter spaces in which there is a natural inner product but where one does not have a natural basis, so they are not naturally identifiable with \mathbb{R}^n .

Def: A (real) inner product space is a real vector space V with a map $V \times V \rightarrow \mathbb{R}$, sending $(v, w) \rightarrow v \cdot w$, and satisfying the usual properties of the dot product in \mathbb{R}^n : $v \cdot w = w \cdot v$, $v \cdot v > 0$ unless $v=0$, $v \cdot (u+w) = v \cdot u + v \cdot w$, $v \cdot (tw) = t(v \cdot w)$, i.e. symmetric, positive definite, and bilinear.

Def: i) We say v and w are “orthogonal” vectors iff $v \cdot w = 0$. The zero vector is thus orthogonal to every vector, and is the only vector that is orthogonal to itself.
ii) The “length” $|v|$ of a vector v is defined to be $|v| = \sqrt{v \cdot v}$. Note that $|v| = 0$ if and only if $v=0$, and $|v| > 0$ for all non zero vectors v .

Lemma: If V is a finite dimensional inner product space, then V is naturally isomorphic to its dual space by the map $V \rightarrow V^*$ taking v to $v(\cdot)$.

Proof: This map is linear since an inner product is bilinear, i.e. linear in each variable. It is injective since if $v \cdot w = 0$ for all w , then $v \cdot v = 0$ and then $v = 0$ by positive definiteness. It is surjective since V and V^* have the same dimension, so every linear injection is also surjective. **QED.**

Since a (finite dimensional) inner product space V is isomorphic to its dual, the transpose of a linear operator $T: V \rightarrow V$ defines a linear operator $T^*: V \rightarrow V$ on the same space. I.e. if v is a vector in V , it corresponds to the linear function $f = v(\cdot)$ in V^* , and thus T^* of that element would be obtained by preceding that function by T , i.e. $T^*(f)$ would be the linear function $v \cdot T(\cdot)$. But we want this to be an element of

V , so we need to identify the element $T^*(v)$ in V , such that the functions $T^*(v)(\cdot)$ and $v.T(\cdot)$ are equal. We know there is exactly one such element by the previous lemma, so we just define $T^*(v)$ to be the unique vector in V such that $T^*(v)(w) = v.T(w)$, for all w in V . Then $T^*:V \rightarrow V$ is a linear operator, called the transpose of T , in case V is an inner product space.

Thus T^* is the composition of the maps $V \rightarrow V^* \rightarrow V^* \rightarrow V$, where the first and last maps are the isomorphisms of V with V^* in the lemma, and the middle map is the abstract transpose of T defined earlier, i.e. it is the map “preceding by T ”.

Def: If V and W are inner product spaces, an “isometric isomorphism” is a linear isomorphism $T:V \rightarrow W$ that carries the inner product of V into the one in W ; i.e. such that $u.v = T(u).T(v)$ for all u,v in V .

We know a basis $\{v_1, \dots, v_n\}$ for a space V defines an isomorphism of V with \mathbb{R}^n , but if V is an inner product space, we would like for the isomorphism to be an isometry for the usual inner product on \mathbb{R}^n . For this to be true we need to find an orthonormal basis for V , i.e. a basis of mutually orthogonal unit vectors. I.e. assume $\{u_1, \dots, u_n\}$ is an orthonormal basis, and $a_1u_1 + \dots + a_nu_n$ and $b_1u_1 + \dots + b_nu_n$ are any vectors, corresponding under this basis to the coordinate vectors (a_1, \dots, a_n) and (b_1, \dots, b_n) in \mathbb{R}^n . Then since $u_j.u_k = 0$ unless $j=k$ when it equals 1, the inner product expands as $(a_1u_1 + \dots + a_nu_n).(b_1u_1 + \dots + b_nu_n) = a_1b_1(u_1.u_1) + \dots + a_nb_n(u_n.u_n) = a_1b_1 + \dots + a_nb_n$, the usual dot product in \mathbb{R}^n .

An orthonormal basis can always be found as in the following proposition. The procedure is often called the Gram - Schmidt process, but it is just the familiar geometric operation of orthogonal projection.

Proposition: Every finite dimensional inner product space has an orthonormal basis, hence is isometrically isomorphic to some \mathbb{R}^n with its usual dot product.

Proof: Start from any basis $\{v_1, \dots, v_n\}$ and, since none of the vectors is zero, we make the first vector unit length by dividing by its length, i.e. set $u_1 = v_1/|v_1|$. Now since we know $u_1.v_2 = |u_1||v_2|\cos(t) = |v_2|\cos(t)$, where t is the angle between u_1 and v_2 , it follows that this number is the oriented length of the projection of v_2 onto the line through u_1 . Hence $|v_2|\cos(t).u_1 = (v_2.u_1)u_1$ is the vector component of v_2 parallel to u_1 , and thus $w_2 = v_2 - (v_2.u_1)u_1$ is the vector component of v_2 perpendicular to u_1 . Since v_2 does not depend on v_1 , this component cannot be zero. Now take $u_2 = w_2/|w_2|$, and we have u_2 perpendicular to u_1 and also of length one. Continue...

I.e. subtract off the components of v_3 that are parallel to both u_1 and u_2 , getting $w_3 = v_3 - (v_3 \cdot u_1)u_1 - (v_3 \cdot u_2)u_2$, and then set $u_3 = w_3/|w_3|$, and so on... until we have found $\{u_1, \dots, u_n\}$. Since no basis vector v depends on the previous v 's, hence not on the corresponding u 's, none of these w 's will be zero. These u 's will then be mutually orthogonal, and all of length one, hence an orthonormal basis of V . This basis defines an isometric isomorphism $\mathbb{R}^n \rightarrow V$. **QED.**

Ex. If $T: V \rightarrow V$ is an operator on an inner product space and we choose an orthonormal basis, then in this basis the matrix of $T^*: V \rightarrow V$ will be the transpose of the matrix for T .

Def: An operator $T: V \rightarrow V$ on an inner product space is called symmetric if $T = T^*$, and orthogonal if $T^* = T^{-1}$.

Rmk: The same proof given above for matrices proves that every symmetric operator on a finite dimensional inner product space has an orthonormal basis of eigenvectors. Next we prove an analogous result for orthogonal operators.

Prop: If T is real orthogonal on V , and $\dim V$ is finite/ \mathbb{R} , then V is an orthogonal direct sum of an eigenspace on which $T = \text{Id}$, an eigenspace on which $T = -\text{Id}$, and a collection of invariant 2-planes, on each of which T is a rotation.

proof: The first step is to produce an “indecomposable” invariant subspace U of dimension ≤ 2 , i.e. a subspace of $\dim \leq 2$ that is not a direct product of smaller invariant subspaces.

Lemma: For every non constant factor f of the minimal polynomial of a linear operator T on a finite dimensional space V , over any field, $\ker(f(T)) \neq \{0\}$.

proof: If the minimal polynomial $m = f \cdot g$, and $f(T)$ is injective, since $f(T)g(T)(x) = 0$ for all x , then $g(T)(x) = 0$ for all x . Then g is a polynomial of lower degree than m that annihilates all of V , a contradiction. **QED.**

Cor: If f is an irreducible factor of the minimal polynomial of any operator T on any finite dimensional space V , over any field, there is an indecomposable invariant subspace U of $\ker f(T)$ in V of dimension equal to $\deg(f)$.

proof: If $x \neq 0$ lies in $\ker f(T)$, the T -cyclic subspace U generated by $\{x, T(x), T^2(x), \dots\}$ is invariant, and has dimension n for the smallest n such that $T^n(x)$ depends on $\{x, T(x), \dots, T^{(n-1)}(x)\}$. Let $n = \dim U$, and assume $T^n(x) = a_0 \cdot x + a_1 \cdot T(x) + \dots + a_{n-1} \cdot T^{(n-1)}(x)$. Then $g(T)(x) = 0$ for $g(t) = t^n - a_{n-1} \cdot t^{(n-1)} - \dots - a_1 \cdot t - a_0$. Since $g(T)$ commutes with powers of T , $g(T)$ annihilates all $T^k(x)$

hence all of U , so we have found a monic polynomial g with $\deg(g) = \dim U$, that annihilates the restriction of T to U . We claim this is the minimal such polynomial.

To see that, if $P(T) = b_0 + b_1T + b_2T^2 + \dots + T^m$ is any polynomial which annihilates T restricted to U , then $b_0x + b_1T(x) + b_2T^2(x) + \dots + T^m(x) = 0$, so $T^m(x) = -b_0x - b_1T(x) - b_2T^2(x) - \dots - T^{(m-1)}(x)$ is a linear combination of the previous powers of $T(x)$, so $\dim(U) \leq m$. Thus the minimal polynomial of the restriction of T to U has degree $= \dim(U)$.

Since U is a subspace of $\ker f(T)$, $f(T)$ vanishes on U , so f is a multiple of that restricted minimal polynomial. Since f is irreducible, f must equal that restricted minimal polynomial, so $\dim(U) = \deg(f)$. Since we could have chosen $x \neq 0$ in any non zero invariant subspace of $\ker f(T)$, this argument shows the dimension of any non trivial invariant subspace of $\ker f(T)$ has dimension $\geq \deg(f)$. In particular U is a minimal $\neq 0$ invariant subspace, hence indecomposable. **QED.**

Since an irreducible polynomial over the reals has degree ≤ 2 , every linear operator on a finite dimensional real vector space has an indecomposable invariant subspace U of dimension 1 or 2, corresponding to a real irreducible linear or quadratic factor of the minimal polynomial. Since T is length preserving, on every one dimensional invariant subspace, $T = \pm \text{Id}$.

If U has dimension 2, since T is length and angle preserving on U , if (x_1, x_2) is an orthonormal basis for U , Tx_1 is length one and orthogonal to Tx_2 . Hence $T(x_1) = \cos(a)x_1 + \sin(a)x_2$, and $Tx_2 = -\sin(a)x_1 + \cos(a)x_2$, or else $Tx_2 = \sin(a)x_1 - \cos(a)x_2$.

But the minimal polynomial on U is irreducible over \mathbb{R} . If $T(x_1) = \cos(a)x_1 + \sin(a)x_2$, and $Tx_2 = \sin(a)x_1 - \cos(a)x_2$, the minimal polynomial equals $t^2 - 1$, which is reducible. So in fact, $T(x_1) = \cos(a)x_1 + \sin(a)x_2$, and $Tx_2 = -\sin(a)x_1 + \cos(a)x_2$, which is a rotation. Since T is orthogonal also on U^\perp , which is T invariant, the theorem is proved by induction on $\dim V$. **QED.**

Cor: Every orthogonal map of \mathbb{R}^2 to itself is a rotation or reflection. Every orthogonal map of \mathbb{R}^3 to itself is either a rotation or a rotation followed by a reflection.

More generally, in every even dimensional space \mathbb{R}^{2n} , an orthogonal operator T admits a decomposition of the space into orthogonal 2 planes, such that either T is composed of a rotation in each of these planes, or of rotations in all but one of them and a reflection in that one.

An orthogonal operator T on an odd dimensional space \mathbb{R}^{2n+1} admits a decomposition of the space into an orthogonal sum of 2 planes and one orthogonal line, such that T is composed of rotations in all the 2 planes and equals $\pm I$ in the remaining line.

In all cases T is called orientation reversing or orientation preserving according to whether there is or is not a reflection present in this decomposition. An orientation preserving orthogonal map is called simply a rotation. (These 2 cases are distinguished by the sign of the determinant \pm , equivalently by the constant term of the characteristic polynomial.)

Some remarks on complex spectral theorems

We have seen that inductive arguments to prove diagonalizability need existence of an eigenvector to get started, and then need a decomposability property to split off a complementary invariant subspace on which one can continue the argument. We can make this precise with a definition.

Definition: An operator $T:V \rightarrow V$ is called “semi simple” if for every T -invariant subspace U , there is a T -invariant complement, i.e. there is a T -invariant subspace W such that the addition map $U \times W \rightarrow V$ taking (u,w) to $u+w$, is an isomorphism.

Lemma: An operator $T:V \rightarrow V$ on a finite dimensional space V over k , is diagonalizable over k if and only if it is semi simple, and its minimal polynomial splits into linear factors in k .

Proof: If the minimal polynomial splits, there is an eigenvector and its span is T -invariant. Assuming semi simplicity there is an invariant complement on which the minimal polynomial is a factor of the original one hence also splits in k , so we can finish by induction on dimension.

Conversely, if T is diagonalizable there is a basis of eigenvectors v_1, \dots, v_n . If U is an invariant subspace with basis w_1, \dots, w_r , then reducing the sequence $w_1, \dots, w_r, v_1, \dots, v_n$ to a basis give us w_1, \dots, w_r , followed by $n-r$ of the basis vectors v_j . Since the v_j are all eigenvectors these span an $n-r$ dimensional T -invariant complement of U . **QED.**

Now in our attempts to prove spectral theorems over \mathbb{R} , we could use perpendicularity to get semi simplicity of both symmetric and orthogonal matrices, but we only had splitting of the minimal polynomial in the symmetric case, and even that took a bit of work. If we consider complex matrices we always have splitting of the minimal polynomial, but we seem to lose the decomposition argument by dot products since in this case a subspace is no longer perpendicular

to the subspace of vectors that dot to zero with it. E.g. recall that with the usual dot product, the vector $(1, i)$ in C^2 is “perpendicular” to itself! But we can fix that by tweaking the definition of the dot product, and hence of perpendicularity.

Complex conjugates: Recall that if $x+iy$ is a complex number, its “complex conjugate” is the number $x-iy$. This has the useful property that for every $x+iy$, the product $(x+iy)(x-iy) = x^2 + y^2 > 0$ unless $x = y = 0$.

Definition: On the complex coordinate space C^n , define the “hermitian product” of two vectors $v = (a_1, \dots, a_n)$ and $w = (b_1, \dots, b_n)$ to be $\langle v, w \rangle = a_1 b_1' + \dots + a_n b_n' = v \cdot w'$, where b' is the complex conjugate of b .

Then we have for all vectors u, v, w and all complex numbers c :

$$\begin{aligned} \langle v, w \rangle &= \langle w, v \rangle', \\ \langle u+v, w \rangle &= \langle u, w \rangle + \langle v, w \rangle, \\ \langle v, u+w \rangle &= \langle v, u \rangle + \langle v, w \rangle, \\ \langle cv, w \rangle &= c \langle v, w \rangle, \\ \langle v, cw \rangle &= c' \langle v, w \rangle, \\ \langle v, v \rangle &\geq 0, \text{ i.e. } \langle v, v \rangle \text{ is both real and non-negative.} \end{aligned}$$

Hermitian orthogonality

Defn: Two vectors v, w in C^n are orthogonal with respect to the hermitian pairing if $\langle v, w \rangle = 0$, and for a subspace U of C^n , its hermitian orthogonal complement $U_{\text{perp}} = \{\text{all vectors } w \text{ in } C^n \text{ such that } \langle v, w \rangle = 0 \text{ for all vectors } v \text{ in } U\}$.

With this new definition we get what we need for decomposability.

Ex. If U is a subspace of C^n and U_{perp} is its hermitian orthocomplement, then the addition map $U \times U_{\text{perp}} \rightarrow C^n$ is an isomorphism.

We can also define a hermitian length for vectors in C^n .

Defn: If v is a vector in C^n , since $\langle v, v \rangle \geq 0$, define $|v| = \sqrt{\langle v, v \rangle} \geq 0$. A vector is normal if it has length one. Note: $|v| = 0$ only for $v = 0$.

Now it is easy to prove some complex spectral theorems by the usual arguments.

Defn: If A is an n by n complex matrix, define its hermitian adjoint A^* as the complex conjugate of its usual transpose.

Ex. With these new definitions then we have $A^{**} = A$, and for all v, w in C^n , $\langle Av, w \rangle = \langle v, A^*w \rangle$.

Defn: A complex n by n matrix A is called hermitian if $A = A^*$, and unitary if $A^* = A^{-1}$.

Ex. Prove that if A is either hermitian or unitary, then there is an orthonormal basis of eigenvectors for A .

There is one more complex spectral theorem we can prove.

Defn: An n by n complex matrix A is called normal if $AA^* = A^*A$, i.e. A commutes with its hermitian adjoint.

Note: Hermitian and unitary operators are special cases of normal operators.

Theorem: A normal complex matrix admits an orthonormal eigenbasis.

Lemma: If A is any complex n by n matrix and A leaves a subspace U of C^n invariant, then A^* leaves U^\perp invariant.

proof: If v is in U and w is in U^\perp , then Av is also in U so $0 = \langle Av, w \rangle = \langle v, A^*w \rangle$, so A^*w is also in U^\perp . **QED.**

Lemma: If A is a normal complex n by n matrix, and v is an eigenvector for A , then A^*v is also an eigenvector of A with the same eigenvalue.

proof: By hypothesis, for some constant c , $Av = cv$, so $A.A^*v = A^*Av = A^*.cv = c.A^*v$. **QED.**

sketch of proof of theorem: Since the minimal polynomial splits over C , there is an eigenvalue c and a non empty space U of eigenvectors for c . By the second lemma above, A^* maps U into itself, so by the first lemma $A = A^{**}$ maps U^\perp into itself. Since U^\perp is A invariant, and $AA^* = A^*A$ still holds on U^\perp , as do all properties of the hermitian product, we can continue the argument. I.e. we can find another eigenspace inside U^\perp and split it off,**QED.**

Remark: To make this argument go through we have to show we get the orthogonal decomposition theory also for subspaces of C^n . The point is that by dimension theory of intersections, if W is a subspace of V , and V is a subspace of C^n , then V is isomorphic to $W \times (W^\perp \cap V)$, where $(W^\perp \cap V)$ is the perp of W relative to V . Or we could define a hermitian inner product space abstractly and not be dependent on coordinates, as we did in the real case. A Gram - Schmidt argument would again show there is actually no gain in generality, i.e.

finding an orthonormal basis would show every n dimensional hermitian inner product space is isometrically isomorphic to C^n with its hermitian product.

Remark: There is a cute trick for deducing the real symmetric spectral theorem from the complex hermitian one. Note that if A is hermitian, i.e. if $A = A^*$, with eigenvector v , then $c\langle v, v \rangle = \langle cv, v \rangle = \langle Av, v \rangle = \langle v, A^*v \rangle = \langle v, Av \rangle = \langle v, cv \rangle = c'\langle v, v \rangle$. So $c|v|^2 = c'|v|^2$, and since $|v|^2 \neq 0$, hence $c = c'$, where c' is the complex conjugate of c . Thus every complex eigenvalue of a hermitian matrix is real. So given a real symmetric matrix, just think of it as a complex matrix, and notice that it is hermitian since all the entries are real. Then its characteristic polynomial $\det(A - X.I)$, which is the same no matter how we regard A , has only real roots. Thus a real symmetric A is not only semi simple, but also its characteristic polynomial splits in R . QED.

Remark: There is also a structure theorem for real operators that are normal in the sense that they commute with their transposes. This decomposes the space into mutually orthogonal invariant subspaces of dimensions one or two. In this case the minimal polynomial factors over R into distinct irreducible factors of degrees one or two. Of course any real polynomial factors over R into irreducible factors of degrees ≤ 2 , but here the point is that the irreducible factors are all distinct.

A fundamental example: linear differential equations with constant coefficients.

An important application of linear algebra is to solving linear ordinary differential equations (with constant coefficients), and systems of them. We discuss this next.

Let $V =$ the set of infinitely differentiable (“smooth”) real valued functions on the real line. Then V is infinite dimensional (it contains all polynomials of all degrees), the derivative map $D: V \rightarrow V$ is linear, and surjective by the fundamental theorem of calculus, and $\ker(D) =$ the one dimensional space of constants by the mean value theorem. If $P(X) = X^n + a_{n-1}X^{(n-1)} + \dots + a_1X + a_0$ is a polynomial with real coefficients a_j , replacing X by D gives another linear operator $P(D)$ on V . Using our dimension theory, elementary facts about polynomials, and basic calculus, we will prove the following:

Proposition: $\dim.\ker P(D) = \deg(P) = n$; and we will exhibit an explicit basis for $\ker P(D)$.

Note: If $f^{(k)}$ denotes the k th derivative of a smooth function f , then $\ker(P(D)) = \{ \text{all } f \text{ in } V \text{ such that: } f^{(n)} + a_{n-1}f^{(n-1)} + \dots + a_1f' + a_0f = 0 \}$. We start with the simplest case, $\text{degree}(P) = \text{one}$.

Lemma: If c is a real constant, then $\ker(D-c) = \{ a.e^{cx}, \text{ for all real scalars } a \}$, i.e. e^{cx} , is a basis.

Pf: If $(D-c)(f) = 0$, then $f' = cf$, so f/e^{cx} , has derivative zero, so is constant, i.e. $f = a.e^{cx}$. **QED.**

Cor: If $P(X) = (X-c_1)^{r_1}(X-c_2)^{r_2}\dots(X-c_t)^{r_t}$ with c_i distinct real numbers, and $r_i > 0$, then $\dim.\ker(P(D)) \leq \text{degree}(P) = r_1 + \dots + r_t = n$.

Pf: Induction on $\text{deg}(P)$. Since all polynomials in D commute, $\ker(D-c_1)$ is contained in $\ker P(D)$, and $(D-c_1)$ injects $\ker P(D)/\ker(D-c_1)$ into $\ker(D-c_1)^{r_1-1}(D-c_2)^{r_2}\dots(D-c_t)^{r_t}$. Hence $\dim.\ker P(D)/\ker(D-c_1) = \dim.\ker P(D) - 1 \leq \dim.\ker(D-c_1)^{r_1-1}(D-c_2)^{r_2}\dots(D-c_t)^{r_t} \leq n-1$. **QED.**

To obtain the full equality, we exhibit $\text{deg}P(D)$ independent elements of $\ker P(D)$.

Lemma: $\dim.\ker(D-c)^r = r$; indeed $\{ e^{ct}, t e^{ct}, t^2/2 e^{ct}, \dots, t^{r-1}/(r-1)! e^{ct} \}$ is a basis.

Pf: By the previous corollary, it suffices to prove $\dim.\ker(D-c)^r \geq r$.

Ex: Show $(D-c)$ takes each element of this basis to the previous element, and annihilates the first element.

By the exercise, these are all in $\ker(D-c)^r$ and since the first k are in $\ker(D-c)^k$, but the $k+1$ st is not, no element depends on the previous ones, hence they are independent. **QED.**

Rmk: We have called a basis with the property in the previous exercise “cyclic”. In this basis the matrix for $(D-c)$ consists of columns with most entries zero, but with “1’s” just above the diagonal: $(0, \dots, 0), (1, 0, \dots, 0), (0, 1, 0, \dots, 0), \dots, (0, \dots, 0, 1, 0)$. The matrix of D adds the constant c along the diagonal: $(c, 0, \dots, 0), (1, c, 0, \dots, 0), (0, 1, c, 0, \dots, 0), \dots, (0, \dots, 0, 1, c)$. If we reverse the basis ordering, we can put the “1’s” just below the diagonal: $(c, 1, 0, \dots, 0), (0, c, 1, 0, \dots, 0), \dots, (0, \dots, 0, c, 1), (0, \dots, 0, c)$, as we did earlier in discussing Jordan form.

Lemma: If all c_i are distinct, $W = \ker(D-c_1)^{r_1}(D-c_2)^{r_2}\dots(D-c_t)^{r_t}$ is isomorphic to

the product of the subspaces $W_i = \ker(D - c_i)^{r_i}$.

Pf: This is our earlier relatively prime decomposition theorem. **QED.**

Since the dimension of a product is the sum of the dimensions of the factors, we are done in case the polynomial $P(X)$ factors completely into linear factors over the reals, i.e. when there are no irreducible quadratic factors. In particular, we see that in the given basis for $\ker P(D)$, the matrix for D consists of blocks of the form described in the previous remark. I.e. the j th factor $(D - c_j)^{r_j}$ gives an r_j by r_j block with the constant c_j along the diagonal, and 1's just above or just below the diagonal. Such a matrix is thus in Jordan form. Note also that it is easy to read off the polynomial P from this matrix. The simplest case occurs when all the exponents $r_j = 1$, i.e. when the polynomial P is a product of distinct linear factors. In this case the basic solutions are all of form $e^{c_j t}$, and the matrix of D in this basis is "diagonal", i.e. the constants c_j appear along the diagonal and the off diagonal entries are zero.

We have shown earlier that a Jordan form matrix can be obtained, not just for the special operator D , but for any linear operator that satisfies a polynomial that factors into linear factors. Recall however that if $Q(X)$ is any factor of a polynomial $P(X)$ satisfied by T , then the operator TxT acting on the product space $\ker P(T) \times \ker Q(T)$ will still satisfy the polynomial $P(X)$, but the matrix will have blocks corresponding to both P and Q . Thus the polynomial P satisfied by the operator does not fully determine the matrix in the general case. I.e. in general there may be several blocks corresponding to the same root c of the polynomial P , so a more general operator resembles a direct product of copies of the differential operator D , acting on spaces corresponding to polynomials all of which divide P . In particular, whereas for the differential operator above, the dimension of $\ker P(D)$ equals the degree of P , for a general space V on which an operator T acts satisfying a polynomial $P(X)$ (of minimal degree), we can only say the dimension of V is $\geq \text{degree}(P)$. I.e. the example of the derivative above corresponds to an operator whose Jordan matrix is a single elementary Jordan block.

Digression: What happens when P has irreducible quadratic factors?

For application to solving differential equations, we will sketch the solution of more general linear constant coefficient differential operators, when the polynomial $P(X)$ has some irreducible quadratic factors. Recall the previous lemma is true if P is any product of powers of irreducible factors, whether or not the factors are linear. If $U =$ all infinitely differentiable complex valued functions on the real line, then $U \approx V \times V$, where (f, g) in $V \times V$ corresponds to $f + ig$ in U , i.e. f

and g are the real and imaginary parts of a function in U . Then a polynomial $P(X)$ with real coefficients again defines a linear operator on U by sending $f+ig$ to $P(D)(f) + i P(D)(g)$. Thus $f+ig$ is in $\ker P(D)$ in U , if and only if both f and g both are in $\ker P(D)$ in V , i.e. $\ker P(D)$ in U is the direct product of two copies of $\ker P(D)$ in V .

By the same arguments as above, if c is a complex constant, the subspace $\ker(D-c)$ of U has complex dimension one, and basis e^{ct} , where e^{ct} is defined by the usual exponential series. If P has complex coefficients, $P(D)$ still acts on U , but the real part of $P(D)(f+ig)$ will usually not equal $P(D)(f)$. Whether $P(X)$ has real or complex coefficients, the arguments above show again that $\ker P(D)$ as a subspace of U , has complex dimension $\leq \deg(P)$.

Ex: Show that in a linear combination with complex coefficients of real valued functions, the real and imaginary parts are linear combinations of those same functions with real coefficients. Hence if some real valued functions are dependent over the complex field, they are also dependent over the real field.

Lemma: As functions of the real variable t , $e^{it} = \cos(t) + i \sin(t)$.

Pf: The function $\cos(t) + i \sin(t)$ lies in $\ker(D-i)$, hence is a constant multiple of e^{it} , and evaluating at $t = 0$, shows the multiplier is one. QED.

Cor: As a subspace of V , $\ker(D^2 + 1)$ has as real basis $\{\cos(t), \sin(t)\}$.

Pf: We know the corresponding complex subspace of U has basis $\{e^{it}, e^{-it}\}$, in particular it has complex dimension two. Thus any three functions in $\ker(D^2+1)$ in V , are dependent as elements of U over the complex field, hence also over R . Thus $\ker(D^2 + 1)$ has real dimension at most two. But if $a.\cos(t)+b.\sin(t) = 0$, then evaluating this function and its derivative at $t = 0$, implies $a = b = 0$. Since $\cos(t)$ and $\sin(t)$ are independent over R , they are a real basis. QED.

Rmk: Note $\ker(D^2 + 1)$ in U has complex bases $\{\cos(t), \sin(t)\}$ and $\{e^{it}, e^{-it}\}$.

Similarly, one can check that if $P(X)$ is an irreducible real quadratic polynomial with complex roots $a \pm bi$, then both $\{e^{(a+bi)t}, e^{(a-bi)t}\}$, $\{e^{at}.\cos(bt), e^{at}.\sin(bt)\}$ are complex bases of $\ker P(D)$ in U , and $\{e^{at}.\cos(bt), e^{at}.\sin(bt)\}$ is a real basis of $\ker P(D)$ in V . Thus for any real polynomial $P(X)$, the subspace $\ker P(D)$ of V has real dimension $= \deg(P)$, and the corresponding subspace of U has complex dimension $= \deg(P)$. Moreover any real basis for the subspace of V serves also as a complex basis for the corresponding subspace of U .

For powers of this irreducible quadratic factor, we get versions of the Jordan matrices. E.g. for $P^2(D)$, if $u = a+ib$, $u' = a-ib$ are the roots of $P(X)$, we get these bases: $\{t.e^{ut}, t.e^{u't}, e^{ut}, e^{u't}\}$; $\{t.e^{at}\sin(bt), t.e^{at}\cos(bt), e^{at}\sin(bt), e^{at}\cos(bt)\}$; which lead to these matrices for the operator D on $\ker.P^2(D)$:

$$\begin{bmatrix} u & 0 & 0 & 0 \\ 0 & u' & 0 & 0 \\ 1 & 0 & u & 0 \\ 0 & 1 & 0 & u' \end{bmatrix}, \text{ (a version of the Jordan matrix with basis ordered differently).}$$

and its real analogue:

$$\begin{bmatrix} a & -b & 0 & 0 \\ b & a & 0 & 0 \\ 1 & 0 & a & -b \\ 0 & 1 & b & a \end{bmatrix}.$$

Note that $P(X)$ here = $X^2 -2aX + (a^2+b^2)$, and the matrix that appears here

$$\begin{bmatrix} a & -b \\ b & a \end{bmatrix},$$

has $P(X)$ as characteristic and minimal polynomial, just like the companion matrix for this polynomial, which appeared in our general Jordan matrices.

End of Digression.

Solving diagonalizable systems of linear differential equations:

Now we ramp up from solving the differential equation $Df = cf$, where c is a constant and $f(t)$ is a real or complex valued function of t , to the system $Dx = Ax$, where $x(t) = (x_1(t), \dots, x_n(t))$ is a (real or complex) vector valued function of t . Note that if A is a diagonal matrix, this is nothing new. I.e. the equation becomes just $Dx_j(t) = c_j.x_j(t)$, so for all j , $x_j(t) = a_j.e^{(c_j.t)}$ for some constant a_j . Thus a basic solution is $e^{(c_j.t)}.e_j$, where e_j is the j th standard basis vector. I.e. these n solutions give a real basis for the solution space. But a diagonal matrix is just one with the standard basis vectors e_j as an eigenbasis. We claim that for any diagonalizable matrix A , a basic solution of $Dx = Ax$ is just $x(t) = e^{(c_j.t)}.v_j$, where v_j is a basic eigenvector. To see this, first observe a tiny extension of the product rule for D .

Lemma: $D(M.x(t)) = M.D(x(t))$, if M is a constant matrix and D is differentiation

of the vector valued function $x(t)$.

Proof: The j th entry of $M \cdot x(t)$ is a linear combination of the entries of $x(t)$ with coefficients the j th row of M . Thus the j th entry of $D(M \cdot x(t))$ is the same linear combination of the derivatives of the entries of $x(t)$, i.e. the j th entry of $M \cdot D(x(t))$. **QED.**

Proposition: If A is a diagonalizable matrix, with eigenbasis $\{v_j\}$, and eigenvalues $\{c_j\}$, the vector differential equation $Dx = Ax$, where $x(t)$ is a vector-valued function, has an n dimensional space of solutions, with basis $e^{(c_j \cdot t)} \cdot v_j$, where here the function $e^{(c_j \cdot t)}$ is multiplied by every (constant) entry in the eigenvector v_j .

Proof: By hypothesis, there is an invertible matrix Q , whose columns are eigenvectors $\{v_j\}$ of A , such that $Q^{-1}AQ = M$ is diagonal with diagonal entries $\{c_j\}$ = the eigenvalues of A . Since we know the basic solutions of $Dy = My$ are $y = e^{(c_j \cdot t)} \cdot e_j$, we put $x = Qy = e^{(c_j \cdot t)} \cdot Qe_j = e^{(c_j \cdot t)} \cdot v_j$, where v_j is a basic eigenvector of A . Then since $Dy = My = Q^{-1}AQy$, we have $DQy = QDy = AQy$, i.e. $Dx = Ax$. Reasoning backwards, these are the only solutions, i.e. $Dx = Ax$ if and only if $x = Qy$ where $Dy = My$. **QED.**

Cor: Given any vector v in \mathbb{R}^n , the solution of $Dx(t) = A \cdot x(t)$ can be chosen uniquely so that at $t = 0$ the value of $x(0)$ is v .

Proof: If $x(t)$ is the j th basic solution, $x(0) = v_j$, and the vectors $\{v_j\}$ form a basis of \mathbb{R}^n . **QED.**

Exponentiating a matrix

In fact there is no need for the matrix to be diagonalizable in order to solve such systems. Indeed there is a way to interpret all solutions of such systems as direct generalizations of the fact that $c \cdot e^{(at)}$, for c an arbitrary constant real number, is the general solution of $f' = a \cdot f$. I.e. for each square matrix A , we can define a matrix valued function $e^{(At)}$, or a matrix with functions as entries, so that for C an arbitrary constant real vector, the vector valued function $x(t) = e^{(At)} \cdot C$ is a general solution of the system $Dx = Ax$. (When A is a one by one matrix, this specializes to the familiar solution $e^{(at)} \cdot c = c \cdot e^{(at)}$, since multiplication of one by one matrices is commutative.) Since $e^{(At)} \cdot C$ is a linear combination of the columns of the matrix $e^{(At)}$ with coefficients from C , this means the columns of $e^{(At)}$ form a basis of solutions of the system $Dx = Ax$. In fact these are exactly the solutions we found for diagonal and diagonalizable matrices.

I.e. recall that the real number $e^{(at)}$ is defined by an infinite series of terms each of which is a constant times a power of (at) . To define $e^{(A \cdot t)}$, where A is a matrix,

we write down the same series with terms which are each a constant times a power of the matrix A . So $e^{(A.t)} = \text{Id} + A.t + (1/2!)A^2t^2 + \dots$, which converges for the same reason the usual series for $e^{(at)}$ does.

If M is diagonal, with diagonal entries a_1, \dots, a_n , M^k is also diagonal with entries k th powers of the a_j , so then this series converges to the diagonal matrix with diagonal entries $e^{(a_j.t)}$. Thus $e^{(M.t)}$ is the matrix whose j th column is $e^{(a_j.t)}.e_j$, the j th basic solution we gave for the system $Dx = Mx$.

If A is only diagonalizable, with $A = QMQ^{-1}$ and M diagonal, then $A^k = Q.M^k.Q^{-1}$, since the Q 's and Q^{-1} 's in the middle cancel, so $e^{(At)} = Q.e^{(Mt)}.Q^{-1}$. I had expected the columns of this matrix to be the basic solutions mentioned above, namely $e^{(a_j.t)}.v_j$, where v_j is an eigenbasis for A , but this would pick out a special such eigenbasis, and there are of course many of them. As above however, the matrix Q carries solutions of $Dx = Mx$ into solutions of $Dx = Ax$, and this is the case here. I.e. the basic solutions of $Dx = Mx$ are $e^{(a_j.t)}.e_j$ and if we apply Q to these we get $Qe^{(a_j.t)}.e_j = e^{(a_j.t)}.Qe_j = e^{(a_j.t)}.v_j$. These are the columns of the matrix $Q.e^{(Mt)} = e^{(At)}.Q$.

Of course the columns of $e^{(At)} = Q.e^{(Mt)}.Q^{-1}$ are also a basis of solutions of $Dx = Ax$, but I don't know how to express those columns in any simpler way. In fact if A is any square matrix at all, the series for $e^{(At)}$ still converges to a matrix whose columns give a basis of solutions of $Dx = Ax$, as one can see just by differentiating the series term by term. But I don't know a simple way to write out the columns of that matrix. However, in case the minimal polynomial of A factors completely into (not necessarily distinct) linear factors there is a way to do this as we will see next, using the Jordan normal form for such matrices.

Recall, if the minimal polynomial of a linear operator T factors completely into linear factors, then T is uniquely expressible as a sum $T = S+N$, of a diagonalizable operator S and a nilpotent operator N such that $SN = NS$. Then since $e^{tA} = e^{t(S+N)} = e^{tS} . e^{tN}$, we can express the matrix e^{tA} as a product of two matrices we can compute, namely e^{tS} and e^{tN} . I.e. we have already shown how to compute $e^{t.S}$ for diagonalizable S , and for nilpotent N the series for e^{tN} is finite, and can thus be summed as a polynomial in N . Thus, over the complex scalar field, we can always find a Jordan matrix for any operator and thus solve the corresponding linear system of differential equations.

Question: Can we also use the rational canonical form to solve systems in terms of the solutions of the equations $P(D)(f)$?

Appendix: Summary of determinants.

Definition: If A is an $n \times n$ matrix over R , define for each (i,j) with $1 \leq i,j \leq n$, $A_{ij} =$ the $(n-1) \times (n-1)$ matrix obtained by deleting from A the i th row and j th column. Then define determinants recursively as follows: If $n=1$, and $A = (a)$ define $D(A) = a$. If we have defined D for all $(n-1) \times (n-1)$ matrices, and if A is in $\text{Mat}_n(R)$, then set $D(A) = a_{11} D(A_{11}) - a_{12} D(A_{12}) + \dots \pm a_{1n} D(A_{1n})$, (expansion by the first row).

Examples: $D \begin{vmatrix} a & b \\ c & d \end{vmatrix} = ad - bc$. $D \begin{vmatrix} a & b & c \\ d & e & f \\ g & h & i \end{vmatrix} = a(ei - hf) - b(di - fg) + c(dh - eg)$.

Theorem: $D(A)$ is n -linear and alternating as a function of the rows and columns of A . Hence,

- (i) If A' is the result of interchanging two rows or columns of A , $D(A') = -D(A)$.
- (ii) If A' is the result of adding to one row (or column) of A , a scalar multiple of another row (or column), then $D(A') = D(A)$.
- (iii) If A' is the result of multiplying through a row or column of A by a scalar c , then $D(A') = c \cdot D(A)$.

Corollary: If A is upper or lower triangular, e.g. diagonal, then $D(A) = \prod a_{ii}$, the product of the diagonal entries. If A is a matrix with two equal rows or columns, then $D(A) = 0$.

Prop: (i) $D(A^*) = D(A)$, where A^* is the transpose of A .

(ii) If A, B are both $n \times n$ matrices, then $D(AB) = D(A)D(B)$; in particular if E is invertible then $D(E^{-1}AE) = D(A)$, so all matrices for the same linear map $f: V \rightarrow V$ with respect to any basis of V , have the same determinant.

Remark: Using these properties of determinants we can make any row the first and expand, and can interchange rows and columns and expand. In particular, we can expand a determinant by any row, not just the first. I.e. for any choice of row, say i , we have $D(A) = \sum_j (-1)^{i+j} a_{ij} D(A_{ij})$.

And we can also expand determinants by columns; i.e. for any choice of column, say j , we have $D(A) = \sum_i (-1)^{i+j} a_{ij} D(A_{ij})$.

Cor (Cramer's rule): Let $B = (b_{ij}) = \text{adj}(A)$, be the matrix such that $b_{ij} = (-1)^{i+j} D(A_{ji})$. (Note the interchange of indices.) Then $AB = D(A).I = BA$. In particular, if $D(A) \neq 0$, then $D(A)^{-1}B = D(A)^{-1}.\text{adj}(A)$ is a (two sided) inverse for A . Thus a matrix A over a field is invertible if and only if $D(A) \neq 0$, and if the entries of A are only from a commutative ring, then A is invertible if and only if $D(A)$ is a unit (invertible element) in that ring.

Cor (Cayley - Hamilton): If A is a square matrix and $\text{ch}_A(X) = D(X.I-A)$ is the characteristic polynomial of A , then $\text{ch}_A(A) = 0$.

Proof: By the non commutative remainder theorem, it suffices to show that $\text{ch}_A(X)$ is divisible as a polynomial with matrix coefficients, either from the right or the left, by the linear polynomial $(X-A)$. But Cramer's rule shows the corresponding matrix $(X.I - A)$ does divide the matrix corresponding to the characteristic polynomial from both left and right, i.e. by Cramer, we have $(X.I - A).\text{adj}(X.I-A) = \text{adj}(X.I-A).(X.I-A) = \text{ch}_A(X).I$. Since the matrix ring with polynomial coefficients is isomorphic to the polynomial ring with matrix coefficients, $(X-A)$ does divide $\text{ch}_A(X)$ both from left and right. Hence A is a (right and left) root of $\text{ch}_A(X)$. **QED.**

Roy Smith