

A Second Course in Linear Algebra

A Second Course in Linear Algebra

Robert A. Beezer
Department of Mathematics and Computer Science
University of Puget Sound
beezer@pugetsound.edu

DRAFT April 4, 2014 DRAFT

© 2004-13 Robert A. Beezer
GFDL License.

Contents

1	The Fundamentals	1
1.1	Introduction	1
1.2	Direct Sums	2
1.2.1	Direct Sums	2
1.3	Orthogonal Complements	7
1.4	Invariant Subspaces	9
1.4.1	Invariant Subspaces	9
1.4.2	Restrictions of Linear Transformations	16
1.5	Reflectors	18
1.6	Projectors	21
1.6.1	Oblique Projectors	21
1.6.2	Orthogonal Projectors	23
1.7	Normal Matrices	25
1.8	Positive Semi-Definite Matrices	26
2	Matrix Decompositions	29
2.1	LU (Triangular) Decomposition	29
2.1.1	LU Decomposition, Nonsingular Case	29
2.1.2	Solving Systems with Triangular Decompositions	32
2.1.3	Computing Triangular Decompositions	34
2.1.4	Triangular Decomposition with Pivoting	36
2.2	QR (Gram-Schmidt) Decomposition	37
2.2.1	QR Decomposition via Gram-Schmidt	37
2.2.2	QR Decomposition via Householder Reflections	39
2.2.3	Solving Systems with a QR Decomposition	40
2.2.4	Uniqueness of QR Decompositions	41
2.2.5	Final Thoughts	41
2.3	Singular Value Decomposition	41
2.3.1	Matrix-Adjoint Products	42
2.3.2	Singular Value Decomposition	44
2.3.3	Visualizing the SVD	45
2.3.4	Properties of the SVD	46
2.4	Cholesky Decomposition	47
2.4.1	The Cholesky Decomposition	47
2.4.2	Computing a Cholesky Decomposition	47
3	Canonical Forms	49
3.1	Generalized Eigenspaces	49
3.1.1	Kernels of Powers of Linear Transformations	50
3.1.2	Generalized Eigenspaces	52
3.2	Nilpotent Linear Transformations	62

3.2.1	Nilpotent Linear Transformations	62
3.2.2	Powers of Kernels of Nilpotent Linear Transformations	64
3.2.3	Restrictions to Generalized Eigenspaces	66
3.2.4	Jordan Blocks	70
3.3	Jordan Canonical Form	73
3.3.1	Canonical Form for Nilpotent Linear Transformations	73
3.3.2	Restrictions to Generalized Eigenspaces	81
3.3.3	Jordan Canonical Form	81
4	Applications	95
4.1	Least Squares	95
4.1.1	Computing Least-Squares Solutions	99
4.2	Curve Fitting	100
4.2.1	Interpolating Polynomials	100
4.2.2	Fitting Curves Through Data	101
5	Topics	103
5.1	Vandermonde Matrices	103
A	GNU Free Documentation License	107

Chapter 1

The Fundamentals

1.1 Introduction

This book is about advanced topics in linear algebra. So we presume you have some experience with matrix algebra, vector spaces (possibly abstract ones), eigenvalues, linear transformations, and matrix representations of linear transformations. All of this material can be found in *A First Course in Linear Algebra*, which we will reference frequently.

Our approach is mathematical, which means we include proofs of our results. However, we are also practical, and will not always be as general as we could be. For example, we will stick to a single inner product throughout (the sesquilinear one that is most useful when employing complex numbers). We will sometimes be careful about our field of scalars, but will not dwell on the distinctions peculiar to the real numbers (versus the algebraically closed complex numbers). This is not a course in numerical linear algebra, but much of what we do provides the mathematical underpinnings of that topic, so this could be a very useful resource for study in that area. We will make mention of algorithmic performance, relying on Trefethen and Bau's excellent *Numerical Linear Algebra* for details.

Many topics we consider are motivated by trying to find simpler versions of matrices. Here "simpler" can be taken to mean many zero entries. Barring a zero entry, then maybe an entry equal to one is second-best. An overall form that is much like a diagonal matrix is also desirable, since diagonal matrices are simple to work with. (forward referenc eto exercise). A familiar example may help to make these ideas more precise.

Example 1. Given an $m \times n$ matrix, A , we know that its reduced row-echelon form is unique (THEOREM RREFU). We also know that we can accomplish row-operations by multiplying A on the left by a (nonsingular) elementary matrix (SUBSECTION DM.EM). Suppose we perform repeated row-operations to transform A into a matrix in reduced row-echelon form, B . Then the product of the elementary matrices is a square nonsingular matrix, J such that

$$B = JA$$

or equivalently

$$A = J^{-1}B$$

We call the second version a **factorization**, or **matrix decomposition**, of A (Though some might use the same terms for the first version, perhaps saying it is a

factorization of B). The pieces of this decomposition have certain properties. The matrix J^{-1} is a nonsingular matrix of size m . The matrix B has an abundance of zero entries, and some strategically placed “leading ones” which signify the pivot columns. The exact structure of B is described by DEFINITION RREF and THEOREM RREF tells us that we can accomplish this decomposition given any matrix A .

If A is not of full rank, then there are many possibilities for the matrix J , even though B is unique. However, results on extended echelon form (SUBSECTION FS.PEEF suggest a choice for J that is unambiguous. We say that choice is **canonical**. This example gives the following theorem, where we have changed the notation slightly.

Again, many of the topics in this book will have a flavor similar to the previous example and theorem. However, we will often need to limit the possibilities for the original matrix (it may need to be square, or its eigenvalues may need certain properties). We may get more specific information about the components of the factorization, or we may get less. We will also be interested in obtaining **canonical forms** of matrices. You can view orthonormal diagonalization (SECTION OD) as a good example of another matrix decomposition, and we will cover it again in some detail in (cite forward) ??.

1.2 Direct Sums

1.2.1 Direct Sums

Some of the more advanced ideas in linear algebra are closely related to decomposing (PROOF TECHNIQUE DC) vector spaces into direct sums of subspaces. A direct sum is a short-hand way to describe the relationship between a vector space and two, or more, of its subspaces. As we will use it, it is not a way to construct new vector spaces from others.

Definition 1. Suppose that V is a vector space with two subspaces U and W such that for every $\mathbf{v} \in V$,

1. There exists vectors $\mathbf{u} \in U$, $\mathbf{w} \in W$ such that $\mathbf{v} = \mathbf{u} + \mathbf{w}$
2. If $\mathbf{v} = \mathbf{u}_1 + \mathbf{w}_1$ and $\mathbf{v} = \mathbf{u}_2 + \mathbf{w}_2$ where $\mathbf{u}_1, \mathbf{u}_2 \in U$, $\mathbf{w}_1, \mathbf{w}_2 \in W$ then $\mathbf{u}_1 = \mathbf{u}_2$ and $\mathbf{w}_1 = \mathbf{w}_2$.

Then V is the direct sum of U and W and we write $V = U \oplus W$.

Informally, when we say V is the direct sum of the subspaces U and W , we are saying that each vector of V can always be expressed as the sum of a vector from U and a vector from W , and this expression can only be accomplished in one way (i.e. uniquely). This statement should begin to feel something like our definitions of nonsingular matrices (DEFINITION NM) and linear independence (DEFINITION LI). It should not be hard to imagine the natural extension of this definition to the case of more than two subspaces. Could you provide a careful definition of $V = U_1 \oplus U_2 \oplus U_3 \oplus \dots \oplus U_m$ (EXERCISE PD.M50)?

Example 2. In \mathbb{C}^3 , define

$$\mathbf{v}_1 = \begin{bmatrix} 3 \\ 2 \\ 5 \end{bmatrix} \quad \mathbf{v}_2 = \begin{bmatrix} -1 \\ 2 \\ 1 \end{bmatrix} \quad \mathbf{v}_3 = \begin{bmatrix} 2 \\ 1 \\ -2 \end{bmatrix}$$

Then $\mathbb{C}^3 = \langle \{\mathbf{v}_1, \mathbf{v}_2\} \rangle \oplus \langle \{\mathbf{v}_3\} \rangle$. This statement derives from the fact that $B = \{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3\}$ is basis for \mathbb{C}^3 . The spanning property of B yields the decomposition of any vector into a sum of vectors from the two subspaces, and the linear independence of B yields the uniqueness of the decomposition. We will illustrate these claims with a numerical example.

Choose $\mathbf{v} = \begin{bmatrix} 10 \\ 1 \\ 6 \end{bmatrix}$. Then

$$\mathbf{v} = 2\mathbf{v}_1 + (-2)\mathbf{v}_2 + 1\mathbf{v}_3 = (2\mathbf{v}_1 + (-2)\mathbf{v}_2) + (1\mathbf{v}_3)$$

where we have added parentheses for emphasis. Obviously $1\mathbf{v}_3 \in \langle \{\mathbf{v}_3\} \rangle$, while $2\mathbf{v}_1 + (-2)\mathbf{v}_2 \in \langle \{\mathbf{v}_1, \mathbf{v}_2\} \rangle$. THEOREM VRRB provides the uniqueness of the scalars in these linear combinations.

EXAMPLE SDS is easy to generalize into a theorem.

Theorem 1.1. *Suppose that V is a vector space with a basis $B = \{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3, \dots, \mathbf{v}_n\}$ and $m \leq n$. Define*

$$U = \langle \{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3, \dots, \mathbf{v}_m\} \rangle \quad W = \langle \{\mathbf{v}_{m+1}, \mathbf{v}_{m+2}, \mathbf{v}_{m+3}, \dots, \mathbf{v}_n\} \rangle$$

Then $V = U \oplus W$.

Proof. Choose any vector $\mathbf{v} \in V$. Then by THEOREM VRRB there are unique scalars, $a_1, a_2, a_3, \dots, a_n$ such that

$$\begin{aligned} \mathbf{v} &= a_1\mathbf{v}_1 + a_2\mathbf{v}_2 + a_3\mathbf{v}_3 + \dots + a_n\mathbf{v}_n \\ &= (a_1\mathbf{v}_1 + a_2\mathbf{v}_2 + a_3\mathbf{v}_3 + \dots + a_m\mathbf{v}_m) + \\ &= (a_{m+1}\mathbf{v}_{m+1} + a_{m+2}\mathbf{v}_{m+2} + a_{m+3}\mathbf{v}_{m+3} + \dots + a_n\mathbf{v}_n) \\ &= \mathbf{u} + \mathbf{w} \end{aligned}$$

where we have implicitly defined \mathbf{u} and \mathbf{w} in the last line. It should be clear that $\mathbf{u} \in U$, and similarly, $\mathbf{w} \in W$ (and not simply by the choice of their names).

Suppose we had another decomposition of \mathbf{v} , say $\mathbf{v} = \mathbf{u}^* + \mathbf{w}^*$. Then we could write \mathbf{u}^* as a linear combination of \mathbf{v}_1 through \mathbf{v}_m , say using scalars $b_1, b_2, b_3, \dots, b_m$. And we could write \mathbf{w}^* as a linear combination of \mathbf{v}_{m+1} through \mathbf{v}_n , say using scalars $c_1, c_2, c_3, \dots, c_{n-m}$. These two collections of scalars would then together give a linear combination of \mathbf{v}_1 through \mathbf{v}_n that equals \mathbf{v} . By the uniqueness of $a_1, a_2, a_3, \dots, a_n$, $a_i = b_i$ for $1 \leq i \leq m$ and $a_{m+i} = c_i$ for $1 \leq i \leq n - m$. From the equality of these scalars we conclude that $\mathbf{u} = \mathbf{u}^*$ and $\mathbf{w} = \mathbf{w}^*$. So with both conditions of Definition 1 fulfilled we see that $V = U \oplus W$. \square

Given one subspace of a vector space, we can always find another subspace that will pair with the first to form a direct sum. The main idea of this theorem, and its proof, is the idea of extending a linearly independent subset into a basis with repeated applications of THEOREM ELIS.

Theorem 1.2. *Suppose that U is a subspace of the vector space V . Then there exists a subspace W of V such that $V = U \oplus W$.*

Proof. If $U = V$, then choose $W = \{\mathbf{0}\}$. Otherwise, choose a basis $B = \{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3, \dots, \mathbf{v}_m\}$ for U . Then since B is a linearly independent set, THEOREM ELIS tells us there is a vector \mathbf{v}_{m+1} in V , but not in U , such that $B \cup \{\mathbf{v}_{m+1}\}$ is linearly independent. Define the subspace $U_1 = \langle B \cup \{\mathbf{v}_{m+1}\} \rangle$.

We can repeat this procedure, in the case were $U_1 \neq V$, creating a new vector \mathbf{v}_{m+2} in V , but not in U_1 , and a new subspace $U_2 = \langle B \cup \{\mathbf{v}_{m+1}, \mathbf{v}_{m+2}\} \rangle$. If we continue repeating this procedure, eventually, $U_k = V$ for some k , and we can no longer apply THEOREM ELIS. No matter, in this case $B \cup \{\mathbf{v}_{m+1}, \mathbf{v}_{m+2}, \dots, \mathbf{v}_{m+k}\}$ is a linearly independent set that spans V , i.e. a basis for V .

Define $W = \langle \{\mathbf{v}_{m+1}, \mathbf{v}_{m+2}, \dots, \mathbf{v}_{m+k}\} \rangle$. We now are exactly in position to apply Theorem 1.1 and see that $V = U \oplus W$. \square

There are several different ways to define a direct sum. Our next two theorems give equivalences (PROOF TECHNIQUE E) for direct sums, and therefore could have been employed as definitions. The first should further cement the notion that a direct sum has some connection with linear independence.

Theorem 1.3. *Suppose U and W are subspaces of the vector space V . Then $V = U \oplus W$ if and only if*

1. *For every $\mathbf{v} \in V$, there exists vectors $\mathbf{u} \in U$, $\mathbf{w} \in W$ such that $\mathbf{v} = \mathbf{u} + \mathbf{w}$.*
2. *Whenever $\mathbf{0} = \mathbf{u} + \mathbf{w}$ with $\mathbf{u} \in U$, $\mathbf{w} \in W$ then $\mathbf{u} = \mathbf{w} = \mathbf{0}$.*

Proof. The first condition is identical in the definition and the theorem, so we only need to establish the equivalence of the second conditions.

Assume that $V = U \oplus W$, according to Definition 1. By PROPERTY Z, $\mathbf{0} \in V$ and $\mathbf{0} = \mathbf{0} + \mathbf{0}$. If we also assume that $\mathbf{0} = \mathbf{u} + \mathbf{w}$, then the uniqueness of the decomposition gives $\mathbf{u} = \mathbf{0}$ and $\mathbf{w} = \mathbf{0}$.

\Leftarrow Suppose that $\mathbf{v} \in V$, $\mathbf{v} = \mathbf{u}_1 + \mathbf{w}_1$ and $\mathbf{v} = \mathbf{u}_2 + \mathbf{w}_2$ where $\mathbf{u}_1, \mathbf{u}_2 \in U$, $\mathbf{w}_1, \mathbf{w}_2 \in W$. Then

$$\begin{aligned} \mathbf{0} &= \mathbf{v} - \mathbf{v} \\ &= (\mathbf{u}_1 + \mathbf{w}_1) - (\mathbf{u}_2 + \mathbf{w}_2) \\ &= (\mathbf{u}_1 - \mathbf{u}_2) + (\mathbf{w}_1 - \mathbf{w}_2) \end{aligned}$$

By PROPERTY AC, $\mathbf{u}_1 - \mathbf{u}_2 \in U$ and $\mathbf{w}_1 - \mathbf{w}_2 \in W$. We can now apply our hypothesis, the second statement of the theorem, to conclude that

$$\begin{array}{ll} \mathbf{u}_1 - \mathbf{u}_2 = \mathbf{0} & \mathbf{w}_1 - \mathbf{w}_2 = \mathbf{0} \\ \mathbf{u}_1 = \mathbf{u}_2 & \mathbf{w}_1 = \mathbf{w}_2 \end{array}$$

which establishes the uniqueness needed for the second condition of the definition. \square

Our second equivalence lends further credence to calling a direct sum a decomposition. The two subspaces of a direct sum have no (nontrivial) elements in common.

Theorem 1.4. *Suppose U and W are subspaces of the vector space V . Then $V = U \oplus W$ if and only if*

1. *For every $\mathbf{v} \in V$, there exists vectors $\mathbf{u} \in U$, $\mathbf{w} \in W$ such that $\mathbf{v} = \mathbf{u} + \mathbf{w}$.*
2. *$U \cap W = \{\mathbf{0}\}$.*

Proof. The first condition is identical in the definition and the theorem, so we only need to establish the equivalence of the second conditions.

Assume that $V = U \oplus W$, according to Definition 1. By PROPERTY Z and DEFINITION SI, $\{\mathbf{0}\} \subseteq U \cap W$. To establish the opposite inclusion, suppose that

$\mathbf{x} \in U \cap W$. Then, since \mathbf{x} is an element of both U and W , we can write two decompositions of \mathbf{x} as a vector from U plus a vector from W ,

$$\mathbf{x} = \mathbf{x} + \mathbf{0} \qquad \mathbf{x} = \mathbf{0} + \mathbf{x}$$

By the uniqueness of the decomposition, we see (twice) that $\mathbf{x} = \mathbf{0}$ and $U \cap W \subseteq \{\mathbf{0}\}$. Applying DEFINITION SE, we have $U \cap W = \{\mathbf{0}\}$.

\Leftarrow Assume that $U \cap W = \{\mathbf{0}\}$. And assume further that $\mathbf{v} \in V$ is such that $\mathbf{v} = \mathbf{u}_1 + \mathbf{w}_1$ and $\mathbf{v} = \mathbf{u}_2 + \mathbf{w}_2$ where $\mathbf{u}_1, \mathbf{u}_2 \in U$, $\mathbf{w}_1, \mathbf{w}_2 \in W$. Define $\mathbf{x} = \mathbf{u}_1 - \mathbf{u}_2$. then by PROPERTY AC, $\mathbf{x} \in U$. Also

$$\begin{aligned} \mathbf{x} &= \mathbf{u}_1 - \mathbf{u}_2 \\ &= (\mathbf{v} - \mathbf{w}_1) - (\mathbf{v} - \mathbf{w}_2) \\ &= (\mathbf{v} - \mathbf{v}) - (\mathbf{w}_1 - \mathbf{w}_2) \\ &= \mathbf{w}_2 - \mathbf{w}_1 \end{aligned}$$

So $\mathbf{x} \in W$ by PROPERTY AC. Thus, $\mathbf{x} \in U \cap W = \{\mathbf{0}\}$ (DEFINITION SI). So $\mathbf{x} = \mathbf{0}$ and

$$\begin{aligned} \mathbf{u}_1 - \mathbf{u}_2 &= \mathbf{0} & \mathbf{w}_2 - \mathbf{w}_1 &= \mathbf{0} \\ \mathbf{u}_1 &= \mathbf{u}_2 & \mathbf{w}_2 &= \mathbf{w}_1 \end{aligned}$$

yielding the desired uniqueness of the second condition of the definition. \square

If the statement of Theorem 1.3 did not remind you of linear independence, the next theorem should establish the connection.

Theorem 1.5. *Suppose U and W are subspaces of the vector space V with $V = U \oplus W$. Suppose that R is a linearly independent subset of U and S is a linearly independent subset of W . Then $R \cup S$ is a linearly independent subset of V .*

Proof. Let $R = \{\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3, \dots, \mathbf{u}_k\}$ and $S = \{\mathbf{w}_1, \mathbf{w}_2, \mathbf{w}_3, \dots, \mathbf{w}_\ell\}$. Begin with a relation of linear dependence (DEFINITION RLD) on the set $R \cup S$ using scalars $a_1, a_2, a_3, \dots, a_k$ and $b_1, b_2, b_3, \dots, b_\ell$. Then,

$$\begin{aligned} \mathbf{0} &= a_1\mathbf{u}_1 + a_2\mathbf{u}_2 + a_3\mathbf{u}_3 + \dots + a_k\mathbf{u}_k + b_1\mathbf{w}_1 + b_2\mathbf{w}_2 + b_3\mathbf{w}_3 + \dots + b_\ell\mathbf{w}_\ell \\ &= (a_1\mathbf{u}_1 + a_2\mathbf{u}_2 + a_3\mathbf{u}_3 + \dots + a_k\mathbf{u}_k) + (b_1\mathbf{w}_1 + b_2\mathbf{w}_2 + b_3\mathbf{w}_3 + \dots + b_\ell\mathbf{w}_\ell) \\ &= \mathbf{u} + \mathbf{w} \end{aligned}$$

where we have made an implicit definition of the vectors $\mathbf{u} \in U$, $\mathbf{w} \in W$.

Applying Theorem 1.3 we conclude that

$$\begin{aligned} \mathbf{u} &= a_1\mathbf{u}_1 + a_2\mathbf{u}_2 + a_3\mathbf{u}_3 + \dots + a_k\mathbf{u}_k = \mathbf{0} \\ \mathbf{w} &= b_1\mathbf{w}_1 + b_2\mathbf{w}_2 + b_3\mathbf{w}_3 + \dots + b_\ell\mathbf{w}_\ell = \mathbf{0} \end{aligned}$$

Now the linear independence of R and S (individually) yields

$$a_1 = a_2 = a_3 = \dots = a_k = 0 \qquad b_1 = b_2 = b_3 = \dots = b_\ell = 0$$

Forced to acknowledge that only a trivial linear combination yields the zero vector, DEFINITION LI says the set $R \cup S$ is linearly independent in V . \square

Our last theorem in this collection will go some ways towards explaining the word “sum” in the moniker “direct sum”.

Theorem 1.6. *Suppose U and W are subspaces of the vector space V with $V = U \oplus W$. Then $\dim(V) = \dim(U) + \dim(W)$.*

Proof. We will establish this equality of positive integers with two inequalities. We will need a basis of U (call it B) and a basis of W (call it C).

First, note that B and C have sizes equal to the dimensions of the respective subspaces. The union of these two linearly independent sets, $B \cup C$ will be linearly independent in V by Theorem 1.5. Further, the two bases have no vectors in common by Theorem 1.4, since $B \cap C \subseteq \{\mathbf{0}\}$ and the zero vector is never an element of a linearly independent set (EXERCISE LI.T10). So the size of the union is exactly the sum of the dimensions of U and W . By THEOREM G the size of $B \cup C$ cannot exceed the dimension of V without being linearly dependent. These observations give us $\dim(U) + \dim(W) \leq \dim(V)$.

Grab any vector $\mathbf{v} \in V$. Then by Theorem 1.4 we can write $\mathbf{v} = \mathbf{u} + \mathbf{w}$ with $\mathbf{u} \in U$ and $\mathbf{w} \in W$. Individually, we can write \mathbf{u} as a linear combination of the basis elements in B , and similarly, we can write \mathbf{w} as a linear combination of the basis elements in C , since the bases are spanning sets for their respective subspaces. These two sets of scalars will provide a linear combination of all of the vectors in $B \cup C$ which will equal \mathbf{v} . The upshot of this is that $B \cup C$ is a spanning set for V . By THEOREM G, the size of $B \cup C$ cannot be smaller than the dimension of V without failing to span V . These observations give us $\dim(U) + \dim(W) \geq \dim(V)$. \square

There is a certain appealing symmetry in the previous proof, where both linear independence and spanning properties of the bases are used, both of the first two conclusions of THEOREM G are employed, and we have quoted both of the two conditions of Theorem 1.4.

One final theorem tells us that we can successively decompose direct sums into sums of smaller and smaller subspaces.

Theorem 1.7. *Suppose V is a vector space with subspaces U and W with $V = U \oplus W$. Suppose that X and Y are subspaces of W with $W = X \oplus Y$. Then $V = U \oplus X \oplus Y$.*

Proof. Suppose that $\mathbf{v} \in V$. Then due to $V = U \oplus W$, there exist vectors $\mathbf{u} \in U$ and $\mathbf{w} \in W$ such that $\mathbf{v} = \mathbf{u} + \mathbf{w}$. Due to $W = X \oplus Y$, there exist vectors $\mathbf{x} \in X$ and $\mathbf{y} \in Y$ such that $\mathbf{w} = \mathbf{x} + \mathbf{y}$. All together,

$$\mathbf{v} = \mathbf{u} + \mathbf{w} = \mathbf{u} + \mathbf{x} + \mathbf{y}$$

which would be the first condition of a definition of a 3-way direct product.

Now consider the uniqueness. Suppose that

$$\mathbf{v} = \mathbf{u}_1 + \mathbf{x}_1 + \mathbf{y}_1 \qquad \mathbf{v} = \mathbf{u}_2 + \mathbf{x}_2 + \mathbf{y}_2$$

Because $\mathbf{x}_1 + \mathbf{y}_1 \in W$, $\mathbf{x}_2 + \mathbf{y}_2 \in W$, and $V = U \oplus W$, we conclude that

$$\mathbf{u}_1 = \mathbf{u}_2 \qquad \mathbf{x}_1 + \mathbf{y}_1 = \mathbf{x}_2 + \mathbf{y}_2$$

From the second equality, an application of $W = X \oplus Y$ yields the conclusions $\mathbf{x}_1 = \mathbf{x}_2$ and $\mathbf{y}_1 = \mathbf{y}_2$. This establishes the uniqueness of the decomposition of \mathbf{v} into a sum of vectors from U , X and Y . \square

Remember that when we write $V = U \oplus W$ there always needs to be a “super-space,” in this case V . The statement $U \oplus W$ is meaningless. Writing $V = U \oplus W$ is simply a shorthand for a somewhat complicated relationship between V , U and W , as described in the two conditions of Definition 1, or Theorem 1.3, or Theorem 1.4.

Theorem 1.1 and Theorem 1.2 gives us sure-fire ways to build direct sums, while Theorem 1.5, Theorem 1.6 and Theorem 1.7 tell us interesting properties of direct sums.

This subsection has been long on theorems and short on examples. If we were to use the term “lemma” we might have chosen to label some of these results as such, since they will be important tools in other proofs, but may not have much interest on their own (see PROOF TECHNIQUE LC). We will be referencing these results heavily in later sections, and will remind you then to come back for a second look.

1.3 Orthogonal Complements

Theorem $\langle\langle$ above on repeated sums $\rangle\rangle$ mentions repeated sums, which are of interest. However, when we begin with a vector space V and a single subspace W , we can ask about the existence of another subspace, W , such that $V = U \oplus W$. The answer is that such a W always exists, and we then refer to it as a **complement** of U .

Definition 2. Suppose that V is a vector space with a subspace U . If W is a subspace such that $V = U \oplus W$, then W is the complement of U .

Every subspace has a complement, and generally it is not unique.

Lemma 1.8. *Suppose that V is a vector space with a subspace U . Then there exists a subspace W such that $V = U \oplus W$, so W is the complement of U .*

Proof. Suppose that $\dim(V) = n$ and $\dim(U) = k$, and let $B = \{\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3, \dots, \mathbf{u}_k\}$ be a basis of U . With $n - k$ applications of THEOREM ELIS we obtain vectors $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3, \dots, \mathbf{v}_{n-k}$ that successively create bases $B_i = \{\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3, \dots, \mathbf{u}_k, \mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3, \dots, \mathbf{v}_i\}$ $0 \leq i \leq n - k$ for subspaces $U = U_0, U_1, \dots, U_{n-k} = V$, where $\dim(U_i) = k + i$.

Define $W = \langle\langle\{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3, \dots, \mathbf{v}_{n-k}\}\rangle\rangle$. Since $\{\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3, \dots, \mathbf{u}_k, \mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3, \dots, \mathbf{v}_i\}$ is a basis for V and $\{\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3, \dots, \mathbf{u}_k\}$ is a basis for U , we can apply Theorem $\langle\langle$ Direct Sum From a Basis (above) $\rangle\rangle$ to see that $V = U \oplus W$, so W is the complement of U . (Compare with $\langle\langle$ Direct Sum From One Subspace (above) $\rangle\rangle$, which has identical content, but a different write-up.) \square

The freedom given when we “extend” a linearly independent set (or basis) to create a basis for the complement means that we can create a complement in many ways, so it is not unique.

Exercise 1. Consider the subspace U of $V = \mathbb{C}^3$,

$$U = \left\langle \left\{ \begin{bmatrix} 1 \\ -6 \\ -8 \end{bmatrix}, \begin{bmatrix} 1 \\ -5 \\ -7 \end{bmatrix} \right\} \right\rangle.$$

Create two different complements of U , being sure to *prove* that your complements are unequal (and not simply have unequal bases). Before reading ahead, can you think of an ideal (or “canonical”) choice for the complement?

Exercise 2. Consider the subspace U of $V = \mathbb{C}^5$,

$$U = \left\langle \left\{ \begin{bmatrix} 1 \\ -4 \\ -2 \\ 6 \\ -5 \end{bmatrix}, \begin{bmatrix} 1 \\ -4 \\ -1 \\ 4 \\ -3 \end{bmatrix} \right\} \right\rangle.$$

Create a complement of U . (If you have read ahead, do not create an orthogonal complement for this exercise.)

With an inner product, and a notion of orthogonality, we can define a canonical, and useful, complement for every subspace.

Definition 3. Suppose that V is a vector space with a subspace U . Then the orthogonal complement of U (relative to V) is

$$U^\perp = \{\mathbf{v} \in V \mid \langle \mathbf{v}, \mathbf{u} \rangle = 0 \text{ for every } \mathbf{u} \in U\}.$$

A matrix formulation of the orthogonal complement will help us establish that the moniker “complement” is deserved.

Theorem 1.9. Suppose that V is a vector space with a subspace U . Let A be a matrix whose columns are a spanning set for U . Then $U^\perp = \mathcal{N}(A^*)$.

Proof. Membership in the orthogonal complement requires a vector to be orthogonal to every vector of U . However, because of the linearity of the inner product (THEOREM IPVA, THEOREM IPSM), it is equivalent to require that a vector be orthogonal to each member of a spanning set for U . So membership in the orthogonal complement is equivalent to being orthogonal to each column of A . We obtain the desired set equality from the equivalences,

$$\mathbf{v} \in U^\perp \iff \mathbf{v}^* A = \mathbf{0}^* \iff A^* \mathbf{v} = \mathbf{0} \iff \mathbf{v} \in \mathcal{N}(A^*).$$

□

Theorem 1.10. Suppose that V is a vector space with a subspace U . Then $V = U \oplus U^\perp$.

Proof. We first establish that $U \cap U^\perp = \{\mathbf{0}\}$. Suppose $\mathbf{u} \in U$ and $\mathbf{u} \in U^\perp$. Then $\langle \mathbf{u}, \mathbf{u} \rangle = 0$ and by THEOREM PIP we conclude that $\mathbf{u} = \mathbf{0}$.

We now show that an arbitrary vector \mathbf{v} can be written as a sum of vectors from U and U^\perp . Without loss of generality, we can assume we have an orthonormal basis for U , for if not, we can apply the Gram-Schmidt process to any basis of U to create an orthogonal spanning set, whose individual vectors can be scaled to have norm one (THEOREM GSP). Denote this basis as $B = \{\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3, \dots, \mathbf{u}_k\}$.

Define the vector \mathbf{v}_1 as a linear combination of the vectors of B , so $\mathbf{v}_1 \in U$.

$$\mathbf{v}_1 = \sum_{i=1}^k \langle \mathbf{u}_i, \mathbf{v} \rangle \mathbf{u}_i$$

. Define $\mathbf{v}_2 = \mathbf{v} - \mathbf{v}_1$, so trivially by construction, $\mathbf{v} = \mathbf{v}_1 + \mathbf{v}_2$. It remains to show that $\mathbf{v}_2 \in U^\perp$. We repeatedly use properties of the inner product. This construction and proof may remind you of the Gram-Schmidt process. For $1 \leq j \leq k$,

$$\begin{aligned} \langle \mathbf{v}_2, \mathbf{u}_j \rangle &= \langle \mathbf{v}, \mathbf{u}_j \rangle - \langle \mathbf{v}_1, \mathbf{u}_j \rangle \\ &= \langle \mathbf{v}, \mathbf{u}_j \rangle - \sum_{i=1}^k \langle \langle \mathbf{u}_i, \mathbf{v} \rangle \mathbf{u}_i, \mathbf{u}_j \rangle \\ &= \langle \mathbf{v}, \mathbf{u}_j \rangle - \sum_{i=1}^k \overline{\langle \mathbf{u}_i, \mathbf{v} \rangle} \langle \mathbf{u}_i, \mathbf{u}_j \rangle \\ &= \langle \mathbf{v}, \mathbf{u}_j \rangle - \overline{\langle \mathbf{u}_j, \mathbf{v} \rangle} \langle \mathbf{u}_j, \mathbf{u}_j \rangle \\ &= \langle \mathbf{v}, \mathbf{u}_j \rangle - \langle \mathbf{v}, \mathbf{u}_j \rangle \\ &= 0 \end{aligned}$$

We have fulfilled the hypotheses of 1.3 and so can say $V = U \oplus U^\perp$.

□

Theorem 1.10 gives us a canonical choice of a complementary subspace, which has useful orthogonality properties. It also allows us to decompose any vector (uniquely) into an element of a subspace, plus an orthogonal vector. This might remind you in some ways of “resolving a vector into components” if you have studied physics some.

Given a matrix, we get a natural vector space decomposition.

Corollary 1.11. *Suppose that A is an $m \times n$ matrix. Then*

$$\mathbb{C}^m = \mathcal{C}(A) \oplus \mathcal{C}(A)^\perp = \mathcal{C}(A) \oplus \mathcal{N}(A^*).$$

Proof. Theorem 1.10 provides the first equality and Theorem 1.9 gives the second. \square

Exercise 3. Compute the orthogonal complement of the subspace $U \subset \mathbb{C}^3$.

$$U = \left\langle \left\{ \begin{bmatrix} 1 \\ -1 \\ 5 \end{bmatrix}, \begin{bmatrix} 3 \\ 1 \\ 3 \end{bmatrix} \right\} \right\rangle$$

Solution. Form the matrix A , whose columns are the two basis vectors given for U and compute the null space $\mathcal{N}(A^*)$ by row-reducing the matrix. (Theorem 1.9)

$$A^* = \begin{bmatrix} 1 & -1 & 5 \\ 3 & 1 & 3 \end{bmatrix} \xrightarrow{\text{RREF}} \begin{bmatrix} 1 & 0 & 2 \\ 0 & 1 & -3 \end{bmatrix}$$

So

$$U^\perp = \mathcal{N}(A^*) = \left\langle \left\{ \begin{bmatrix} -2 \\ 3 \\ 1 \end{bmatrix} \right\} \right\rangle$$

Exercise 4. Compute the orthogonal complements of the two subspaces from Exercises 1 and 2. For the subspace of \mathbb{C}^5 verify that your first complement was not the orthogonal complement (or return to the exercise and find a complement that is not orthogonal).

1.4 Invariant Subspaces

1.4.1 Invariant Subspaces

Definition 4. Suppose that $T: V \rightarrow V$ is a linear transformation and W is a subspace of V . Suppose further that $T(\mathbf{w}) \in W$ for every $\mathbf{w} \in W$. Then W is an **invariant subspace** of V relative to T .

We do not have any special notation for an invariant subspace, so it is important to recognize that an invariant subspace is always relative to both a superspace (V) and a linear transformation (T), which will sometimes not be mentioned, yet will be clear from the context. Note also that the linear transformation involved must have an equal domain and codomain—the definition would not make much sense if our outputs were not of the same type as our inputs.

As is our habit, we begin with an example that demonstrates the existence of invariant subspaces, while leaving other questions unanswered for the moment. We will return later to understand how this example was constructed, but for now, just understand how we check the existence of the invariant subspaces.

Example 3. Consider the linear transformation $T: \mathbb{C}^4 \rightarrow \mathbb{C}^4$ defined by $T(\mathbf{x}) = A\mathbf{x}$ where A is given by

$$A = \begin{bmatrix} -8 & 6 & -15 & 9 \\ -8 & 14 & -10 & 18 \\ 1 & 1 & 3 & 0 \\ 3 & -8 & 2 & -11 \end{bmatrix}$$

Define (with zero motivation),

$$\mathbf{w}_1 = \begin{bmatrix} -7 \\ -2 \\ 3 \\ 0 \end{bmatrix} \qquad \mathbf{w}_2 = \begin{bmatrix} -1 \\ -2 \\ 0 \\ 1 \end{bmatrix}$$

and set $W = \langle \{\mathbf{w}_1, \mathbf{w}_2\} \rangle$. We verify that W is an invariant subspace of \mathbb{C}^4 with respect to T . By the definition of W , any vector chosen from W can be written as a linear combination of \mathbf{w}_1 and \mathbf{w}_2 . Suppose that $\mathbf{w} \in W$, and then check the details of the following verification,

$$\begin{aligned} T(\mathbf{w}) &= T(a_1\mathbf{w}_1 + a_2\mathbf{w}_2) \\ &= a_1T(\mathbf{w}_1) + a_2T(\mathbf{w}_2) \\ &= a_1 \begin{bmatrix} -1 \\ -2 \\ 0 \\ 1 \end{bmatrix} + a_2 \begin{bmatrix} 5 \\ -2 \\ -3 \\ 2 \end{bmatrix} \\ &= a_1\mathbf{w}_2 + a_2((-1)\mathbf{w}_1 + 2\mathbf{w}_2) \\ &= (-a_2)\mathbf{w}_1 + (a_1 + 2a_2)\mathbf{w}_2 \in W \end{aligned}$$

So, by DEFINITION IS, W is an invariant subspace of \mathbb{C}^4 relative to T .

In an entirely similar manner we construct another invariant subspace of T . With zero motivation, define

$$\mathbf{x}_1 = \begin{bmatrix} -3 \\ -1 \\ 1 \\ 0 \end{bmatrix} \qquad \mathbf{w}_2 = \begin{bmatrix} -1 \\ -2 \\ 0 \\ 1 \end{bmatrix}$$

and set $W = \langle \{\mathbf{w}_1, \mathbf{w}_2\} \rangle$. We verify that W is an invariant subspace of \mathbb{C}^4 with respect to T . By the definition of W , any vector chosen from W can be written as a linear combination of \mathbf{w}_1 and \mathbf{w}_2 . Suppose that $\mathbf{w} \in W$, and then check the details of the following verification,

$$\begin{aligned} T(\mathbf{w}) &= T(a_1\mathbf{w}_1 + a_2\mathbf{w}_2) \\ &= a_1T(\mathbf{w}_1) + a_2T(\mathbf{w}_2) \\ &= a_1 \begin{bmatrix} -1 \\ -2 \\ 0 \\ 1 \end{bmatrix} + a_2 \begin{bmatrix} 5 \\ -2 \\ -3 \\ 2 \end{bmatrix} \\ &= a_1\mathbf{w}_2 + a_2((-1)\mathbf{w}_1 + 2\mathbf{w}_2) \\ &= (-a_2)\mathbf{w}_1 + (a_1 + 2a_2)\mathbf{w}_2 \in W \end{aligned}$$

So, by DEFINITION IS, W is an invariant subspace of \mathbb{C}^4 relative to T .

In an entirely similar manner we construct another invariant subspace of T . With zero motivation, define

$$\mathbf{x}_1 = \begin{bmatrix} -3 \\ -1 \\ 1 \\ 0 \end{bmatrix} \qquad \mathbf{x}_2 = \begin{bmatrix} 0 \\ -1 \\ 0 \\ 1 \end{bmatrix}$$

and set $X = \langle \{\mathbf{x}_1, \mathbf{x}_2\} \rangle$. We verify that X is an invariant subspace of \mathbb{C}^4 with respect to T . By the definition of X , any vector chosen from X can be written as a linear combination of \mathbf{x}_1 and \mathbf{x}_2 . Suppose that $\mathbf{x} \in X$, and then check the details of the following verification,

$$\begin{aligned} T(\mathbf{x}) &= T(b_1\mathbf{x}_1 + b_2\mathbf{x}_2) \\ &= b_1T(\mathbf{x}_1) + b_2T(\mathbf{x}_2) \\ &= b_1 \begin{bmatrix} 3 \\ 0 \\ -1 \\ 1 \end{bmatrix} + b_2 \begin{bmatrix} 3 \\ 4 \\ -1 \\ -3 \end{bmatrix} \\ &= b_1((-1)\mathbf{x}_1 + \mathbf{x}_2) + b_2((-1)\mathbf{x}_1 + (-3)\mathbf{x}_2) \\ &= (-b_1 - b_2)\mathbf{x}_1 + (b_1 - 3b_2)\mathbf{x}_2 \in X \end{aligned}$$

So, by DEFINITION IS, X is an invariant subspace of \mathbb{C}^4 relative to T .

There is a bit of magic in each of these verifications where the two outputs of T happen to equal linear combinations of the two inputs. But this is the essential nature of an invariant subspace. We'll have a peek under the hood later, and it won't look so magical after all.

Verify that $B = \{\mathbf{w}_1, \mathbf{w}_2, \mathbf{x}_1, \mathbf{x}_2\}$ is linearly independent, and hence a basis of \mathbb{C}^4 . Splitting this basis in half, Theorem 1.1 tells us that $\mathbb{C}^4 = W \oplus X$. To see exactly why a decomposition of a vector space into a direct sum of invariant subspaces might be interesting work Exercise 5 now.

Exercise 5. Construct a matrix representation of the linear transformation T of Exercise 3 relative to the basis formed as the union of the bases of the two invariant subspaces, $M_{B,B}^T$. Comment on your observations, perhaps after computing a few powers of the matrix representation (which represent repeated compositions of T with itself). Hmmmmm.

Solution. Our basis is

$$B = \{\mathbf{w}_1, \mathbf{w}_2, \mathbf{x}_1, \mathbf{x}_2\} = \left\{ \begin{bmatrix} -7 \\ -2 \\ 3 \\ 0 \end{bmatrix}, \begin{bmatrix} -1 \\ -2 \\ 0 \\ 1 \end{bmatrix}, \begin{bmatrix} -3 \\ -1 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ -1 \\ 0 \\ 1 \end{bmatrix} \right\}$$

Now we perform the necessary computations for the matrix representation of T

relative to B

$$\begin{aligned}\rho_B(T(\mathbf{w}_1)) &= \rho_B \left(\begin{bmatrix} -1 \\ -2 \\ 0 \\ 1 \end{bmatrix} \right) = \rho_B((0)\mathbf{w}_1 + (1)\mathbf{w}_2) = \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix} \\ \rho_B(T(\mathbf{w}_2)) &= \rho_B \left(\begin{bmatrix} 5 \\ -2 \\ -3 \\ 2 \end{bmatrix} \right) = \rho_B((-1)\mathbf{w}_1 + (2)\mathbf{w}_2) = \begin{bmatrix} -1 \\ 2 \\ 0 \\ 0 \end{bmatrix} \\ \rho_B(T(\mathbf{x}_1)) &= \rho_B \left(\begin{bmatrix} 3 \\ 0 \\ -1 \\ 1 \end{bmatrix} \right) = \rho_B((-1)\mathbf{x}_1 + (1)\mathbf{x}_2) = \begin{bmatrix} 0 \\ 0 \\ -1 \\ 1 \end{bmatrix} \\ \rho_B(T(\mathbf{x}_2)) &= \rho_B \left(\begin{bmatrix} 3 \\ 4 \\ -1 \\ -3 \end{bmatrix} \right) = \rho_B((-1)\mathbf{x}_1 + (-3)\mathbf{x}_2) = \begin{bmatrix} 0 \\ 0 \\ -1 \\ -3 \end{bmatrix}\end{aligned}$$

Applying DEFINITION MR, we have

$$M_{B,B}^T = \begin{bmatrix} 0 & -1 & 0 & 0 \\ 1 & 2 & 0 & 0 \\ 0 & 0 & -1 & -1 \\ 0 & 0 & 1 & -3 \end{bmatrix}$$

The interesting feature of this representation is the two 2×2 blocks on the diagonal that arise from the decomposition of \mathbb{C}^4 into a direct sum of invariant subspaces. Or maybe the interesting feature of this matrix is the two 2×2 submatrices in the “other” corners that are all zero. You can decide.

Exercise 6. Prove that the subspaces $U, V \subseteq \mathbb{C}^5$ are invariant with respect to the linear transformation $R: \mathbb{C}^5 \rightarrow \mathbb{C}^5$ defined by $R(\mathbf{x}) = B\mathbf{x}$.

$$B = \begin{bmatrix} 4 & 47 & 3 & -46 & 20 \\ 10 & 61 & 8 & -56 & 10 \\ -10 & -69 & -7 & 67 & -20 \\ 11 & 70 & 9 & -64 & 12 \\ 3 & 19 & 3 & -16 & 1 \end{bmatrix}$$

$$U = \left\langle \left\{ \begin{bmatrix} 1 \\ 0 \\ -1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ -1 \\ 1 \\ 0 \end{bmatrix} \right\} \right\rangle \quad V = \left\langle \left\{ \begin{bmatrix} 1 \\ 1 \\ -1 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 3 \\ 3 \\ -2 \\ 4 \\ 2 \end{bmatrix}, \begin{bmatrix} -2 \\ 3 \\ -2 \\ 3 \\ 1 \end{bmatrix} \right\} \right\rangle$$

```
B = matrix(QQ, [[4, 47, 3, -46, 20],
                [10, 61, 8, -56, 10],
                [-10, -69, -7, 67, -20],
                [11, 70, 9, -64, 12],
                [3, 19, 3, -16, 1]])
```

Prove that the union of U and V is a basis of \mathbb{C}^5 , and then provide a matrix representation of R relative to this basis.

Example 3 and Exercise 6 are a bit mysterious at this stage. Do we know any other examples of invariant subspaces? Yes, as it turns out, we have already seen quite a few. We will give some specific examples, and for more general situations, describe broad classes of invariant subspaces by theorems. First up is eigenspaces.

Theorem 1.12. *Suppose that $T: V \rightarrow V$ is a linear transformation with eigenvalue λ and associated eigenspace $\mathcal{E}_T(\lambda)$. Let W be any subspace of $\mathcal{E}_T(\lambda)$. Then W is an invariant subspace of V relative to T .*

Proof. Choose $\mathbf{w} \in W$. Then

$$T(\mathbf{w}) = \lambda\mathbf{w} \in W.$$

So by Definition 4, W is an invariant subspace of V relative to T . \square

Theorem 1.12 is general enough to determine that an entire eigenspace is an invariant subspace, or that simply the span of a single eigenvector is an invariant subspace. It is not always the case that any subspace of an invariant subspace is again an invariant subspace, but eigenspaces do have this property. Here is an example of the theorem, which also allows us to very quickly build several invariant subspaces.

Example 4. Define the linear transformation $S: M_{22} \rightarrow M_{22}$ by

$$S \left(\begin{bmatrix} a & b \\ c & d \end{bmatrix} \right) = \begin{bmatrix} -2a + 19b - 33c + 21d & -3a + 16b - 24c + 15d \\ -2a + 9b - 13c + 9d & -a + 4b - 6c + 5d \end{bmatrix}$$

Build a matrix representation of S relative to the standard basis (DEFINITION MR) and compute eigenvalues and eigenspaces of S with the computational techniques of CHAPTER E in concert with THEOREM EER. Then

$$\mathcal{E}_S(1) = \left\langle \left\{ \begin{bmatrix} 4 & 3 \\ 2 & 1 \end{bmatrix} \right\} \right\rangle \quad \mathcal{E}_S(2) = \left\langle \left\{ \begin{bmatrix} 6 & 3 \\ 1 & 0 \end{bmatrix}, \begin{bmatrix} -9 & -3 \\ 0 & 1 \end{bmatrix} \right\} \right\rangle$$

So by Theorem 1.12, both $\mathcal{E}_S(1)$ and $\mathcal{E}_S(2)$ are invariant subspaces of M_{22} relative to S .

However, Theorem 1.12 provides even more invariant subspaces. Since $\mathcal{E}_S(1)$ has dimension 1, it has no interesting subspaces, however $\mathcal{E}_S(2)$ has dimension 2 and has a plethora of subspaces. For example, set

$$\mathbf{u} = 2 \begin{bmatrix} 6 & 3 \\ 1 & 0 \end{bmatrix} + 3 \begin{bmatrix} -9 & -3 \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} -6 & -3 \\ 2 & 3 \end{bmatrix}$$

and define $U = \langle \{\mathbf{u}\} \rangle$. Then since U is a subspace of $\mathcal{E}_S(2)$, Theorem 1.12 says that U is an invariant subspace of M_{22} (or we could check this claim directly based simply on the fact that \mathbf{u} is an eigenvector of S).

For every linear transformation there are some obvious, trivial invariant subspaces. Suppose that $T: V \rightarrow V$ is a linear transformation. Then simply because T is a function, the subspace V is an invariant subspace of T . In only a minor twist on this theme, the range of T , $\mathcal{R}(T)$, is an invariant subspace of T by DEFINITION RLT. Finally, THEOREM LTTZZ provides the justification for claiming that $\{\mathbf{0}\}$ is an invariant subspace of T .

That the trivial subspace is always an invariant subspace is a special case of the next theorem. But first, work the following straightforward exercise *before* reading the next theorem. We'll wait.

Exercise 7. Prove that the kernel of a linear transformation (DEFINITION KLT), $\mathcal{K}(T)$, is an invariant subspace.

Theorem 1.13. Suppose that $T: V \rightarrow V$ is a linear transformation and k is a non-negative integer. Then $\mathcal{K}(T^k)$ is an invariant subspace of V .

Proof. Suppose that $\mathbf{z} \in \mathcal{K}(T^k)$. Then

$$T^k(T(\mathbf{z})) = T^{k+1}(\mathbf{z}) = T(T^k(\mathbf{z})) = T(\mathbf{0}) = \mathbf{0}.$$

So by DEFINITION KLT, we see that $T(\mathbf{z}) \in \mathcal{K}(T^k)$. Thus $\mathcal{K}(T^k)$ is an invariant subspace of V relative to T by Definition 4. \square

Two special cases of Theorem 1.13 occur when we choose $k = 0$ and $k = 1$. The next example is unusual, but a good illustration.

Example 5. Consider the 10×10 matrix A below as defining a linear transformation $T: \mathbb{C}^{10} \rightarrow \mathbb{C}^{10}$. We also provide a Sage version of the matrix for use online.

$$A = \begin{bmatrix} -1 & -9 & -24 & -16 & -40 & -36 & 72 & 66 & 21 & 59 \\ 19 & 4 & 7 & -18 & 2 & -12 & 67 & 69 & 16 & -35 \\ -1 & 1 & 2 & 5 & 5 & 6 & -17 & -17 & -5 & -8 \\ 2 & -2 & -7 & -8 & -13 & -13 & 32 & 28 & 11 & 14 \\ -11 & -2 & -1 & 12 & 6 & 11 & -50 & -44 & -16 & 13 \\ 4 & -1 & -5 & -14 & -16 & -16 & 55 & 43 & 24 & 17 \\ -14 & 1 & 7 & 20 & 19 & 26 & -82 & -79 & -21 & -1 \\ 12 & 0 & -4 & -17 & -14 & -20 & 68 & 64 & 19 & -2 \\ 10 & -2 & -9 & -16 & -20 & -24 & 68 & 65 & 17 & 9 \\ -1 & -2 & -5 & -3 & -8 & -7 & 13 & 12 & 4 & 14 \end{bmatrix}$$

```
matrix(QQ, [[-1, -9, -24, -16, -40, -36, 72, 66, 21, 59],
            [19, 4, 7, -18, 2, -12, 67, 69, 16, -35],
            [-1, 1, 2, 5, 5, 6, -17, -17, -5, -8],
            [2, -2, -7, -8, -13, -13, 32, 28, 11, 14],
            [-11, -2, -1, 12, 6, 11, -50, -44, -16, 13],
            [4, -1, -5, -14, -16, -16, 55, 43, 24, 17],
            [-14, 1, 7, 20, 19, 26, -82, -79, -21, -1],
            [12, 0, -4, -17, -14, -20, 68, 64, 19, -2],
            [10, -2, -9, -16, -20, -24, 68, 65, 17, 9],
            [-1, -2, -5, -3, -8, -7, 13, 12, 4, 14]
            ])
```

The matrix A has rank 9 and so T has a nontrivial kernel. But it gets better. T has been constructed specially so that the nullity of T^k is exactly k , for $0 \leq k \leq 10$. This is an extremely unusual situation, but is a good illustration of a very general theorem about kernels of null spaces, coming next. We compute the invariant subspace $\mathcal{K}(T^5)$, you can practice with others.

We work with the matrix, recalling that null spaces and column spaces of matrices correspond to kernels and ranges of linear transformations once we understand

matrix representations of linear transformations (SECTION MR).

$$A^5 = \begin{bmatrix} 37 & 24 & 65 & -35 & 77 & 32 & 80 & 98 & 23 & -125 \\ 19 & 11 & 37 & -21 & 46 & 19 & 29 & 49 & 6 & -70 \\ -7 & -5 & -15 & 8 & -18 & -8 & -15 & -19 & -6 & 26 \\ 14 & 9 & 27 & -15 & 33 & 14 & 26 & 37 & 7 & -50 \\ -8 & -7 & -25 & 14 & -33 & -16 & -10 & -23 & -5 & 37 \\ 12 & 11 & 35 & -19 & 45 & 22 & 22 & 35 & 11 & -52 \\ -27 & -18 & -56 & 31 & -69 & -30 & -49 & -72 & -15 & 100 \\ 20 & 14 & 45 & -25 & 56 & 25 & 35 & 54 & 12 & -77 \\ 24 & 16 & 49 & -27 & 60 & 26 & 45 & 64 & 14 & -88 \\ 8 & 5 & 13 & -7 & 15 & 6 & 18 & 21 & 5 & -26 \end{bmatrix}$$

$$\mathcal{K}(T^5) = \mathcal{N}(A^5) = \left\langle \begin{bmatrix} 1 \\ -1 \\ 0 \\ 0 \\ -1 \\ 2 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} -1 \\ -3 \\ -3 \\ -2 \\ 2 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} -2 \\ -1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ -1 \\ -4 \\ -2 \\ 2 \\ 0 \\ 0 \\ 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 5 \\ -3 \\ 1 \\ 1 \\ 1 \\ 0 \\ 0 \\ 0 \\ 1 \\ 2 \end{bmatrix} \right\rangle$$

As an illustration of $\mathcal{K}(T^5)$ as a subspace invariant under T , we form a linear combination of the five basis vectors (named \mathbf{z}_i , in order), which will be then be an element of the invariant subspace. We apply T , so the output should again be in the subspace, which we verify by giving an expression for the output as a linear combination of the basis vectors for the subspace.

$$\mathbf{z} = 3\mathbf{z}_1 - \mathbf{z}_2 + 3\mathbf{z}_3 + 2\mathbf{z}_4 - 2\mathbf{z}_5 = \begin{bmatrix} -10 \\ 1 \\ -9 \\ -6 \\ -3 \\ 6 \\ -1 \\ 3 \\ 2 \\ -4 \end{bmatrix}$$

$$T(\mathbf{z}) = A\mathbf{z} = \begin{bmatrix} 149 \\ 93 \\ -28 \\ 68 \\ -73 \\ 94 \\ -136 \\ 110 \\ 116 \\ 28 \end{bmatrix}$$

$$= 47\mathbf{z}_1 - 136\mathbf{z}_2 + 110\mathbf{z}_3 + 116\mathbf{z}_4 + 14\mathbf{z}_5$$

Exercise 8. Reprise Example 5 using the same linear transformation. Use a different power (not $k = 0, 1, 5, 9, 10$ on your first attempt), form a vector in the kernel

of your chosen power, then apply T to it. Your output should be in the kernel. (Check this with Sage by using the `in` Python operator.) Thus, you should be able to write the output as a linear combination of the basis vectors. Rinse, repeat.

1.4.2 Restrictions of Linear Transformations

A primary reason for our interest in invariant subspaces is they provide us with another method for creating new linear transformations from old ones.

Definition 5. Suppose that $T: V \rightarrow V$ is a linear transformation, and U is an invariant subspace of V relative to T . Define the **restriction** of T to U by

$$T|_U: U \rightarrow U \qquad T|_U(\mathbf{u}) = T(\mathbf{u})$$

It might appear that this definition has not accomplished anything, as $T|_U$ would appear to take on exactly the same values as T . And this is true. However, $T|_U$ differs from T in the choice of domain and codomain. We tend to give little attention to the domain and codomain of functions, while their defining rules get the spotlight. But the restriction of a linear transformation is all about the choice of domain and codomain. We are *restricting* the rule of the function to a smaller subspace. Notice the importance of only using this construction with an invariant subspace, since otherwise we cannot be assured that the outputs of the function are even contained in the codomain. This observation should be the key step in the proof of a theorem saying that $T|_U$ is also a linear transformation, but leave that as an exercise.

Example 6. In Exercise 6 you verified that the subspaces $U, V \subseteq \mathbb{C}^5$ are invariant with respect to the linear transformation $R: \mathbb{C}^5 \rightarrow \mathbb{C}^5$ defined by $R(\mathbf{x}) = B\mathbf{x}$.

$$B = \begin{bmatrix} 4 & 47 & 3 & -46 & 20 \\ 10 & 61 & 8 & -56 & 10 \\ -10 & -69 & -7 & 67 & -20 \\ 11 & 70 & 9 & -64 & 12 \\ 3 & 19 & 3 & -16 & 1 \end{bmatrix}$$

$$U = \left\langle \left\{ \begin{bmatrix} 1 \\ 0 \\ -1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ -1 \\ 1 \\ 0 \end{bmatrix} \right\} \right\rangle \qquad V = \left\langle \left\{ \begin{bmatrix} 1 \\ 1 \\ -1 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 3 \\ 3 \\ -2 \\ 4 \\ 2 \end{bmatrix}, \begin{bmatrix} -2 \\ 3 \\ -2 \\ 3 \\ 1 \end{bmatrix} \right\} \right\rangle$$

It is a simple matter to define two new linear transformations, $R|_U, R|_V$,

$$\begin{aligned} R|_U: U &\rightarrow U & R|_U(\mathbf{x}) &= B\mathbf{x} \\ R|_V: V &\rightarrow V & R|_V(\mathbf{x}) &= B\mathbf{x} \end{aligned}$$

It should not look like we have accomplished much. Worse, it might appear that $R = R|_U = R|_V$ since each is described by the same rule for computing the image of an input. The difference is that the domains are all different: \mathbb{C}^5, U, V . Since U and V are invariant subspaces, we can then use these subspaces for the codomains of the restrictions.

We will frequently need the matrix representations of linear transformation restrictions, so let's compute those now for this example. Let

$$C = \{\mathbf{u}_1, \mathbf{u}_2\} \qquad D = \{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3\}$$

be the bases for U and V , respectively.

For U

$$\rho_C(R|_U(\mathbf{u}_1)) = \rho_C \left(\begin{bmatrix} 1 \\ 2 \\ -3 \\ 2 \\ 0 \end{bmatrix} \right) = \rho_C(1\mathbf{u}_1 + 2\mathbf{u}_2) = \begin{bmatrix} 1 \\ 2 \end{bmatrix}$$

$$\rho_C(R|_U(\mathbf{u}_2)) = \rho_C \left(\begin{bmatrix} -2 \\ -3 \\ 5 \\ -3 \\ 0 \end{bmatrix} \right) = \rho_C(-2\mathbf{u}_1 - 3\mathbf{u}_2) = \begin{bmatrix} -2 \\ -3 \end{bmatrix}$$

Applying DEFINITION MR, we have

$$M_{C,C}^{R|_U} = \begin{bmatrix} 1 & -2 \\ 2 & -3 \end{bmatrix}$$

For V

$$\rho_D(R|_V(\mathbf{v}_1)) = \rho_D([2, 7, -5, 8, 3]) = \rho_D(\mathbf{v}_1 + \mathbf{v}_2 + \mathbf{v}_3) = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}$$

$$\rho_D(R|_V(\mathbf{v}_2)) = \rho_D([3, -7, 5, -7, -2]) = \rho_D(-\mathbf{v}_1 + 2\mathbf{v}_3) = \begin{bmatrix} -1 \\ 0 \\ 2 \end{bmatrix}$$

$$\rho_D(R|_V(\mathbf{v}_3)) = \rho_D([9, -11, 8, -10, -2]) = \rho_D(-2\mathbf{v}_1 + \mathbf{v}_2 + 4\mathbf{v}_3) = \begin{bmatrix} -2 \\ 1 \\ 4 \end{bmatrix}$$

Applying DEFINITION MR, we have

$$M_{D,D}^{R|_V} = \begin{bmatrix} 1 & -1 & -2 \\ 1 & 0 & 1 \\ 1 & 2 & 4 \end{bmatrix}$$

It is no accident that these two square matrices are the diagonal blocks of the matrix representation you built for R relative to the basis $C \cup D$ in Exercise 6.

The key observation of the previous example is worth stating very clearly: A basis derived from a direct sum decomposition into subspaces that are invariant relative to a linear transformation will provide a matrix representation of the linear transformation with a block diagonal form.

Diagonalizing a linear transformation is the most extreme example of decomposing a vector space into invariant subspaces. When a linear transformation is diagonalizable, then there is a basis composed of eigenvectors (THEOREM DC). Each of these basis vectors can be used individually as the lone element of a basis for an invariant subspace (THEOREM EIS). So the domain decomposes into a direct sum of one-dimensional invariant subspaces via Theorem 1.1. The corresponding matrix representation is then block diagonal with all the blocks of size 1, i.e. the matrix is diagonal. Section [⟨section-jordan-canonical-form⟩](#) is devoted to generalizing this extreme situation when there are not enough eigenvectors available to make such a complete decomposition and arrive at such an elegant matrix representation.

You can find more examples of invariant subspaces, linear transformation restrictions and matrix representations in Sections 3.1, [⟨section-nilpotent-linear-transformations⟩](#), [⟨section-jordan-canonical-form⟩](#).

1.5 Reflectors

When we decompose matrices into products of other matrices, we often wish to create matrices with many zero entries. A **Householder matrix** is a unitary matrix which transforms a vector into a vector of the same size that is a nonzero multiple of the first column of an identity matrix, thus creating a vector with just a single nonzero entry. A typical application is to “zero out” entries of a matrix below the diagonal, column-by-column, in order to achieve a triangular matrix.

Definition 6. Given a nonzero vector $\mathbf{v} \in \mathbb{C}^n$, the Householder matrix for \mathbf{v} is

$$P = I_n - \left(\frac{2}{\langle \mathbf{v}, \mathbf{v} \rangle} \right) \mathbf{v} \mathbf{v}^*$$

. The vector \mathbf{v} is called the Householder vector.

A Householder matrix is both Hermitian and unitary.

Lemma 1.14. *The Householder matrix for $\mathbf{v} \in \mathbb{C}^n$ is Hermitian.*

Proof.

$$\begin{aligned} P^* &= \left(I_n - \left(\frac{2}{\langle \mathbf{v}, \mathbf{v} \rangle} \right) \mathbf{v} \mathbf{v}^* \right)^* \\ &= I_n^* - \left(\frac{2}{\langle \mathbf{v}, \mathbf{v} \rangle} \right) (\mathbf{v} \mathbf{v}^*)^* \\ &= I_n - \left(\frac{2}{\langle \mathbf{v}, \mathbf{v} \rangle} \right) (\mathbf{v}^*)^* \mathbf{v}^* \\ &= I_n - \left(\frac{2}{\langle \mathbf{v}, \mathbf{v} \rangle} \right) \mathbf{v} \mathbf{v}^* \\ &= P \end{aligned}$$

□

Lemma 1.15. *The Householder matrix for $\mathbf{v} \in \mathbb{C}^n$ is unitary.*

Proof.

$$\begin{aligned} P^* P &= P P \\ &= \left(I_n - \left(\frac{2}{\langle \mathbf{v}, \mathbf{v} \rangle} \right) \mathbf{v} \mathbf{v}^* \right) \left(I_n - \left(\frac{2}{\langle \mathbf{v}, \mathbf{v} \rangle} \right) \mathbf{v} \mathbf{v}^* \right) \\ &= I_n^2 - \left(\frac{2}{\langle \mathbf{v}, \mathbf{v} \rangle} \right) \mathbf{v} \mathbf{v}^* - \left(\frac{2}{\langle \mathbf{v}, \mathbf{v} \rangle} \right) \mathbf{v} \mathbf{v}^* + \left(\frac{4}{\langle \mathbf{v}, \mathbf{v} \rangle^2} \right) \mathbf{v} \mathbf{v}^* \mathbf{v} \mathbf{v}^* \\ &= I_n - \left(\frac{4}{\langle \mathbf{v}, \mathbf{v} \rangle} \right) \mathbf{v} \mathbf{v}^* + \left(\frac{4}{\langle \mathbf{v}, \mathbf{v} \rangle^2} \right) \mathbf{v} \langle \mathbf{v}, \mathbf{v} \rangle \mathbf{v}^* \\ &= I_n - \left(\frac{4}{\langle \mathbf{v}, \mathbf{v} \rangle} \right) \mathbf{v} \mathbf{v}^* + \left(\frac{4}{\langle \mathbf{v}, \mathbf{v} \rangle} \right) \mathbf{v} \mathbf{v}^* \\ &= I_n \end{aligned}$$

□

Our aim with a Householder matrix is to convert a vector \mathbf{x} into a scalar multiple of the first column of the identity matrix, \mathbf{e}_1 . Which Householder vector should we choose for constructing such a matrix, and which multiple will we get? It is an instructive exercise to reverse-engineer the choice by setting $P\mathbf{x} = \alpha\mathbf{e}_1$ (Exercise 12). Instead, we give the answer and prove that it does the desired job.

Theorem 1.16. *Given a vector $\mathbf{x} \in \mathbb{R}^n$, define $\mathbf{v} = \mathbf{x} \pm \|\mathbf{x}\| \mathbf{e}_1$ and let P be the Householder matrix for the Householder vector \mathbf{v} . Then $P\mathbf{x} = \mp \|\mathbf{x}\| \mathbf{e}_1$.*

Proof. We first establish an unmotivated identity.

$$\begin{aligned} \langle \mathbf{v}, \mathbf{v} \rangle &= (\mathbf{x} \pm \|\mathbf{x}\| \mathbf{e}_1)^* (\mathbf{x} \pm \|\mathbf{x}\| \mathbf{e}_1) \\ &= \mathbf{x}^* \mathbf{x} \pm \mathbf{x}^* \|\mathbf{x}\| \mathbf{e}_1 \pm (\|\mathbf{x}\| \mathbf{e}_1)^* \mathbf{x} + (\|\mathbf{x}\| \mathbf{e}_1)^* (\|\mathbf{x}\| \mathbf{e}_1) \\ &= \mathbf{x}^* \mathbf{x} \pm \|\mathbf{x}\| \mathbf{e}_1^* \mathbf{x} \pm \|\mathbf{x}\| \mathbf{e}_1^* \mathbf{x} + \mathbf{x}^* \mathbf{x} \mathbf{e}_1^* \mathbf{e}_1 \\ &= 2(\mathbf{x}^* \mathbf{x} \pm \|\mathbf{x}\| \mathbf{e}_1^* \mathbf{x}) \\ &= 2(\mathbf{x}^* \pm \|\mathbf{x}\| \mathbf{e}_1^*) \mathbf{x} \\ &= 2(\mathbf{x} \pm \|\mathbf{x}\| \mathbf{e}_1)^* \mathbf{x} \\ &= 2\mathbf{v}^* \mathbf{x} \end{aligned}$$

Then

$$\begin{aligned} P\mathbf{x} &= \left(I_n - \left(\frac{2}{\langle \mathbf{v}, \mathbf{v} \rangle} \right) \mathbf{v} \mathbf{v}^* \right) \mathbf{x} \\ &= I_n \mathbf{x} - \left(\frac{2}{\langle \mathbf{v}, \mathbf{v} \rangle} \right) \mathbf{v} \mathbf{v}^* \mathbf{x} \\ &= \mathbf{x} - \left(\frac{2}{2\mathbf{v}^* \mathbf{x}} \right) \mathbf{v} \mathbf{v}^* \mathbf{x} \\ &= \mathbf{x} - \mathbf{v} \\ &= \mathbf{x} - (\mathbf{x} \pm \|\mathbf{x}\| \mathbf{e}_1) \\ &= \mp \|\mathbf{x}\| \mathbf{e}_1 \end{aligned}$$

□

Example 7. Consider the vector \mathbf{x} and construct the Householder vector \mathbf{v} .

$$\mathbf{x} = \begin{bmatrix} 4 \\ 4 \\ 7 \end{bmatrix} \qquad \mathbf{v} = \mathbf{x} - 9\mathbf{e}_1 = \begin{bmatrix} -5 \\ 4 \\ 7 \end{bmatrix}$$

Then the Householder matrix for \mathbf{v} is

$$P = \begin{bmatrix} \frac{4}{9} & \frac{4}{9} & \frac{7}{9} \\ \frac{4}{9} & \frac{29}{45} & -\frac{28}{45} \\ \frac{7}{9} & -\frac{28}{45} & -\frac{4}{45} \end{bmatrix}$$

We can check that the matrix behaves as we expect.

$$P\mathbf{x} = \begin{bmatrix} 9 \\ 0 \\ 0 \end{bmatrix}$$

A Householder matrix is often called a **Householder reflection**. Why? Any Householder matrix, when thought of as a mapping of points in a physical space, will

fix the elements of a hyperplane and reflect any other points about that hyperplane. To see this, consider any vector \mathbf{w} and compare it with its image, $P\mathbf{w}$

$$\begin{aligned} P\mathbf{w} - \mathbf{w} &= I_n \mathbf{w} - \left(\frac{2}{\langle \mathbf{v}, \mathbf{v} \rangle} \right) \mathbf{v} \mathbf{v}^* \mathbf{w} - \mathbf{w} \\ &= -\frac{2}{\langle \mathbf{v}, \mathbf{v} \rangle} \mathbf{v} \langle \mathbf{v}, \mathbf{w} \rangle \\ &= -\frac{2 \langle \mathbf{v}, \mathbf{w} \rangle}{\langle \mathbf{v}, \mathbf{v} \rangle} \mathbf{v} \end{aligned}$$

So this difference is always a scalar multiple of the Householder vector \mathbf{v} , and thus every point gets moved in the same direction, the direction described by \mathbf{v} . Which points are fixed by P ? The difference above is the zero vector precisely when the scalar multiple is zero, which will happen when $\langle \mathbf{v}, \mathbf{w} \rangle = 0$. So the set of points/vectors which are orthogonal to \mathbf{v} will be unmoved. This is a subspace of dimension one less than the full space, which are typically described by the term **hyperplanes**.

To be more specific, consider the specific situation of Example 7, viewed in \mathbb{R}^3 . The hyperplane is the subspace orthogonal to \mathbf{v} , or the two-dimensional plane with \mathbf{v} as its normal vector, and equation $-5x + 4y + 7z = 0$. The points $(4, 4, 7)$ and $(9, 0, 0)$ lie on either side of the plane and are a reflection of each other in the plane, by which we mean the vector $(4, 4, 7) - (9, 0, 0) = (-5, 4, 7)$ is perpendicular (orthogonal, normal) to the plane.

Our choice of \mathbf{v} can be replaced by $\mathbf{v} = \mathbf{x} + \|\mathbf{x}\| \mathbf{e}_1$, so in the previous example we would have $\mathbf{v} = \begin{bmatrix} 13 \\ 4 \\ 7 \end{bmatrix}$, and then P would take \mathbf{x} to $\begin{bmatrix} -9 \\ 0 \\ 0 \end{bmatrix}$. This would be a reflection across the plane with equation $13x + 4y + 7z = 0$. Notice that the normal vector of this plane is orthogonal to the normal vector of the previous plane, which is not an accident (Exercise 9).

As a practical matter, we would choose the Householder vector which moves \mathbf{x} the furthest, to get better numerical behavior. So in our example above, the second choice would be better, since \mathbf{x} will be moved a distance $2\|\mathbf{v}\|$ and the second \mathbf{v} has a larger norm.

Exercise 9. In the real case, we have two choices for a Householder vector which will “zero out” most of a vector. Show that these two vectors, $\mathbf{x} + \|\mathbf{x}\| \mathbf{e}_1$ and $\mathbf{x} - \|\mathbf{x}\| \mathbf{e}_1$, are orthogonal to each other.

Exercise 10. Prove the following generalization of Theorem 1.16. Given a vector $\mathbf{x} \in \mathbb{C}^n$, define $\rho = [\mathbf{x}]_1 / \|\mathbf{x}\|$ and $\mathbf{v} = \mathbf{x} \pm \rho \|\mathbf{x}\| \mathbf{e}_1$ and let P be the Householder matrix for the Householder vector \mathbf{v} . Then $P\mathbf{x} = \mp \rho \|\mathbf{x}\| \mathbf{e}_1$.

Exercise 11. Suppose that P is a Householder matrix of size n and $\mathbf{b} \in \mathbb{C}^n$ is any vector. Find an expression for the matrix-vector product $P\mathbf{b}$ which will suggest a way to compute this vector with fewer than the $\sim 2n^2$ operations required for a general matrix-vector product.

Exercise 12. Begin with the condition that a Householder matrix will accomplish $P\mathbf{x} = \alpha \mathbf{e}_1$ and “discover” the choice for the Householder vector described in Theorem 1.16. Hint: The condition implies that the Householder vector \mathbf{v} is in the span of $\{\mathbf{x}, \mathbf{e}_1\}$.

1.6 Projectors

When we multiply a vector by a matrix, we form a linear combination of the columns of the matrix. Said differently, the result of the product is in the column space of the matrix. So we can think of a matrix as moving a vector into a subspace, and we call that subspace the column space of the matrix $\mathcal{C}(A)$. In the case of a linear transformation, we call this subspace the range, $\mathcal{R}(T)$, or we might call it the image. A **projector** is a square matrix which moves vectors into a subspace (like any matrix can), but fixes vectors already in the subspace. This property earns a projector the moniker **idempotent**. We will see that projectors have a variety of interesting properties.

1.6.1 Oblique Projectors

Definition 7. A square matrix P is a projector if $P^2 = P$.

A projector fixes vectors in its column space.

Lemma 1.17. Suppose P is a projector and $\mathbf{x} \in \mathcal{C}(P)$. Then $P\mathbf{x} - \mathbf{x} = \mathbf{0}$.

Proof. Since $\mathbf{x} \in \mathcal{C}(P)$, there is a vector \mathbf{w} such that $P\mathbf{w} = \mathbf{x}$. Then

$$P\mathbf{x} - \mathbf{x} = P(P\mathbf{w}) - P\mathbf{w} = P^2\mathbf{w} - P\mathbf{w} = P\mathbf{w} - P\mathbf{w} = \mathbf{0}.$$

□

For a general vector, the difference between the vector and its image under a projector may not always be the zero vector, but it will be a vector in the null space of the projector.

Lemma 1.18. Suppose P is a projector of size n and $\mathbf{x} \in \mathbb{C}^n$ is any vector. Then $P\mathbf{x} - \mathbf{x} \in \mathcal{N}(P)$. Furthermore, $\mathcal{N}(P) = \{P\mathbf{x} - \mathbf{x} \mid \mathbf{x} \in \mathbb{C}^n\}$.

Proof. First,

$$P(P\mathbf{x} - \mathbf{x}) = P^2\mathbf{x} - P\mathbf{x} = P\mathbf{x} - P\mathbf{x} = \mathbf{0}.$$

To establish the second half of the claimed set equality, suppose $\mathbf{z} \in \mathcal{N}(P)$, then

$$\mathbf{z} = \mathbf{0} - (-\mathbf{z}) = P(-\mathbf{z}) - (-\mathbf{z})$$

which establishes that $\mathbf{z} \in \{P\mathbf{x} - \mathbf{x} \mid \mathbf{x} \in \mathbb{C}^n\}$. □

When the null space of a projector has dimension one, it is easy to understand the choice of the term “projector”. Imagine the setting in three dimensions where the column space of the projector is a subspace of dimension two, which is physically a plane through the origin. Imagine some vector as an arrow from the origin, or as just the point that would be at the tip of the arrow. A light shines on the vector and casts a shadow onto the plane (either another arrow, or just a point). This shadow is the projection, the image of the projector. The image of the shadow is unchanged, since shining the light on the vector that is the shadow will not move it. What direction does the light come from? What is the vector that describes the change from the vector to its shadow (projection)? For a vector \mathbf{x} , this direction is $P\mathbf{x} - \mathbf{x}$, an element of the null space of P . So if $\mathcal{N}(P)$ has dimension one, then every vector is moved in the same direction, a multiple of a lone basis vector for $\mathcal{N}(P)$. This matches our assumptions about physical light from a distant source, with rays all moving parallel to each other. Here is a simple example of just this scenario.

Example 8. Verify the following facts about the matrix P to understand that it is a projector and to understand its geometry.

$$P = \frac{1}{13} \begin{bmatrix} 11 & -3 & -5 \\ -4 & 7 & -10 \\ -2 & -3 & 8 \end{bmatrix}$$

1. $P^2 = P$

2. $\mathcal{C}(P) = \left\langle \left\{ \begin{bmatrix} 1 \\ 0 \\ -\frac{2}{5} \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ -\frac{3}{5} \end{bmatrix} \right\} \right\rangle$

3. $\mathcal{N}(P) = \left\langle \left\{ \begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix} \right\} \right\rangle$

So P sends every vector onto a two-dimensional subspace, with an equation we might write as $2x + 3y + 5z = 0$ in Cartesian coordinates, or which we might describe as the plane through the origin with normal vector $\mathbf{n} = 2\vec{i} + 3\vec{j} + 5\vec{k}$. Vectors, or points, are always moved in the direction of the vector $\mathbf{d} = \vec{i} + 2\vec{j} + 1\vec{k}$ —this is the direction the light is shining. Exercise 13 asks you to experiment further.

Exercise 13. Continue experimenting with Example 8 by constructing a vector *not* in the column space of P . Compute its image under P and verify that it is a linear combination of the basis vectors given in the example. Compute the direction your vector moved and verify that it is a scalar multiple of the basis vector for the null space given in the example. Finally, construct a new vector in the column space and verify that it is unmoved by P .

Given a projector, we can define a complementary projector, which has some interesting properties.

Definition 8. Given a projector P , the complementary projector to P is $I - P$.

The next lemma justifies calling $I - P$ a projector.

Lemma 1.19. *If P is a projector then $I - P$ is also a projector.*

Proof.

$$(I - P)^2 = I^2 - P - P + P^2 = I - P - P + P = I - P$$

□

The complementary projector to P projects onto the null space of P .

Lemma 1.20. *Suppose P is a projector. Then $\mathcal{C}(I - P) = \mathcal{N}(P)$ and therefore $\mathcal{N}(I - P) = \mathcal{C}(P)$.*

Proof. First, suppose $\mathbf{x} \in \mathcal{N}(P)$. Then

$$(I - P)\mathbf{x} = I\mathbf{x} - P\mathbf{x} = \mathbf{x}$$

demonstrating that \mathbf{x} is a linear combination of the columns of $I - P$. So $\mathcal{N}(P) \subseteq \mathcal{C}(I - P)$.

Now, suppose $\mathbf{x} \in \mathcal{C}(I - P)$. Then there is a vector \mathbf{w} such that $\mathbf{x} = (I - P)\mathbf{w}$. Then

$$P\mathbf{x} = P(I - P)\mathbf{w} = (P - P^2)\mathbf{w} = \mathcal{O}\mathbf{w} = \mathbf{0}.$$

So $\mathcal{C}(I - P) \subseteq \mathcal{N}(P)$.

To establish the second conclusion, replace the projector P in the first conclusion by the projector $I - P$. □

Using these facts about complementary projectors we find a simple direct sum decomposition.

Theorem 1.21. *Suppose P is a projector of size n . Then $\mathbb{C}^n = \mathcal{C}(P) \oplus \mathcal{N}(P)$.*

Proof. First, we show that $\mathcal{C}(P) \cap \mathcal{N}(P) = \{\mathbf{0}\}$. Suppose $\mathbf{x} \in \mathcal{C}(P) \cap \mathcal{N}(P)$. Since $\mathbf{x} \in \mathcal{C}(P)$, Lemma 1.20 implies that $\mathbf{x} \in \mathcal{N}(I - P)$. So

$$\mathbf{x} = \mathbf{x} - \mathbf{0} = \mathbf{x} - P\mathbf{x} = (I - P)\mathbf{x} = \mathbf{0}.$$

Using Lemma 1.20 again, $\mathcal{N}(P) = \mathcal{C}(I - P)$. We show that an arbitrary vector $\mathbf{w} \in \mathbb{C}^n$ can be written as a sum of two vectors from the two column spaces,

$$\mathbf{w} = I\mathbf{w} - P\mathbf{w} + P\mathbf{w} = (I - P)\mathbf{w} + P\mathbf{w}.$$

So \mathbb{C}^n , $\mathcal{C}(P)$ and $\mathcal{N}(P)$ meet the hypotheses of Theorem 1.3, allowing us to establish the direct sum. \square

1.6.2 Orthogonal Projectors

The projectors of the previous section would be termed **oblique projectors** since no assumption was made about the direction that a vector was moved when projected. We remedy that situation now by defining an **orthogonal projector** to be a projector where the complementary subspace is orthogonal to the space the projector projects onto.

Definition 9. A projector P is orthogonal if $\mathcal{N}(P) = (\mathcal{C}(P))^\perp$.

We know from Theorem 1.21 that for a projector P , $\mathbb{C}^n = \mathcal{C}(P) \oplus \mathcal{N}(P)$. We also know by Corollary 1.11, that for any $m \times n$ matrix A , $\mathbb{C}^m = \mathcal{C}(A) \oplus \mathcal{C}(A)^\perp = \mathcal{C}(A) \oplus \mathcal{N}(A^*)$. So, superficially, we might expect orthogonal projectors to be Hermitian. And so it is.

Theorem 1.22. *Suppose P is a projector. Then P is an orthogonal projector if and only if P is Hermitian.*

Proof. THEOREM HMIP says that Hermitian matrices are characterized by the property that $\langle A\mathbf{x}, \mathbf{y} \rangle = \langle \mathbf{x}, A\mathbf{y} \rangle$ for every choice of the vectors \mathbf{x}, \mathbf{y} . We will use this result in both halves of the proof.

Suppose that $\mathbf{x} \in \mathcal{N}(P)$. Then for any $\mathbf{y} \in \mathcal{C}(P)$, there is a vector \mathbf{w} that allows us to write

$$\langle \mathbf{x}, \mathbf{y} \rangle = \langle \mathbf{x}, P\mathbf{w} \rangle = \langle P\mathbf{x}, \mathbf{w} \rangle = \langle \mathbf{0}, \mathbf{w} \rangle = 0.$$

So $\mathcal{N}(P) \subseteq \mathcal{C}(A)^\perp$.

Now suppose that $\mathbf{x} \in \mathcal{C}(P)^\perp$. Consider,

$$\langle P\mathbf{x}, P\mathbf{x} \rangle = \langle P^2\mathbf{x}, \mathbf{x} \rangle = \langle P\mathbf{x}, \mathbf{x} \rangle = 0.$$

By THEOREM PIP, we conclude that $P\mathbf{x} = \mathbf{0}$ and $\mathbf{x} \in \mathcal{N}(P)$. So $\mathcal{C}(A)^\perp \subseteq \mathcal{N}(P)$ and we have established the set equality of Definition 9.

Let $\mathbf{u}, \mathbf{v} \in \mathbb{C}^n$ be any two vectors. Decompose each into two pieces, the first from the column space, the second from the null space, according to Theorem 1.21. So

$$\mathbf{u} = \mathbf{u}_1 + \mathbf{u}_2$$

$$\mathbf{v} = \mathbf{v}_1 + \mathbf{v}_2$$

with $\mathbf{u}_1, \mathbf{v}_1 \in \mathcal{C}(P)$ and $\mathbf{u}_2, \mathbf{v}_2 \in \mathcal{N}(P)$. Then

$$\begin{aligned} \langle P\mathbf{u}, \mathbf{v} \rangle &= \langle P\mathbf{u}_1 + P\mathbf{u}_2, \mathbf{v}_1 + \mathbf{v}_2 \rangle = \langle P\mathbf{u}_1, \mathbf{v}_1 + \mathbf{v}_2 \rangle \\ &= \langle \mathbf{u}_1, \mathbf{v}_1 + \mathbf{v}_2 \rangle = \langle \mathbf{u}_1, \mathbf{v}_1 \rangle + \langle \mathbf{u}_1, \mathbf{v}_2 \rangle = \langle \mathbf{u}_1, \mathbf{v}_1 \rangle \\ \langle \mathbf{u}, P\mathbf{v} \rangle &= \langle \mathbf{u}_1 + \mathbf{u}_2, P\mathbf{v}_1 + P\mathbf{v}_2 \rangle = \langle \mathbf{u}_1 + \mathbf{u}_2, P\mathbf{v}_1 \rangle \\ &= \langle \mathbf{u}_1 + \mathbf{u}_2, \mathbf{v}_1 \rangle = \langle \mathbf{u}_1, \mathbf{v}_1 \rangle + \langle \mathbf{u}_2, \mathbf{v}_1 \rangle = \langle \mathbf{u}_1, \mathbf{v}_1 \rangle \end{aligned}$$

Since $\langle P\mathbf{u}, \mathbf{v} \rangle = \langle \mathbf{u}, P\mathbf{v} \rangle$ for all choices of $\mathbf{u}, \mathbf{v} \in \mathbb{C}^n$, THEOREM HMIP, establishes that P is Hermitian. \square

There is an easy recipe for creating orthogonal projectors onto a given subspace. We will first informally motivate the construction, then give the formal proof. Suppose U is a subspace with a basis $\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3, \dots, \mathbf{u}_k$ and let A be a matrix with these basis vectors as the columns. Let P denote the desired orthogonal projector, and consider its action on an arbitrary vector \mathbf{x} . To project onto U , we must have $P\mathbf{x} \in \mathcal{C}(A)$, so there is a vector \mathbf{w} such that $P\mathbf{x} = A\mathbf{w}$. The orthogonality condition will be satisfied if $P\mathbf{x} - \mathbf{x}$ is orthogonal to every vector of U . It is enough to require orthogonality to each basis vector of U , and hence to each column of A . So we have

$$\begin{aligned} A^* (P\mathbf{x} - \mathbf{x}) &= \mathbf{0} \\ A^* A\mathbf{w} - A^* \mathbf{x} &= \mathbf{0} \\ A^* A\mathbf{w} &= A^* \mathbf{x} \end{aligned}$$

As A has full rank, A^*A is nonsingular ($\langle\langle$ adjoint- A is nonsingular result $\rangle\rangle$), so we can employ its inverse to find

$$P\mathbf{x} = A\mathbf{w} = A(A^*A)^{-1}A^*\mathbf{x}$$

This suggests that $P = A(A^*A)^{-1}A^*$. And so it is.

Theorem 1.23. *Suppose U is a subspace and A is a matrix whose columns form a basis of U . Then $P = A(A^*A)^{-1}A^*$ is an orthogonal projector onto U .*

Proof. Because A is the leftmost term in the product for P , $\mathcal{C}(P) \subseteq \mathcal{C}(A)$. Because $(A^*A)^{-1}A^*$ has full (column) rank, $\mathcal{C}(A) \subseteq \mathcal{C}(P)$. So the image of the projector is exactly U .

Now we verify that P is a projector.

$$\begin{aligned} P^2 &= \left(A(A^*A)^{-1}A^* \right) \left(A(A^*A)^{-1}A^* \right) \\ &= A(A^*A)^{-1}(A^*A)(A^*A)^{-1}A^* \\ &= A(A^*A)^{-1}A^* \\ &= P \end{aligned}$$

And lastly, orthogonality against a basis of U .

$$\begin{aligned} A^* (P\mathbf{x} - \mathbf{x}) &= A^* A(A^*A)^{-1}A^*\mathbf{x} - A^*\mathbf{x} \\ &= A^*\mathbf{x} - A^*\mathbf{x} \\ &= \mathbf{0} \end{aligned}$$

\square

Suppose the basis vectors of U described in Theorem 1.23 form an orthonormal set, and in acknowledgment we denote the matrix with these vectors as columns by Q . Then the projector simplifies to $P = Q(Q^*Q)^{-1}Q^* = QQ^*$. The other interesting special case is when U is 1-dimensional (a “line”). Then A^*A is just the square of the norm of the lone basis vector. With this scalar moved out of the way, the remaining computation, AA^* , is an outer product that results in a rank 1 matrix (as we would expect).

Exercise 14. Illustrate Theorem 1.22 by proving directly that the orthogonal projector described in Theorem 1.23 is Hermitian.

Exercise 15. Construct the orthogonal projector onto the line spanned by

$$\mathbf{v} = \begin{bmatrix} 4 \\ 2 \\ 1 \\ 10 \end{bmatrix}.$$

Illustrate its use by projecting some vector not on the line, and verifying that the difference between the vector and its projection is orthogonal to the line.

Exercise 16. Construct the orthogonal projector onto the subspace

$$U = \left\langle \left\{ \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 4 \\ 2 \\ 2 \\ 5 \end{bmatrix} \right\} \right\rangle.$$

Illustrate its use by projecting some vector not in the subspace, and verifying that the difference between the vector and its projection is orthogonal to the line.

Exercise 17. Redo Exercise 16 but first convert the basis for U to an orthonormal basis via the Gram-Schmidt process THEOREM GSP and then use the simpler construction applicable to the case of an orthonormal basis.

1.7 Normal Matrices

Normal matrices comprise a broad class of interesting matrices, many of which you probably already know by other names. But they are most interesting since they define exactly which matrices we can diagonalize via a unitary matrix. This is the upcoming Theorem $\langle\langle$ orthonormal diagonalization $\rangle\rangle$.

Definition 10. The square matrix A is normal if $A^*A = AA^*$.

So a normal matrix commutes with its adjoint. Part of the beauty of this definition is that it includes many other types of matrices. A diagonal matrix will commute with its adjoint, since the adjoint is again diagonal and the entries are just conjugates of the entries of the original diagonal matrix. A Hermitian (self-adjoint) matrix (DEFINITION HM) will trivially commute with its adjoint, since the two matrices are the same. A real, symmetric matrix is Hermitian, so these matrices are also normal. A unitary matrix (DEFINITION UM) has its adjoint as its inverse, and inverses commute (THEOREM OSIS), so unitary matrices are normal. Another class of normal matrices is the skew-symmetric matrices. However, these broad classes of matrices do not capture all of the normal matrices, as the next example shows.

Example 9. Consider the matrix $\begin{bmatrix} 1 & -1 \\ 1 & 1 \end{bmatrix}$. Then

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

so we see by Definition 10 that A is normal. However, notice that A is not symmetric (hence, as a real matrix, not Hermitian), not unitary, and not skew-symmetric.

1.8 Positive Semi-Definite Matrices

Positive semi-definite matrices (and their cousins, positive definite matrices) are square matrices which in many ways behave like non-negative (respectively, positive) real numbers. These results will be useful as we study various matrix decompositions in Chapter 2.

Definition 11. A square matrix A of size n is **positive semi-definite** if A is Hermitian and for all $\mathbf{x} \in \mathbb{C}^n$, $\langle \mathbf{x}, A\mathbf{x} \rangle \geq 0$.

For a definition of **positive definite** replace the inequality in the definition with a strict inequality, and exclude the zero vector from the vectors \mathbf{x} required to meet the condition. Similar variations allow definitions of **negative definite** and **negative semi-definite**.

Our first theorem in this section gives us an easy way to build positive semi-definite matrices.

Theorem 1.24. *Suppose that A is any $m \times n$ matrix. Then the matrices A^*A and AA^* are positive semi-definite matrices.*

Proof. We will give the proof for the first matrix, the proof for the second is entirely similar. First we check that A^*A is Hermitian, with an appeal to DEFINITION HM,

$$(A^*A)^* = A^*(A^*)^* = A^*A$$

Second, for any $\mathbf{x} \in \mathbb{C}^n$, THEOREM AIP and THEOREM PIP give,

$$\langle \mathbf{x}, A^*A\mathbf{x} \rangle = \langle (A^*)^*\mathbf{x}, A\mathbf{x} \rangle = \langle A\mathbf{x}, A\mathbf{x} \rangle \geq 0$$

which is the second condition for a positive semi-definite matrix. \square

A statement very similar to the converse of this theorem is also true. Any positive semi-definite matrix can be realized as the product of a square matrix, B , with its adjoint, B^* . (See ?? after studying this entire section.) The matrices A^*A and AA^* will be important later when we define singular values in Section ((section-SVD)).

Positive semi-definite matrices can also be characterized by their eigenvalues, without any mention of inner products. This next result further reinforces the notion that positive semi-definite matrices behave like non-negative real numbers.

Theorem 1.25. *Suppose that A is a Hermitian matrix. Then A is a positive semi-definite matrix if and only if every eigenvalue λ of A has $\lambda \geq 0$.*

Proof. First notice first that since we are considering only Hermitian matrices in this theorem, it is always possible to compare eigenvalues with the real number zero, since eigenvalues of Hermitian matrices are all real numbers (THEOREM HMRE).

(\Rightarrow) Let $\mathbf{x} \neq 0$ be an eigenvector of A for λ . Since A is positive semi-definite,

$$\lambda \langle \mathbf{x}, \mathbf{x} \rangle = \langle \mathbf{x}, \lambda \mathbf{x} \rangle = \langle \mathbf{x}, A\mathbf{x} \rangle \geq 0$$

By THEOREM PIP we know $\langle \mathbf{x}, \mathbf{x} \rangle > 0$, so we conclude that $\lambda \geq 0$.

(\Leftarrow) Let n denote the size of A . Suppose that $\lambda_1, \lambda_2, \lambda_3, \dots, \lambda_n$ are the eigenvalues of the Hermitian matrix A , each of which is non-negative. Let $B = \{\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, \dots, \mathbf{x}_n\}$ be a set of associated eigenvectors for these eigenvalues. Since a Hermitian matrix is normal (10), THEOREM OBNM allows us to choose B to also be an orthonormal basis of \mathbb{C}^n . Choose any $\mathbf{x} \in \mathbb{C}^n$ and let $a_1, a_2, a_3, \dots, a_n$ be the scalars guaranteed by the spanning property of the basis B , so $\mathbf{x} = \sum_{i=1}^n a_i \mathbf{x}_i$. Since we have presumed A is Hermitian, we need only check the second condition of the definition. The use of an orthonormal basis provides the simplification for the last equality.

$$\begin{aligned} \langle \mathbf{x}, A\mathbf{x} \rangle &= \left\langle \sum_{i=1}^n a_i \mathbf{x}_i, A \sum_{j=1}^n a_j \mathbf{x}_j \right\rangle \\ &= \left\langle \sum_{i=1}^n a_i \mathbf{x}_i, \sum_{j=1}^n a_j A\mathbf{x}_j \right\rangle \\ &= \left\langle \sum_{i=1}^n a_i \mathbf{x}_i, \sum_{j=1}^n a_j \lambda_j \mathbf{x}_j \right\rangle \\ &= \sum_{i=1}^n \sum_{j=1}^n \langle a_i \mathbf{x}_i, a_j \lambda_j \mathbf{x}_j \rangle \\ &= \sum_{i=1}^n \sum_{j=1}^n \bar{a}_i a_j \lambda_j \langle \mathbf{x}_i, \mathbf{x}_j \rangle \\ &= \sum_{i=1}^n \bar{a}_i a_i \lambda_i \langle \mathbf{x}_i, \mathbf{x}_i \rangle + \sum_{i=1}^n \sum_{\substack{j=1 \\ j \neq i}}^n \bar{a}_i a_j \lambda_j \langle \mathbf{x}_i, \mathbf{x}_j \rangle \\ &= \sum_{i=1}^n \bar{a}_i a_i \lambda_i \end{aligned}$$

The expression $\bar{a}_i a_i$ is the modulus of a_i squared, hence is always non-negative. With the eigenvalues assumed non-negative, this final sum is clearly non-negative as well, as desired. \square

As positive semi-definite matrices are defined to be Hermitian, they are then normal and subject to orthonormal diagonalization (THEOREM OD). Now consider the interpretation of orthonormal diagonalization as a rotation to principal axes, a stretch by a diagonal matrix and a rotation back (SUBSECTION OD.OD). For a positive semi-definite matrix, the diagonal matrix has diagonal entries that are the non-negative eigenvalues of the original positive semi-definite matrix. So the “stretching” along each axis is never a reflection.

Chapter 2

Matrix Decompositions

A **matrix decomposition** is a way to express a matrix as a combination of other matrices, which are ostensibly simpler matrices. The combination is often a product of two or three matrices, though it can be a sum, as in the case of a rank one decomposition ($\langle\langle$ rank one decompositions $\rangle\rangle$). The constituent matrices are simpler because they have many zero entries, or have many strategically placed entries that are one, or the nonzero entries lie on the diagonal (or close by) or \dots . Furthermore, the constituent matrices may be simpler because they have desirable properties that make them easier to work with, such as being nonsingular or triangular or Hermitian or \dots . We will see examples of all of this behavior in this chapter.

There is a “Big Five” of matrix decompositions, which you will come to know as the LU, QR, SVD, Schur and Cholesky. Every student of advanced linear algebra should become intimately familiar with these basic decompositions. There are many other ways to decompose a matrix, and we will see these at other junctures. Encyclopedic texts like Horn & Johnson, [?] or Watkins [?] are good places to begin exploring more.

2.1 LU (Triangular) Decomposition

The **LU decomposition** begins with any matrix and describes it as a product of a lower-triangular matrix (L) and an upper-triangular matrix (U). Hence the customary shorthand name, “LU”. The term **triangular decomposition** might be more evocative, if not more verbose.

You will notice that the LU decomposition feels very much like reduced row-echelon form, and in some ways could be considered an improvement. Triangular matrices are defined in Subsection SUBSECTION OD.TM and two basic facts are that the product of two triangular matrices “of the same type” (i.e. both upper or both lower) is again of that type (THEOREM PTMT) and the inverse of a nonsingular triangular matrix will be triangular of the same type (THEOREM ITMT).

2.1.1 LU Decomposition, Nonsingular Case

Theorem 2.1. *Suppose A is a square matrix of size n . Let A_k be the $k \times k$ matrix formed from A by taking the first k rows and the first k columns. Suppose that A_k is nonsingular for all $1 \leq k \leq n$. Then there is a lower triangular matrix L with all of its diagonal entries equal to 1 and an upper triangular matrix U such that $A = LU$. Furthermore, this decomposition is unique.*

Proof. We will row reduce A to a row-equivalent upper triangular matrix through a series of row operations, forming intermediate matrices A'_j , $1 \leq j \leq n$, that denote the state of the conversion after working on column j . First, the lone entry of A_1 is $[A]_{11}$ and this scalar must be nonzero if A_1 is nonsingular (THEOREM SMZD). We can use row operations DEFINITION RO of the form $\alpha R_1 + R_k$, $2 \leq k \leq n$, where $\alpha = -[A]_{1k} / [A]_{11}$ to place zeros in the first column below the diagonal. The first two rows and columns of A'_1 are a 2×2 upper triangular matrix whose determinant is equal to the determinant of A_2 , since the matrices are row-equivalent through a sequence of row operations strictly of the third type (THEOREM DRCMA). As such the diagonal entries of this 2×2 submatrix of A'_1 are nonzero. We can employ this nonzero diagonal element with row operations of the form $\alpha R_2 + R_k$, $3 \leq k \leq n$ to place zeros below the diagonal in the second column. We can continue this process, column by column. The key observations are that our hypothesis on the nonsingularity of the A_k will guarantee a nonzero diagonal entry for each column when we need it, that the row operations employed are always of the third type using a multiple of a row to transform another row *with a greater row index*, and that the final result will be a nonsingular upper triangular matrix. This is the desired matrix U .

Each row operation described in the previous paragraph can be accomplished with matrix multiplication by the appropriate elementary matrix (THEOREM EM-DRO). Since every row operation employed is adding a multiple of a row to a subsequent row these elementary matrices are of the form $E_{j,k}(\alpha)$ with $j < k$. By DEFINITION ELEM, these matrices are lower triangular with every diagonal entry equal to 1. We know that the product of two such matrices will again be lower triangular (THEOREM PTMT), but also, as you can also easily check using a proof with a style similar to one above, that the product maintains all 1's on the diagonal. Let $E_1, E_2, E_3, \dots, E_m$ denote the elementary matrices for this sequence of row operations. Then

$$U = E_m E_{m-1} \dots E_3 E_2 E_1 A = L' A$$

where L' is the product of the elementary matrices, and we know L' is lower triangular with all 1's on the diagonal. Our desired matrix L is then $L = (L')^{-1}$. By THEOREM ITMT, L is lower triangular with all 1's on the diagonal and $A = LU$, as desired.

The process just described is deterministic. That is, the proof is constructive, with no freedom for each of us to walk through it differently. But could there be other matrices with the same properties as L and U that give such a decomposition of A ? In other words, is the decomposition unique (PROOF TECHNIQUE U)? Suppose that we have two triangular decompositions, $A = L_1 U_1$ and $A = L_2 U_2$. Since A is nonsingular, two applications of THEOREM NPNT imply that L_1, L_2, U_1, U_2 are all nonsingular. We have

$$\begin{aligned} L_2^{-1} L_1 &= L_2^{-1} I_n L_1 \\ &= L_2^{-1} A A^{-1} L_1 \\ &= L_2^{-1} L_2 U_2 (L_1 U_1)^{-1} L_1 \\ &= L_2^{-1} L_2 U_2 U_1^{-1} L_1^{-1} L_1 \\ &= I_n U_2 U_1^{-1} I_n \\ &= U_2 U_1^{-1} \end{aligned}$$

THEOREM ITMT tells us that L_2^{-1} is lower triangular and has 1's as the diagonal entries. By THEOREM PTMT, the product $L_2^{-1} L_1$ is again lower triangular, and it is simple to check (as before) that the diagonal entries of the product are again

all 1's. By the entirely similar process we can conclude that the product $U_2U_1^{-1}$ is upper triangular. Because these two products are equal, their common value is a matrix that is both lower triangular *and* upper triangular, with all 1's on the diagonal. The only matrix meeting these three requirements is the identity matrix (DEFINITION IM). So, we have,

$$I_n = L_2^{-1}L_1 \Rightarrow L_2 = L_1 \qquad I_n = U_2U_1^{-1} \Rightarrow U_1 = U_2$$

which establishes the uniqueness of the decomposition. \square

Studying the proofs of some previous theorems will perhaps give you an idea for an approach to computing a triangular decomposition. In the proof of THEOREM CINM we augmented a nonsingular matrix with an identity matrix of the same size, and row-reduced until the original matrix became the identity matrix (as we knew in advance would happen, since we knew THEOREM NMRRI). THEOREM PEEF tells us about properties of extended echelon form, and in particular, that $B = JA$, where A is the matrix that begins on the left, and B is the reduced row-echelon form of A . The matrix J is the result on the right side of the augmented matrix, which is the result of applying the same row operations to the identity matrix. We should recognize now that J is just the product of the elementary matrices (SUBSECTION DM.EM) that perform these row operations. THEOREM ITMT used the extended echelon form to discern properties of the inverse of a triangular matrix. THEOREM TD proves the existence of a triangular decomposition by applying specific row operations, and tracking the relevant elementary row operations. It is not a great leap to combine these observations into a computational procedure.

To find the triangular decomposition of A , augment A with the identity matrix of the same size and call this new $2n \times n$ matrix, M . Perform row operations on M that convert the first n columns to an upper triangular matrix. Do this using only row operations that add a scalar multiple of one row to another row *with higher index* (i.e. lower down). In this way, the last n columns of M will be converted into a lower triangular matrix with 1's on the diagonal (since M has 1's in these locations initially). We could think of this process as doing about half of the work required to compute the inverse of A . Take the first n columns of the row-equivalent version of M and call this matrix U .

Take the final n columns of the row-equivalent version of M and call this matrix L' . Then by a proof employing elementary matrices, or a proof similar in spirit to the one used to prove THEOREM PEEF, we arrive at a result similar to the second assertion of THEOREM PEEF. Namely, $U = L'A$. Multiplication on the left, by the inverse of L' , will give us a decomposition of A (which we know to be unique). Ready? Lets try it.

Example 10. In this example, we will illustrate the process for computing a triangular decomposition, as described in the previous paragraphs. Consider the nonsingular square matrix A of size 4,

$$A = \begin{bmatrix} -2 & 6 & -8 & 7 \\ -4 & 16 & -14 & 15 \\ -6 & 22 & -23 & 26 \\ -6 & 26 & -18 & 17 \end{bmatrix}$$

We form M by augmenting A with the size 4 identity matrix I_4 . We will perform the allowed operations, column by column, only reporting intermediate results as we finish converting each column. It is easy to determine exactly which row operations we perform, since the final four columns contain a record of each such operation.

We will not verify our hypotheses about the nonsingularity of the A_k , since if we do not have these conditions, we will reach a stage where a diagonal entry is zero and we cannot create the row operations we need to zero out the bottom portion of the associated column. In other words, we can boldly proceed and the necessity of our hypotheses will become apparent.

$$\begin{aligned}
 M &= \begin{bmatrix} -2 & 6 & -8 & 7 & 1 & 0 & 0 & 0 \\ -4 & 16 & -14 & 15 & 0 & 1 & 0 & 0 \\ -6 & 22 & -23 & 26 & 0 & 0 & 1 & 0 \\ -6 & 26 & -18 & 17 & 0 & 0 & 0 & 1 \end{bmatrix} \\
 &\rightarrow \begin{bmatrix} -2 & 6 & -8 & 7 & 1 & 0 & 0 & 0 \\ 0 & 4 & 2 & 1 & -2 & 1 & 0 & 0 \\ 0 & 4 & 1 & 5 & -3 & 0 & 1 & 0 \\ 0 & 8 & 6 & -4 & -3 & 0 & 0 & 1 \end{bmatrix} \\
 &\rightarrow \begin{bmatrix} -2 & 6 & -8 & 7 & 1 & 0 & 0 & 0 \\ 0 & 4 & 2 & 1 & -2 & 1 & 0 & 0 \\ 0 & 0 & -1 & 4 & -1 & -1 & 1 & 0 \\ 0 & 0 & 2 & -6 & 1 & -2 & 0 & 1 \end{bmatrix} \\
 &\rightarrow \begin{bmatrix} -2 & 6 & -8 & 7 & 1 & 0 & 0 & 0 \\ 0 & 4 & 2 & 1 & -2 & 1 & 0 & 0 \\ 0 & 0 & -1 & 4 & -1 & -1 & 1 & 0 \\ 0 & 0 & 0 & 2 & -1 & -4 & 2 & 1 \end{bmatrix}
 \end{aligned}$$

So at this point, we have U and L' ,

$$U = \begin{bmatrix} -2 & 6 & -8 & 7 \\ 0 & 4 & 2 & 1 \\ 0 & 0 & -1 & 4 \\ 0 & 0 & 0 & 2 \end{bmatrix} \qquad L' = \begin{bmatrix} 1 & 0 & 0 & 0 \\ -2 & 1 & 0 & 0 \\ -1 & -1 & 1 & 0 \\ -1 & -4 & 2 & 1 \end{bmatrix}$$

Then by whatever procedure we like (such as THEOREM CINM), we find

$$L = (L')^{-1} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 2 & 1 & 0 & 0 \\ 3 & 1 & 1 & 0 \\ 3 & 2 & -2 & 1 \end{bmatrix}$$

It is instructive to verify that indeed $LU = A$.

2.1.2 Solving Systems with Triangular Decompositions

In this section we give an explanation of why you might be interested in a triangular decomposition for a matrix. Many of the computational problems in linear algebra revolve around solving large systems of equations, or nearly equivalently, finding inverses of large matrices. Suppose we have a system of equations with coefficient matrix A and vector of constants \mathbf{b} , and suppose further that A has the triangular decomposition $A = LU$.

Let \mathbf{y} be the solution to the linear system $\mathcal{LS}(L, \mathbf{b})$, so that by THEOREM SLEMM, we have $L\mathbf{y} = \mathbf{b}$. Notice that since L is nonsingular, this solution is unique, and the form of L makes it trivial to solve the system. The first component of \mathbf{y} is determined easily, and we can continue on through determining the components of \mathbf{y} , without even ever dividing. Now, with \mathbf{y} in hand, consider the linear system, $\mathcal{LS}(U, \mathbf{y})$. Let \mathbf{x} be the unique solution to this system, so by THEOREM

SLEMM we have $U\mathbf{x} = \mathbf{y}$. Notice that a system of equations with U as a coefficient matrix is also straightforward to solve, though we will compute the bottom entries of \mathbf{x} first, and we will need to divide. The upshot of all this is that \mathbf{x} is a solution to $\mathcal{LS}(A, \mathbf{b})$, as we now show,

$$A\mathbf{x} = LU\mathbf{x} = L(U\mathbf{x}) = L\mathbf{y} = \mathbf{b}$$

An application of THEOREM SLEMM demonstrates that \mathbf{x} is a solution to $\mathcal{LS}(A, \mathbf{b})$.

Example 11. Here we illustrate the previous discussion, recycling the decomposition found previously in EXAMPLE TD4. Consider the linear system $\mathcal{LS}(A, \mathbf{b})$ with

$$A = \begin{bmatrix} -2 & 6 & -8 & 7 \\ -4 & 16 & -14 & 15 \\ -6 & 22 & -23 & 26 \\ -6 & 26 & -18 & 17 \end{bmatrix} \quad \mathbf{b} = \begin{bmatrix} -10 \\ -2 \\ -1 \\ -8 \end{bmatrix}$$

First we solve the system $\mathcal{LS}(L, \mathbf{b})$ (see EXAMPLE TD4 for L),

$$\begin{aligned} y_1 &= -10 \\ 2y_1 + y_2 &= -2 \\ 3y_1 + y_2 + y_3 &= -1 \\ 3y_1 + 2y_2 - 2y_3 + y_4 &= -8 \end{aligned}$$

Then

$$\begin{aligned} y_1 &= -10 \\ y_2 &= -2 - 2y_1 = -2 - 2(-10) = 18 \\ y_3 &= -1 - 3y_1 - y_2 = -1 - 3(-10) - 18 = 11 \\ y_4 &= -8 - 3y_1 - 2y_2 + 2y_3 = -8 - 3(-10) - 2(18) + 2(11) = 8 \end{aligned}$$

so

$$\mathbf{y} = \begin{bmatrix} -10 \\ 18 \\ 11 \\ 8 \end{bmatrix}$$

Then we solve the system $\mathcal{LS}(U, \mathbf{y})$ (see EXAMPLE TD4 for U),

$$\begin{aligned} -2x_1 + 6x_2 - 8x_3 + 7x_4 &= -10 \\ 4x_2 + 2x_3 + x_4 &= 18 \\ -x_3 + 4x_4 &= 11 \\ 2x_4 &= 8 \end{aligned}$$

Then

$$\begin{aligned} x_4 &= 8/2 = 4 \\ x_3 &= (11 - 4x_4)/(-1) = (11 - 4(4))/(-1) = 5 \\ x_2 &= (18 - 2x_3 - x_4)/4 = (18 - 2(5) - 4)/4 = 1 \\ x_1 &= (-10 - 6x_2 + 8x_3 - 7x_4)/(-2) = (-10 - 6(1) + 8(5) - 7(4))/(-2) = 2 \end{aligned}$$

And so

$$\mathbf{x} = \begin{bmatrix} 2 \\ 1 \\ 5 \\ 4 \end{bmatrix}$$

is the solution to $\mathcal{LS}(U, \mathbf{y})$ and consequently is the unique solution to $\mathcal{LS}(A, \mathbf{b})$, as you can easily verify.

2.1.3 Computing Triangular Decompositions

It would be a simple matter to adjust the algorithm for converting a matrix to reduced row-echelon form and obtain an algorithm to compute the triangular decomposition of the matrix, along the lines of EXAMPLE TD4 and the discussion preceding this example. However, it is possible to obtain relatively simple formulas for the entries of the decomposition, and if computed in the proper order, an implementation will be straightforward. We will state the result as a theorem and then give an example of its use.

Theorem 2.2. *Suppose that A is a square matrix of size n with a triangular decomposition $A = LU$, where L is lower triangular with diagonal entries all equal to 1, and U is upper triangular. Then*

$$\begin{aligned} [U]_{ij} &= [A]_{ij} - \sum_{k=1}^{i-1} [L]_{ik} [U]_{kj} & 1 \leq i \leq j \leq n \\ [L]_{ij} &= \frac{1}{[U]_{jj}} \left([A]_{ij} - \sum_{k=1}^{j-1} [L]_{ik} [U]_{kj} \right) & 1 \leq j < i \leq n \end{aligned}$$

Proof. Consider a single scalar product of an entry of L with an entry of U of the form $[L]_{ik} [U]_{kj}$. By DEFINITION LTM, if $k > i$ then $[L]_{ik} = 0$, while DEFINITION UTM, says that if $k > j$ then $[U]_{kj} = 0$. So we can combine these two facts to assert that if $k > \min(i, j)$, $[L]_{ik} [U]_{kj} = 0$ since at least one term of the product will be zero. Employing this observation,

$$\begin{aligned} [A]_{ij} &= \sum_{k=1}^n [L]_{ik} [U]_{kj} \\ &= \sum_{k=1}^{\min(i, j)} [L]_{ik} [U]_{kj} \end{aligned}$$

Now, assume that $1 \leq i \leq j \leq n$,

$$\begin{aligned}
[U]_{ij} &= [A]_{ij} - [A]_{ij} + [U]_{ij} \\
&= [A]_{ij} - \sum_{k=1}^{\min(i,j)} [L]_{ik} [U]_{kj} + [U]_{ij} \\
&= [A]_{ij} - \sum_{k=1}^i [L]_{ik} [U]_{kj} + [U]_{ij} \\
&= [A]_{ij} - \sum_{k=1}^{i-1} [L]_{ik} [U]_{kj} - [L]_{ii} [U]_{ij} + [U]_{ij} \\
&= [A]_{ij} - \sum_{k=1}^{i-1} [L]_{ik} [U]_{kj} - [U]_{ij} + [U]_{ij} \\
&= [A]_{ij} - \sum_{k=1}^{i-1} [L]_{ik} [U]_{kj}
\end{aligned}$$

And for $1 \leq j < i \leq n$,

$$\begin{aligned}
[L]_{ij} &= \frac{1}{[U]_{jj}} \left([L]_{ij} [U]_{jj} \right) \\
&= \frac{1}{[U]_{jj}} \left([A]_{ij} - [A]_{ij} + [L]_{ij} [U]_{jj} \right) \\
&= \frac{1}{[U]_{jj}} \left([A]_{ij} - \sum_{k=1}^{\min(i,j)} [L]_{ik} [U]_{kj} + [L]_{ij} [U]_{jj} \right) \\
&= \frac{1}{[U]_{jj}} \left([A]_{ij} - \sum_{k=1}^j [L]_{ik} [U]_{kj} + [L]_{ij} [U]_{jj} \right) \\
&= \frac{1}{[U]_{jj}} \left([A]_{ij} - \sum_{k=1}^{j-1} [L]_{ik} [U]_{kj} - [L]_{ij} [U]_{jj} + [L]_{ij} [U]_{jj} \right) \\
&= \frac{1}{[U]_{jj}} \left([A]_{ij} - \sum_{k=1}^{j-1} [L]_{ik} [U]_{kj} \right)
\end{aligned}$$

□

At first glance, these formulas may look exceedingly complex. Upon closer examination, it looks even worse. We have expressions for entries of U that depend on other entries of U and also on entries of L . But then the formula for entries of L depend on entries from L and entries from U . Do these formula have circular dependencies? Or perhaps equivalently, how do we get started? The key is to be organized about the computations and employ these two (similar) formulas in a specific order. First compute the first row of L , followed by the first column of U . Then the second row of L , followed by the second column of U . And so on. In this way, all of the values required for each new entry will have already been computed previously.

Of course, the formula for entries of L require division by diagonal entries of U . These entries might be zero, but in this case A is nonsingular and does not have a triangular decomposition. So we need not check the hypothesis carefully and can launch into the arithmetic dictated by the formulas, confident that we will

be reminded when a decomposition is not possible. Note that these formula give us all of the values that we need for the decomposition, since we require that L has 1's on the diagonal. If we replace the 1's on the diagonal of L by zeros, and add the matrix U , we get an $n \times n$ matrix containing all the information we need to resurrect the triangular decomposition. This is mostly a notational convenience, but it is a frequent way of presenting the information. We'll employ it in the next example.

Example 12. We illustrate the application of the formulas in THEOREM TDEE for the 6×6 matrix A .

$$A = \begin{bmatrix} 3 & 3 & -3 & -2 & -1 & 0 \\ -6 & -4 & 5 & 2 & 4 & 2 \\ 9 & 9 & -7 & -7 & 0 & 1 \\ -6 & -10 & 8 & 10 & -1 & -7 \\ 6 & 4 & -9 & -2 & -10 & 1 \\ 9 & 3 & -12 & -3 & -21 & -2 \end{bmatrix}$$

Using the notational convenience of packaging the two triangular matrices into one matrix, and using the ordering of the computations mentioned above, we display the results after computing a single row and column of each of the two triangular matrices.

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

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

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

Splitting out the pieces of this matrix, we have the decomposition,

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

2.1.4 Triangular Decomposition with Pivoting

The hypotheses of THEOREM TD can be weakened slightly to include matrices where not every A_k is nonsingular. This introduces a rearrangement of the rows and

columns of A to force as many as possible of the smaller submatrices to be nonsingular. Then permutation matrices also enter into the decomposition. We will not present the details here, but instead suggest consulting a more advanced text on matrix analysis.

2.2 QR (Gram-Schmidt) Decomposition

The **QR decomposition** begins with any matrix and describes it as a product of a unitary matrix (Q) and an upper triangular matrix (R). Hence the customary shorthand name, “QR”. If the LU decomposition is reminiscent of reduced row-echelon form, then the QR decomposition is reminiscent of the **Gram-Schmidt** process (see Subsection SUBSECTION 0.GSP).

2.2.1 QR Decomposition via Gram-Schmidt

The Gram-Schmidt procedure is based on THEOREM GSP. We begin with a set of linearly independent vectors and progressively convert them into a new set of nonzero vectors that form an orthogonal set, which can be easily converted to an orthonormal set. An orthonormal set is a wondrous thing, so the other portion of the conclusion is often overlooked. The new set of vectors spans the same subspace as the space spanned by the original set. This is half the reason that the Gram-Schmidt procedure is so useful.

As a preview of our main theorem, let’s convert this idea into the language of matrices for a special case. Let A be a square matrix of size n . Take the columns of A as a set of vectors, which will form a basis of \mathbb{C}^n . Apply the Gram-Schmidt procedure to this set of n vectors. The manufacture of the i -th vector of this new set is the i -th vector of the original set, added to a linear combination of vectors 1 through $i - 1$ of the new set. If we recursively unpack these linear combinations, we can express each new vector as a linear combination of vectors 1 through i of the original set, where the i -th vector has a coefficient of 1. Record the scalars of this linear combination in a column vector, whose last nonzero entry is 1. Make these column vectors the columns of a square matrix, R' , of size n . Define $Q = AR'$.

By the Gram-Schmidt procedure, the columns of Q are an orthogonal set of nonzero vectors, and so Q^*Q will be a diagonal matrix with nonzero entries. The matrix R' is square, upper-triangular, and each diagonal entry is 1. Hence R' is invertible, so let R denote the inverse, which is again upper-triangular with diagonal entries equal to 1. We then obtain $A = QR$. It is a simple matter to scale the columns of Q to form an orthonormal set, and the requisite scaling of the columns of R' will not impede the existence of R , though we can only claim diagonal entries are nonzero. In this way, we can claim that Q is a unitary matrix.

A QR decomposition can be created for any matrix —it need not be square and it need not have full rank. The matrix Q is unitary, and R is upper triangular. Thus, each column of A can be expressed as a linear combination of the columns of Q , which form an orthonormal basis. So the column space of A is spanned by an orthonormal subset of the columns of Q , giving us the essence of the Gram-Schmidt procedure without the hypothesis that our original set is linearly independent. For the statement of THEOREM GSP, it was a convenience to hypothesize that S is linearly independent. Can you examine the proof and see what changes are required if we lift this hypothesis?

We now state, and prove, a sequence of theorems which solidify the discussion above and generalize to rectangular matrices that may not have full rank.

Theorem 2.3. *Suppose that A is an $m \times n$ matrix of rank n . Then there exists an $m \times n$ matrix Q whose columns form an orthonormal set, and an upper-triangular matrix R of size n with positive diagonal entries, such that $A = QR$.*

Proof. Outline: Use Gram-Schmidt to successively build the columns. Scale each column by positive/negative norm to get positive diagonal entries in R . \square

Notice that the rank n condition of Theorem 2.3 necessarily implies $m \geq n$. We can expand Q and R simultaneously to get a decomposition where Q is a unitary matrix.

The column space of a matrix is an important property of a matrix. For example, the column space of the coefficient matrix of a system of equations is the set of vectors that can be used to form a consistent system when paired with the coefficient matrix (THEOREM CSCS). Not only does Q have a column space equal to that of A , the first i columns of Q are a basis for the space spanned by the first i columns of A , for $1 \leq i \leq n$.

Theorem 2.4. *Suppose that A is an $m \times n$ matrix of rank n . Then there exists a unitary matrix Q of size m and an upper-triangular $m \times n$ matrix R with positive diagonal entries such that $A = QR$.*

Proof. Begin with a decomposition $A = Q'R'$ as given by Theorem 2.3. Create the matrix Q by adding $m - n$ columns to Q' by the following process. Find a vector outside the span of the current set of columns (THEOREM ELIS). Apply the Gram-Schmidt procedure (THEOREM GSP) to the set of columns plus this one new vector. The current columns will be unaffected by an additional application of the Gram-Schmidt procedure (so in a practical application it is unnecessary). The one additional vector will be orthogonal to the others, and can be scaled to have norm 1. add this vector as a new column of Q , preserving the property that the columns are an orthonormal set.

For each of the $m - n$ new columns, add a new zero row to R' creating the $m \times n$ matrix R , with $QR = Q'R' = A$. \square

The decomposition of Theorem 2.3 is referred to as a **thin** decomposition, while the decomposition of Theorem 2.4 is referred to as a **full** decomposition.

What happens if A does not have full rank? Then there will be some relation of linear dependence on the columns of the matrix. In the course of working the Gram-Schmidt procedure, as in the construction given for the proof of Theorem 2.3, this will be discovered when the newly created vector is the zero vector. No matter, we simply add any vector as the next column of Q which will preserve the columns as an orthonormal set. This vector can be determined in a fashion entirely similar to the device used in the proof of Theorem 2.4. However, this nearly arbitrary choice outside the span of the current set of columns, requires that we add a row of zeros in R and lose our positive diagonal entry. The ultimate price for this is that certain uniqueness results ($\langle\langle$ uniqueness of QR $\rangle\rangle$) no longer hold.

For the most general case of a QR decomposition for a rectangular matrix of arbitrary rank, we could fashion a proof based on the discussion of the previous paragraph. However, while Gram-Schmidt provides a good theoretical grounding for the QR decomposition, in numerical contexts its performance is weak ($\langle\langle$ TB, Lecture 7 to 10 $\rangle\rangle$). Instead we give a constructive proof based on Householder reflections.

2.2.2 QR Decomposition via Householder Reflections

Householder reflections provide a useful tool for creating the zero entries required in each column of R , and will provide an algorithm with better performance (speed and accuracy), compared to the more intuitive approach of the Gram-Schmidt procedure.

Theorem 2.5. *Suppose that A is an $m \times n$ matrix. Then there exists a unitary matrix Q of size m and an upper-triangular $m \times n$ matrix R such that $A = QR$.*

Proof. We proceed by induction on n . When $n = 1$ form the Householder matrix, P , which will convert all of the entries of the lone column of A into zeros, except the first one (Theorem 1.16). Denote the resulting column vector as the matrix R , which is upper triangular. Define $Q = P^*$, which is unitary. Then $A = P^*PA = QR$.

Now consider general n . Let \hat{A} be the first $n - 1$ columns of A , so by induction there is a unitary matrix \hat{Q} of size m and an upper-triangular matrix \hat{R} providing a QR decomposition of \hat{A} . Partition \hat{R} into a square upper triangular matrix R_1 comprised of the first $n - 1$ rows of \hat{R} , leaving a second matrix with $m - n + 1$ rows and zero in every entry. Let \mathbf{v} denote the final column of A and compute $\mathbf{w} = \hat{Q}^*\mathbf{v}$. Partition \mathbf{w} into two pieces by denoting the first $n - 1$ entries as \mathbf{w}_1 , and entries n through m as \mathbf{w}_2 . Compute the Householder matrix, P , of size $m - n + 1$ that takes \mathbf{w}_2 , to a multiple of \mathbf{e}_1 (Theorem 1.16). We have all the pieces in place, so now observe:

$$\begin{aligned} \left[\begin{array}{c|c} I_{n-1} & 0 \\ \hline 0 & P \end{array} \right] \hat{Q}^*A &= \left[\begin{array}{c|c} I_{n-1} & 0 \\ \hline 0 & P \end{array} \right] \hat{Q}^* \left[\hat{A} \mid \mathbf{v} \right] \\ &= \left[\begin{array}{c|c} I_{n-1} & 0 \\ \hline 0 & P \end{array} \right] \left[\hat{Q}^*\hat{A} \mid \hat{Q}^*\mathbf{v} \right] \\ &= \left[\begin{array}{c|c} I_{n-1} & 0 \\ \hline 0 & P \end{array} \right] \left[\hat{R} \mid \mathbf{w} \right] \\ &= \left[\begin{array}{c|c} I_{n-1} & 0 \\ \hline 0 & P \end{array} \right] \left[\begin{array}{c|c} R_1 & \mathbf{w}_1 \\ \hline 0 & \mathbf{w}_2 \end{array} \right] \\ &= \left[\begin{array}{c|c} R_1 & \mathbf{w}_1 \\ \hline 0 & P\mathbf{w}_2 \end{array} \right] \end{aligned}$$

Notice that the first two matrices in these equations are unitary, and hence so is their product. Because of the action of the Householder matrix, the final matrix is upper triangular. If we move the unitary matrix to the other side of the equation, with an inverse (adjoint), we arrive at a QR decomposition of A and complete the induction step. \square

The inductive proof of Theorem 2.5 will automatically provide a recipe for a recursive procedure to compute a QR decomposition. But recursion is rarely, if ever, a good idea for efficiency. Fortunately, the proof suggests a better procedure. Work on A column-by-column, progressively using Householder matrices to “zero out” each column below the diagonal entry. The construction of such a Householder matrix will require a nonzero entry in the diagonal entry. A row swap, accomplished by a (unitary) permutation matrix, can move a nonzero entry from elsewhere in the column. Of course, if the whole remainder of the column is all zeros, then we can just move on to the next column.

Notice how the unitary matrices change over the course of these iterations. In later steps each unitary matrix has a larger identity matrix as the block in the upper-left corner of the matrix. So if a product of these matrices is required, it can

be more efficient to begin building up the product starting with the last unitary matrix. Of course, order matters in matrix multiplication, and maybe you need the product in the reverse order. No matter, compute the transpose of the desired matrix, which necessarily will reverse the order of the product. Once the product is concluded, then transpose the result.

Example 13. We illustrate the algorithm suggested by the proof of Theorem 2.5 on a 4×4 nonsingular matrix A . All of our computations were performed in Sage using algebraic numbers, so there were no approximations of irrational square roots as floating point numbers. However, we do give each matrix here with the final exact entries displayed as approximations with a limited number of places. At step i , $1 \leq i \leq 3$, we display the unitary Householder matrix, Q_i , and the partially upper triangular matrix $R_i = Q_i \dots Q_1 A$.

$$A = \begin{bmatrix} 4 & -5 & -7 & -4 \\ 1 & -1 & -1 & -2 \\ 0 & 0 & 1 & -1 \\ -1 & 5 & 8 & -8 \end{bmatrix}$$

$$\begin{array}{l} i = 1 \\ i = 2 \\ i = 3 \end{array} \quad \begin{array}{l} Q_1 = \\ Q_2 = \\ Q_3 = \end{array} \begin{array}{l} \begin{bmatrix} 0.94 & 0.24 & 0.0 & -0.24 \\ 0.24 & 0.03 & 0.0 & 0.97 \\ 0.0 & 0.0 & 1.0 & 0.0 \\ -0.24 & 0.97 & 0.0 & 0.03 \end{bmatrix} \\ \begin{bmatrix} 1.0 & 0.0 & 0.0 & 0.0 \\ 0.0 & 0.99 & 0.0 & 0.01 \\ 0.0 & 0.0 & 1.0 & 0.0 \\ 0.0 & 0.01 & 0.0 & -0.99 \end{bmatrix} \\ \begin{bmatrix} 1.0 & 0.0 & 0.0 & 0.0 \\ 0.0 & 1.0 & 0.0 & 0.0 \\ 0.0 & 0.0 & 0.95 & -0.31 \\ 0.0 & 0.0 & -0.31 & -0.95 \end{bmatrix} \end{array} \quad \begin{array}{l} R_1 = \\ R_2 = \\ R_3 = \end{array} \begin{array}{l} \begin{bmatrix} 4.24 & -6.13 & -8.72 & -2.36 \\ 0.0 & 3.65 & 6.09 & -8.77 \\ 0.0 & 0.0 & 1.0 & -1.0 \\ 0.0 & 0.35 & 0.91 & -1.23 \end{bmatrix} \\ \begin{bmatrix} 4.24 & -6.13 & -8.72 & -2.36 \\ 0.0 & 3.67 & 6.15 & -8.85 \\ 0.0 & 0.0 & 1.0 & -1.0 \\ 0.0 & 0.0 & -0.32 & 0.39 \end{bmatrix} \\ \begin{bmatrix} 4.24 & -6.13 & -8.72 & -2.36 \\ 0.0 & 3.67 & 6.15 & -8.85 \\ 0.0 & 0.0 & 1.05 & -1.07 \\ 0.0 & 0.0 & 0.0 & -0.06 \end{bmatrix} \end{array}$$

So,

$$R = R_3 = \begin{bmatrix} 4.2426 & -6.1283 & -8.721 & -2.357 \\ 0.0 & 3.6667 & 6.1515 & -8.8485 \\ 0.0 & 0.0 & 1.0504 & -1.0701 \\ 0.0 & 0.0 & 0.0 & -0.0612 \end{bmatrix}$$

and then $R = Q_3 Q_2 Q_1 A$, so

$$Q = Q_1^* Q_2^* Q_3^* = \begin{bmatrix} 0.9428 & 0.2121 & -0.0787 & -0.2448 \\ 0.2357 & 0.1212 & 0.2951 & 0.918 \\ 0.0 & 0.0 & 0.952 & -0.306 \\ -0.2357 & 0.9697 & -0.0197 & -0.0612 \end{bmatrix}$$

Notice how the product for Q involves progressively simpler matrices (bigger identity matrix blocks) moving from right to left.

2.2.3 Solving Systems with a QR Decomposition

Consider the problem of solving a linear system $A\mathbf{x} = \mathbf{b}$. Replace A by a QR decomposition, to obtain $QR\mathbf{x} = \mathbf{b}$. Now, the inverse of the unitary matrix Q is its adjoint, so we have the new system $R\mathbf{x} = Q^*\mathbf{b}$. Since R is upper-triangular, the new

system can be solved via back-solving, much as we did with an LU decomposition (Subsection 2.1.2).

A primary advantage of this approach is that a unitary matrix is generally well-behaved in numerical computations. Since the columns have unit norm, the entries can never be too large. Many of the problems of floating-point arithmetic come from combining numbers of grossly different magnitudes, and this is avoided in computations with a unitary matrix, such as the matrix-vector product in the rearrangement of a linear system in the previous paragraph.

What about the number of operations required? For a QR decomposition of an $m \times n$ matrix via the Gram-Schmidt procedure the operation count is $\sim 2mn^2$. When $m \geq n$ the determination of R can be accomplished in $\sim 2mn^2 - \frac{2}{3}n^3$ operations when using a sequence of Householder reflections ([?]). This includes storing the Householder vector for each iteration, but not the computation of the Householder matrix itself, or the accumulated product of all the Householder matrices into one grand unitary matrix Q . Notice that for a square matrix, when $m = n$, the count is $\sim \frac{4}{3}n^3$, which is twice the cost of an LU factorization

If we use Householder reflections, then we must decide what we want to do with Householder matrices. If we are solving a linear system, we can successively multiply each new Householder matrix times the column vector on the right hand side of the equation. Thus, we modify both sides of the equation in the same way, producing equivalent systems as we go. In practice, the product of a column vector by a Householder matrix can be accomplished very efficiently, and not by explicitly forming the matrix. (See Exercise 11.) It should *never* be necessary to explicitly form a Householder matrix, but instead the Householder vector should be enough information to perform whatever computation is at hand.

2.2.4 Uniqueness of QR Decompositions

Generally, when A has full rank and we require that diagonal entries of R be positive, we get a unique QR decomposition. Proving this will be easier once we learn about Cholesky decompositions (Section <<a section about Cholesky decompositions>>), so we will defer a proof until then.

2.2.5 Final Thoughts

Exercise 18. Suppose that A is a nonsingular matrix with *real* entries. Then A has a QR decomposition where (1) $A = QR$, (2) Q is unitary, and (3) R is upper triangular with positive entries, (4) Q and R are also matrices with real entries. Prove that this decomposition is unique. (So you may assume such a decomposition exists, you are just being asked to establish uniqueness.)

Solution. Assume there are two such decompositions, so $R_1Q_1 = A = R_2Q_2$ and rearrange to obtain $R_2^{-1}R_1 = Q_2Q_1^*$. The left-hand side will be an upper triangular matrix with positive diagonal entries. The right-hand side is a product of two unitary matrices, hence is unitary, and its columns form an orthonormal set. Now, viewing the columns of the left-hand matrix as an orthonormal set will allow you to progressively conclude that the columns (from left to right) are the columns of an identity matrix. Thus, $R_2^{-1}R_1 = I$ and $R_1 = R_2$, and similarly, $Q_1 = Q_2$.

2.3 Singular Value Decomposition

The singular value decomposition is one of the more useful ways to represent any matrix, even rectangular ones. We can also view the singular values of a (rectan-

gular) matrix as analogues of the eigenvalues of a square matrix.

2.3.1 Matrix-Adjoint Products

Our definitions and theorems in this section rely heavily on the properties of the matrix-adjoint products (A^*A and AA^*). We start by examining some of the basic properties of these two positive semi-definite matrices. Now would be a good time to review the basic facts about positive semi-definite matrices in Section 1.8.

Theorem 2.6. *Suppose that A is an $m \times n$ matrix and A^*A has rank r . Let $\lambda_1, \lambda_2, \lambda_3, \dots, \lambda_p$ be the nonzero distinct eigenvalues of A^*A and let $\rho_1, \rho_2, \rho_3, \dots, \rho_q$ be the nonzero distinct eigenvalues of AA^* . Then,*

1. $p = q$.
2. The distinct nonzero eigenvalues can be ordered such that $\lambda_i = \rho_i, 1 \leq i \leq p$.
3. Properly ordered, the algebraic multiplicities of the nonzero eigenvalues are identical, $\alpha_{A^*A}(\lambda_i) = \alpha_{AA^*}(\rho_i), 1 \leq i \leq p$.
4. The rank of A^*A is equal to the rank of AA^* .
5. There is an orthonormal basis, $\{\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, \dots, \mathbf{x}_n\}$ of \mathbb{C}^n composed of eigenvectors of A^*A and an orthonormal basis, $\{\mathbf{y}_1, \mathbf{y}_2, \mathbf{y}_3, \dots, \mathbf{y}_m\}$ of \mathbb{C}^m composed of eigenvectors of AA^* with the following properties. Order the eigenvectors so that $\mathbf{x}_i, r+1 \leq i \leq n$ are the eigenvectors of A^*A for the zero eigenvalue. Let $\delta_i, 1 \leq i \leq r$ denote the nonzero eigenvalues of A^*A . Then $A\mathbf{x}_i = \sqrt{\delta_i}\mathbf{y}_i, 1 \leq i \leq r$ and $A\mathbf{x}_i = \mathbf{0}, r+1 \leq i \leq n$. Finally, $\mathbf{y}_i, r+1 \leq i \leq m$, are eigenvectors of AA^* for the zero eigenvalue.

Proof. Suppose that $\mathbf{x} \in \mathbb{C}^n$ is any eigenvector of A^*A for a nonzero eigenvalue λ . We will show that $A\mathbf{x}$ is an eigenvector of AA^* for the same eigenvalue, λ . First, we ascertain that $A\mathbf{x}$ is not the zero vector.

$$\langle A\mathbf{x}, A\mathbf{x} \rangle = \langle \mathbf{x}, A^*A\mathbf{x} \rangle = \langle \mathbf{x}, \lambda\mathbf{x} \rangle = \lambda \langle \mathbf{x}, \mathbf{x} \rangle$$

Since \mathbf{x} is an eigenvector, $\mathbf{x} \neq \mathbf{0}$, and by THEOREM PIP, $\langle \mathbf{x}, \mathbf{x} \rangle \neq 0$. As λ was assumed to be nonzero, we see that $\langle A\mathbf{x}, A\mathbf{x} \rangle \neq 0$. Again, THEOREM PIP tells us that $A\mathbf{x} \neq \mathbf{0}$.

Much of the sequel turns on the following simple computation. If you ever wonder what all the fuss is about adjoints, Hermitian matrices, square roots, and singular values, return to this brief computation, as it holds the key. There is much more to do in this proof, but after this it is mostly bookkeeping. Here we go. We check that $A\mathbf{x}$ functions as an eigenvector of AA^* for the eigenvalue λ ,

$$(AA^*)A\mathbf{x} = A\lambda\mathbf{x} = \lambda(A\mathbf{x})$$

That's it. If \mathbf{x} is an eigenvector of A^*A , for a *nonzero* eigenvalue, then $A\mathbf{x}$ is an eigenvector for AA^* for the same nonzero eigenvalue. Let's see what this buys us.

A^*A and AA^* are Hermitian matrices (DEFINITION HM), and hence are normal (Definition 10). This provides the existence of orthonormal bases of eigenvectors for each matrix by THEOREM OBNM. Also, since each matrix is diagonalizable (DEFINITION DZM) by THEOREM OD the algebraic and geometric multiplicities of each eigenvalue (zero or not) are equal by THEOREM DMFE.

Our first step is to establish that a nonzero eigenvalue λ has the same geometric multiplicity for both A^*A and AA^* . Suppose $\{\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, \dots, \mathbf{x}_s\}$ is an

orthonormal basis of eigenvectors of A^*A for the eigenspace $\mathcal{E}_{A^*A}(\lambda)$. Then for $1 \leq i < j \leq s$,

$$\langle A\mathbf{x}_i, A\mathbf{x}_j \rangle = \langle \mathbf{x}_i, A^*A\mathbf{x}_j \rangle = \langle \mathbf{x}_i, \lambda\mathbf{x}_j \rangle = \lambda \langle \mathbf{x}_i, \mathbf{x}_j \rangle = \lambda(0)$$

So the set $E = \{A\mathbf{x}_1, A\mathbf{x}_2, A\mathbf{x}_3, \dots, A\mathbf{x}_s\}$ is an orthogonal set of nonzero eigenvectors of AA^* for the eigenvalue λ . By THEOREM OS LI, the set E is linearly independent and so the geometric multiplicity of λ as an eigenvalue of AA^* is s or greater. We have

$$\alpha_{A^*A}(\lambda) = \gamma_{A^*A}(\lambda) \leq \gamma_{AA^*}(\lambda) = \alpha_{AA^*}(\lambda)$$

This inequality applies to any matrix for any of its nonzero eigenvalues. We now apply it to the matrix A^* ,

$$\alpha_{AA^*}(\lambda) = \alpha_{(A^*)^*A^*}(\lambda) \leq \alpha_{A^*(A^*)^*}(\lambda) = \alpha_{A^*A}(\lambda)$$

With the twin inequalities, we see that the four multiplicities of a common nonzero eigenvalue of A^*A and AA^* are all equal. This is enough to establish that $p = q$, since we cannot have an eigenvalue of one of the matrix-adjoint products along with a zero algebraic multiplicity for the other matrix-adjoint product. Then the eigenvalues can be ordered such that $\lambda_i = \rho_i$ for $1 \leq i \leq p$.

For any matrix, the null space is identical to the eigenspace of the zero eigenvalue, and thus the nullity of the matrix is equal to the geometric multiplicity of the zero eigenvalue. As our matrix-adjoint products are diagonalizable, the nullity is equal to the algebraic multiplicity of the zero eigenvalue. The algebraic multiplicities of all the eigenvalues sum to the size of the matrix (THEOREM NEM), as does the rank and nullity (THEOREM RPNC). So the rank of our matrix-adjoint products is equal to the sum of the algebraic multiplicities of the nonzero eigenvalues. So the ranks of A^*A and AA^* are equal,

$$r(A^*A) = \sum_{i=1}^p \alpha_{A^*A}(\lambda_i) = \sum_{i=1}^q \alpha_{AA^*}(\rho_i) = r(AA^*)$$

When A is rectangular, the square matrices A^*A and AA^* have different sizes. With equal algebraic and geometric multiplicities for their common nonzero eigenvalues, the difference in their sizes is manifest in different algebraic multiplicities for the zero eigenvalue and different nullities. Specifically,

$$n(A^*A) = n - r \qquad n(AA^*) = m - r$$

Suppose that $\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, \dots, \mathbf{x}_n$ is an orthonormal basis of \mathbb{C}^n composed of eigenvectors of A^*A and ordered so that $\mathbf{x}_i, r+1 \leq i \leq n$ are eigenvectors of A^*A for the zero eigenvalue. Denote the associated nonzero eigenvalues of A^*A for these eigenvectors by $\delta_i, 1 \leq i \leq r$. Then define

$$\mathbf{y}_i = \frac{1}{\sqrt{\delta_i}} A\mathbf{x}_i, \quad 1 \leq i \leq r$$

Let $\mathbf{y}_{r+1}, \mathbf{y}_{r+2}, \mathbf{y}_{r+3}, \dots, \mathbf{y}_m$ be an orthonormal basis for the eigenspace $\mathcal{E}_{AA^*}(0)$, whose existence is guaranteed by the Gram-Schmidt process (THEOREM GSP). As scalar multiples of demonstrated eigenvectors of AA^* , $\mathbf{y}_i, 1 \leq i \leq r$ are also eigenvectors of AA^* , and $\mathbf{y}_i, r+1 \leq i \leq m$ have been chosen as eigenvectors of AA^* . These eigenvectors also have norm 1, as we now show.

For $1 \leq i \leq r$,

$$\begin{aligned}\|\mathbf{y}_i\|^2 &= \left\| \frac{1}{\sqrt{\delta_i}} A\mathbf{x}_i \right\|^2 = \left\langle \frac{1}{\sqrt{\delta_i}} A\mathbf{x}_i, \frac{1}{\sqrt{\delta_i}} A\mathbf{x}_i \right\rangle \\ &= \frac{1}{\sqrt{\delta_i}} \frac{1}{\sqrt{\delta_i}} \langle A\mathbf{x}_i, A\mathbf{x}_i \rangle = \frac{1}{\delta_i} \langle A\mathbf{x}_i, A\mathbf{x}_i \rangle \\ &= \frac{1}{\delta_i} \langle \mathbf{x}_i, A^* A\mathbf{x}_i \rangle = \frac{1}{\delta_i} \langle \mathbf{x}_i, \delta_i \mathbf{x}_i \rangle \\ &= \frac{1}{\delta_i} \delta_i \langle \mathbf{x}_i, \mathbf{x}_i \rangle = 1\end{aligned}$$

For $r+1 \leq i \leq n$, the \mathbf{y}_i have been chosen to have norm 1.

Finally we check orthogonality. Consider two eigenvectors \mathbf{y}_i and \mathbf{y}_j with $1 \leq i < j \leq m$. If these two vectors have different eigenvalues, then THEOREM HMOE establishes that the two eigenvectors are orthogonal. If the two eigenvectors have a zero eigenvalue, then they are orthogonal by the choice of the orthonormal basis of $\mathcal{E}_{AA^*}(0)$. If the two eigenvectors have identical eigenvalues, which are nonzero, then

$$\begin{aligned}\langle \mathbf{y}_i, \mathbf{y}_j \rangle &= \left\langle \frac{1}{\sqrt{\delta_i}} A\mathbf{x}_i, \frac{1}{\sqrt{\delta_j}} A\mathbf{x}_j \right\rangle = \frac{\overline{1}}{\sqrt{\delta_i}} \frac{1}{\sqrt{\delta_j}} \langle A\mathbf{x}_i, A\mathbf{x}_j \rangle \\ &= \frac{1}{\sqrt{\delta_i \delta_j}} \langle A\mathbf{x}_i, A\mathbf{x}_j \rangle = \frac{1}{\sqrt{\delta_i \delta_j}} \langle \mathbf{x}_i, A^* A\mathbf{x}_j \rangle \\ &= \frac{1}{\sqrt{\delta_i \delta_j}} \langle \mathbf{x}_i, \delta_j \mathbf{x}_j \rangle = \frac{\delta_j}{\sqrt{\delta_i \delta_j}} \langle \mathbf{x}_i, \mathbf{x}_j \rangle \\ &= \frac{\delta_j}{\sqrt{\delta_i \delta_j}} (0) = 0\end{aligned}$$

So $\{\mathbf{y}_1, \mathbf{y}_2, \mathbf{y}_3, \dots, \mathbf{y}_m\}$ is an orthonormal set of eigenvectors for AA^* . The critical relationship between these two orthonormal bases is present by design. For $1 \leq i \leq r$,

$$A\mathbf{x}_i = \sqrt{\delta_i} \frac{1}{\sqrt{\delta_i}} A\mathbf{x}_i = \sqrt{\delta_i} \mathbf{y}_i$$

For $r+1 \leq i \leq n$ we have

$$\langle A\mathbf{x}_i, A\mathbf{x}_i \rangle = \langle \mathbf{x}_i, A^* A\mathbf{x}_i \rangle = \langle \mathbf{x}_i, \mathbf{0} \rangle = 0$$

So by THEOREM PIP, $A\mathbf{x}_i = \mathbf{0}$. □

2.3.2 Singular Value Decomposition

The square roots of the eigenvalues of A^*A (or almost equivalently, AA^* !) are known as the singular values of A . Here is the definition.

Definition 12. Suppose A is an $m \times n$ matrix. If the eigenvalues of A^*A are $\delta_1, \delta_2, \delta_3, \dots, \delta_n$, then the **singular values** of A are

$$\sqrt{\delta_1}, \sqrt{\delta_2}, \sqrt{\delta_3}, \dots, \sqrt{\delta_n}$$

Theorem 2.6 is a total setup for the singular value decomposition. This remarkable theorem says that *any* matrix can be broken into a product of three matrices. Two are square and unitary. In light of THEOREM UMPIP, we can view these matrices as transforming vectors or coordinates in a rotational fashion. The middle

matrix of this decomposition is rectangular, but is as close to being diagonal as a rectangular matrix can be. Viewed as a transformation, this matrix effects, reflections, contractions, or expansions along axes—it stretches vectors. So any matrix, viewed as a geometric transformation is the product of a rotation, a stretch and a rotation.

The singular value theorem can also be viewed as an application of our most general statement about matrix representations of linear transformations relative to different bases. THEOREM MRCB concerns linear transformations $T: U \rightarrow V$ where U and V are possibly different vector spaces. When U and V have different dimensions, the resulting matrix representation will be rectangular. In SECTION CB we quickly specialized to the case where $U = V$, and the matrix representations are square, with one of our most central results, THEOREM SCB. Theorem 2.7, next, is an application of the full generality of THEOREM MRCB where the relevant bases are now orthonormal sets.

Theorem 2.7. *Suppose A is an $m \times n$ matrix of rank r with nonzero singular values $s_1, s_2, s_3, \dots, s_r$. Then $A = USV^*$ where U is a unitary matrix of size m , V is a unitary matrix of size n and S is an $m \times n$ matrix given by*

$$[S]_{ij} = \begin{cases} s_i & \text{if } 1 \leq i = j \leq r \\ 0 & \text{otherwise} \end{cases}$$

Proof. Let $\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, \dots, \mathbf{x}_n$ and $\mathbf{y}_1, \mathbf{y}_2, \mathbf{y}_3, \dots, \mathbf{y}_m$ be the orthonormal bases described by the conclusion of Theorem 2.6. Define U to be the $m \times m$ matrix whose columns are \mathbf{y}_i , $1 \leq i \leq m$, and define V to be the $n \times n$ matrix whose columns are \mathbf{x}_i , $1 \leq i \leq n$. With orthonormal sets of columns, both U and V are unitary matrices by THEOREM CUMOS.

Then for $1 \leq i \leq m$, $1 \leq j \leq n$,

$$\begin{aligned} [AV]_{ij} &= [A\mathbf{x}_j]_i = \left[\sqrt{\delta_j} \mathbf{y}_j \right]_i = s_j [\mathbf{y}_j]_i \\ &= [U]_{ij} [S]_{jj} = \sum_{k=1}^m [U]_{ik} [S]_{kj} = [US]_{ij} \end{aligned}$$

So by THEOREM ME, AV and US are equal matrices, $AV = US$. V is unitary, so applying its inverse yields the decomposition $A = USV^*$. \square

Typically, the singular values of a matrix are ordered from largest to smallest, so this convention is used for the diagonal elements of the matrix S in the decomposition, and then the columns of U and V will be ordered similarly.

2.3.3 Visualizing the SVD

It is helpful to think of the orthonormal bases that are the columns of U and V as “coordinate systems,” since they are pairwise “perpendicular” unit vectors, like the $\vec{i}, \vec{j}, \vec{k}$ often used in describing the geometry of space. We would then call each basis vector an “axis”.

Now think of an $m \times n$ matrix as a function from \mathbb{C}^n to \mathbb{C}^m . For an input vector $\mathbf{w} \in \mathbb{C}^n$, we have the output vector $\mathbf{y} = A\mathbf{w} \in \mathbb{C}^m$. If we write the output vector using the SVD decomposition of A as $A\mathbf{w} = USV^*\mathbf{w}$ we can consider the output as a three-step process that is more formally a composition of three linear transformations. Recall that unitary matrices preserve inner products, and thus preserve norms (length) and relative positions of two vectors (the “angle” between vectors), which is why unitary matrices are sometimes called “isometries” (THEOREM UMPIP).

1. V^* is the inverse of V so it will take any basis vector \mathbf{x}_i to a column of the identity matrix \mathbf{e}_i , an element of the standard basis of \mathbb{C}^n , or an axis of the “usual” coordinate system. Any other vector, such as our \mathbf{w} , will have its length preserved, but changing its position, though its position relative to the axes is unchanged.
2. S converts the “rotated” \mathbf{w} from \mathbb{C}^n to \mathbb{C}^m . But it does so in a very simple way. It simply scales each entry of the vector by a positive amount using a different singular value for each entry. If $m > n$, then the extra entries are simply new zeros. If $m < n$, then some entries get discarded.
3. U will convert the standard basis vectors (the usual axes) to the new orthonormal basis given by the \mathbf{y}_i . The twice-transformed version of \mathbf{w} will have its length preserved, but change position, though its position relative to the axes is unchanged.

So, *every* matrix is a rotation, a stretch, and a rotation. That is a simple, but accurate, understanding of what the SVD tells us.

Here is another look at the same idea. Consider the columns of U and V again as the axes of new coordinate systems. Then their adjoints are their inverses and each take the new axes to the standard unit vectors (columns of the identity matrix), the axes of the usual coordinate system. Consider an input vector \mathbf{w} , and its output $\mathbf{y} = A\mathbf{w}$, relative to these new bases and convert each to the standard coordinate system, $\mathbf{w}' = V^*\mathbf{w}$ and $\mathbf{y}' = U^*\mathbf{y}$. Then

$$\mathbf{y}' = U^*\mathbf{y} = U^*A\mathbf{w} = U^*USV^*\mathbf{w} = S\mathbf{y}'$$

In the “primed” spaces, the action of A is very simple. Or in other words, every linear transformation has a matrix representation that is basically diagonal, if only we pick the right bases for the domain and codomain.

2.3.4 Properties of the SVD

The appeal of the singular value decomposition is two-fold. First it is applicable to *any* matrix. That should be obvious enough. Second, components of the SVD provide a wealth of information about a matrix, and in the case of numerical matrices, they are well-behaved. In this subsection we collect various theorems about the SVD, and explore the consequences in a section about applications, Section [\(section-applications-of-SVD\)](#).

The SVD gives a decomposition of a matrix of a sum of rank one matrices, and the magnitude of the singular values tells us which of these rank one matrices is the most “important”.

Theorem 2.8. *Suppose the singular value decomposition of an $m \times n$ matrix A is given by $A = USV^*$, where the nonzero singular values in the first r entries of the diagonal of S are $s_1, s_2, s_3, \dots, s_r$, and the columns of U and V are, respectively, $\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, \dots, \mathbf{x}_n$ and $\mathbf{y}_1, \mathbf{y}_2, \mathbf{y}_3, \dots, \mathbf{y}_m$. Then*

$$A = \sum_{i=1}^r s_i \mathbf{x}_i \mathbf{y}_i^*$$

Proof. As usual, let \mathbf{e}_i be column i of the identity matrix I_m . Define S_i , for $1 \leq i \leq r$ to be the $m \times n$ matrix where every entry is zero, except column i is $s_i \mathbf{e}_i$. Then,

by design, $S = \sum_{i=1}^r S_i$. We have,

$$\begin{aligned} A &= USV^* = U \sum_{i=1}^r S_i V^* = \sum_{i=1}^r US_i V^* \\ &= \sum_{i=1}^r U [\mathbf{0} \dots |s_i \mathbf{e}_i| \dots | \mathbf{0}] V^* = \sum_{i=1}^r s_i [\mathbf{0} \dots |U \mathbf{e}_i| \dots | \mathbf{0}] V^* \\ &= \sum_{i=1}^r s_i [\mathbf{0} \dots | \mathbf{x}_i | \dots | \mathbf{0}] V^* = \sum_{i=1}^r s_i \mathbf{x}_i \mathbf{y}_i^* \end{aligned}$$

□

Be sure to recognize $\mathbf{x}_i \mathbf{y}_i^*$ as the outer product, an $m \times n$ matrix of rank one (every row is a multiple of every other row, and similarly for columns). See Subsection ?? for a good example of the utility of this result.

2.4 Cholesky Decomposition

An LU decomposition of a matrix is obtained by repeated row operations and produces a result with some symmetry of sorts. The “L” matrix is lower triangular and the “U” matrix is upper triangular, so $[L]_{ij} = 0 = \overline{[U]_{ji}}$ for $i < j$, which should be reminiscent of the definition of the adjoint of a matrix (DEFINITION AM). If we begin with a positive definite matrix, then we can do better. By beginning with a Hermitian matrix, we can do row operations, *and* identical column operations and maintain the symmetry of the entries. We arrive at a decomposition of the form U^*U , where U is upper-triangular.

2.4.1 The Cholesky Decomposition

Recall that a Hermitian matrix A is positive definite if $\langle \mathbf{x}, A\mathbf{x} \rangle > 0$ for all $\mathbf{x} \neq \mathbf{0}$. This is just the variant of positive semi-definiteness (Definition 11) where we replace the inequality by a strict inequality.

Theorem 2.9. *Suppose that A is a positive definite matrix. Then there exists a unique upper triangular matrix, U , with positive diagonal matrices such that $A = U^*U$.*

Proof. Coming soon. Algorithm below contains the essential ideas. Uniqueness is an exercise. □

Exercise 19. Prove that the upper triangular matrix U in the conclusion of Theorem 2.9 is unique.

2.4.2 Computing a Cholesky Decomposition

To create an LU decomposition, we used row operations to “zero out” entries below the diagonal of a matrix A . If we represented these row operations as elementary matrices, we could accumulate their net effect in a lower triangular matrix that operates on the left of the matrix. For a Cholesky decomposition, we do the same thing, but also perform the analogous column operation, which can be represented as the adjoint of the same elementary matrix, and then applied from the right.

Here is the same idea, but expressed as intermediate steps leading to the eventual Cholesky decomposition. Recall that Hermitian matrices necessarily have real diagonal entries. Suppose that A is an $n \times n$ positive definite matrix.

$$\begin{aligned} A &= \left[\begin{array}{c|c} a & \mathbf{y}^* \\ \mathbf{y} & B \end{array} \right] \\ &= \left[\begin{array}{c|c} \sqrt{a} & \mathbf{0}^* \\ \frac{1}{\sqrt{a}}\mathbf{y} & I \end{array} \right] \left[\begin{array}{c|c} 1 & \mathbf{0}^* \\ \mathbf{0} & B - \frac{1}{a}\mathbf{y}\mathbf{y}^* \end{array} \right] \left[\begin{array}{c|c} \sqrt{a} & \frac{1}{\sqrt{a}}\mathbf{y}^* \\ \mathbf{0} & I \end{array} \right] \\ &= U_1^* A_1 U_1 \end{aligned}$$

The only obstacle to this computation is the square root of the entry in the top left corner of A , and the result should be positive. If we apply the positive definite condition, with $\mathbf{x} = \mathbf{e}_1$ (the first column of the identity matrix) then we have

$$a = \langle \mathbf{e}_1, A\mathbf{e}_1 \rangle > 0$$

Can we repeat this decomposition on the $(n-1) \times (n-1)$ matrix $B - \frac{1}{a}\mathbf{y}\mathbf{y}^*$? As before we just need a strictly positive entry in the upper left corner of this slightly smaller matrix. Similar to before, employ the positive definite condition for A using $\mathbf{x} = U_1^{-1}\mathbf{e}_2$ and employ the version of A defining A_1 (see Exercise 20). What is the result after n iterations?

$$A = U_n^* \dots U_2^* U_1^* I U_1 U_2 \dots U_n = U^* U$$

Here we have used the observation that a product of upper triangular matrices is again upper triangular, and you should notice the appearance of the positive diagonal entries. So we have our desired factorization.

Exercise 20. In the discussion of a recursive algorithm for computing a Cholesky decomposition in Section 2.4.2, verify that the matrix A_1 has a strictly positive value in the second diagonal entry.

Chapter 3

Canonical Forms

You will know that some matrices are diagonalizable and some are not. Some authors refer to a non-diagonalizable matrix as **defective**, but we will study them carefully anyway. Examples of such matrices include EXAMPLE EMMS4, EXAMPLE HMEM5, and EXAMPLE CEMS6. Each of these matrices has at least one eigenvalue with geometric multiplicity strictly less than its algebraic multiplicity, and therefore THEOREM DMFE tells us these matrices are not diagonalizable.

Given a square matrix A , it is likely similar to many, many other matrices. Of all these possibilities, which is the best? “Best” is a subjective term, but we might agree that a diagonal matrix is certainly a very nice choice. Unfortunately, as we have seen, this will not always be possible. What form of a matrix is “next-best”? Our goal, which will take us several sections to reach, is to show that every matrix is similar to a matrix that is “nearly-diagonal” (Section 3.3). More precisely, every matrix is similar to a matrix with elements on the diagonal, and zeros and ones on the diagonal just above the main diagonal (the “super diagonal”), with zeros everywhere else. In the language of equivalence relations (see THEOREM SER), we are determining a systematic representative for each equivalence class, where the relation is similarity. Such a representative for a set of similar matrices is called a **canonical form**.

We have just discussed the determination of a canonical form as a question about matrices. However, we know that every square matrix creates a natural linear transformation (THEOREM MBLT) and every linear transformation with identical domain and codomain has a square matrix representation for each choice of a basis, with a change of basis creating a similarity transformation (THEOREM SCB). So we will state, and prove, theorems using the language of linear transformations on abstract vector spaces, while most of our examples will work with square matrices. You can, and should, mentally translate between the two settings frequently and easily.

3.1 Generalized Eigenspaces

In this section we will define a new type of invariant subspace and explore its key properties. This generalization of eigenvalues and eigenspaces will allow us to move from diagonal matrix representations of diagonalizable matrices to nearly diagonal matrix representations of arbitrary matrices.

3.1.1 Kernels of Powers of Linear Transformations

With Section $\langle\langle$ section-jordan-canonical-form $\rangle\rangle$ as our goal, we will become increasingly interested in kernels of powers of linear transformations, which will go a long way to helping us understand the structure of a linear transformation, or its matrix representation. We need the next theorem to help us understand generalized eigenspaces, though its specialization in Theorem $\langle\langle$ theorem-about-powers-nilpotent $\rangle\rangle$ to nilpotent linear transformations will be the real workhorse.

Theorem 3.1. *Suppose $T: V \rightarrow V$ is a linear transformation, where $\dim(V) = n$. Then there is an integer m , $0 \leq m \leq n$, such that*

$$\{\mathbf{0}\} = \mathcal{K}(T^0) \subsetneq \mathcal{K}(T^1) \subsetneq \mathcal{K}(T^2) \subsetneq \cdots \subsetneq \mathcal{K}(T^m) = \mathcal{K}(T^{m+1}) = \mathcal{K}(T^{m+2}) = \cdots$$

Proof. There are several items to verify in the conclusion as stated. First, we show that $\mathcal{K}(T^k) \subseteq \mathcal{K}(T^{k+1})$ for any k . Choose $\mathbf{z} \in \mathcal{K}(T^k)$. Then

$$T^{k+1}(\mathbf{z}) = T(T^k(\mathbf{z})) = T(\mathbf{0}) = \mathbf{0}$$

so $\mathbf{z} \in \mathcal{K}(T^{k+1})$

Second, we demonstrate the existence of a power m where consecutive powers result in equal kernels. A by-product will be the condition that m can be chosen so that $m \leq n$. To the contrary, suppose that

$$\{\mathbf{0}\} = \mathcal{K}(T^0) \subsetneq \mathcal{K}(T^1) \subsetneq \mathcal{K}(T^2) \subsetneq \cdots \subsetneq \mathcal{K}(T^{n-1}) \subsetneq \mathcal{K}(T^n) \subsetneq \mathcal{K}(T^{n+1}) \subsetneq \cdots$$

Since $\mathcal{K}(T^k) \subsetneq \mathcal{K}(T^{k+1})$, THEOREM PSSD implies that $\dim(\mathcal{K}(T^{k+1})) \geq \dim(\mathcal{K}(T^k)) + 1$. Repeated application of this observation yields

$$\begin{aligned} \dim(\mathcal{K}(T^{n+1})) &\geq \dim(\mathcal{K}(T^n)) + 1 \\ &\geq \dim(\mathcal{K}(T^{n-1})) + 2 \\ &\geq \dim(\mathcal{K}(T^{n-2})) + 3 \\ &\vdots \\ &\geq \dim(\mathcal{K}(T^0)) + (n+1) \\ &= n+1 \end{aligned}$$

As $\mathcal{K}(T^{n+1})$ is a subspace of V , of dimension n , this is a contradiction.

The contradiction yields the existence of an integer k such that $\mathcal{K}(T^k) = \mathcal{K}(T^{k+1})$, so we can define m to be smallest such integer with this property. From the argument above about dimensions resulting from a strictly increasing chain of subspaces, we conclude that $m \leq n$.

It remains to show that once two consecutive kernels are equal, then all of the remaining kernels are equal. More formally, if $\mathcal{K}(T^m) = \mathcal{K}(T^{m+1})$, then $\mathcal{K}(T^m) = \mathcal{K}(T^{m+j})$ for all $j \geq 1$. The proof is by induction on j . The base case ($j = 1$) is precisely our defining property for m .

For the induction step, our hypothesis is that $\mathcal{K}(T^m) = \mathcal{K}(T^{m+j})$. We want to establish that $\mathcal{K}(T^m) = \mathcal{K}(T^{m+j+1})$. At the outset of this proof we showed that $\mathcal{K}(T^m) \subseteq \mathcal{K}(T^{m+j+1})$. So we need only show the subset inclusion in the opposite direction. To wit, choose $\mathbf{z} \in \mathcal{K}(T^{m+j+1})$. Then

$$T^{m+j}(T(\mathbf{z})) = T^{m+j+1}(\mathbf{z}) = \mathbf{0}$$

so $T(\mathbf{z}) \in \mathcal{K}(T^{m+j}) = \mathcal{K}(T^m)$. Thus

$$T^{m+1}(\mathbf{z}) = T^m(T(\mathbf{z})) = \mathbf{0}$$

so $\mathbf{z} \in \mathcal{K}(T^{m+1}) = \mathcal{K}(T^m)$, as desired. \square

Example 14. As an illustration of Theorem 3.1 consider the linear transformation $T: \mathbb{C}^{10} \rightarrow \mathbb{C}^{10}$ defined by $T(\mathbf{x}) = A\mathbf{x}$.

$$A = \begin{bmatrix} -27 & 17 & -12 & -1 & 24 & 16 & 26 & -2 & -1 & -3 \\ -66 & 45 & -55 & 11 & 73 & 44 & 45 & -6 & 15 & 1 \\ -85 & 58 & -65 & 13 & 94 & 56 & 61 & -6 & 16 & -4 \\ -81 & 58 & -55 & 4 & 83 & 52 & 70 & -7 & 6 & -6 \\ -33 & 21 & -22 & 6 & 37 & 21 & 23 & -1 & 6 & -4 \\ -38 & 28 & -25 & 1 & 39 & 25 & 34 & -3 & 1 & -4 \\ 20 & -15 & 23 & -8 & -28 & -15 & -8 & 1 & -9 & 0 \\ 41 & -30 & 39 & -12 & -53 & -29 & -23 & 2 & -13 & 3 \\ 58 & -43 & 43 & -6 & -64 & -38 & -47 & 4 & -6 & 6 \\ -69 & 46 & -47 & 7 & 72 & 44 & 54 & -5 & 9 & -4 \end{bmatrix}$$

```
A = matrix(QQ, [
  [-27, 17, -12, -1, 24, 16, 26, -2, -1, -3],
  [-66, 45, -55, 11, 73, 44, 45, -6, 15, 1],
  [-85, 58, -65, 13, 94, 56, 61, -6, 16, -4],
  [-81, 58, -55, 4, 83, 52, 70, -7, 6, -6],
  [-33, 21, -22, 6, 37, 21, 23, -1, 6, -4],
  [-38, 28, -25, 1, 39, 25, 34, -3, 1, -4],
  [20, -15, 23, -8, -28, -15, -8, 1, -9, 0],
  [41, -30, 39, -12, -53, -29, -23, 2, -13, 3],
  [58, -43, 43, -6, -64, -38, -47, 4, -6, 6],
  [-69, 46, -47, 7, 72, 44, 54, -5, 9, -4]
])
```

This linear transformation is engineered to illustrate the full generality of the theorem. The kernels of the powers (null spaces of the matrix powers) increase, with the nullity incrementing first by twos, then by ones, until we top out to find the maximum nullity of 8 at the $m = 6$ power, well less than the maximum of $n = 10$.

$$\begin{array}{lll} \dim(\mathcal{K}(T^0)) = 0 & \dim(\mathcal{K}(T^1)) = 2 & \dim(\mathcal{K}(T^2)) = 4 \\ \dim(\mathcal{K}(T^3)) = 5 & \dim(\mathcal{K}(T^4)) = 6 & \dim(\mathcal{K}(T^5)) = 7 \\ \dim(\mathcal{K}(T^6)) = 8 & \dim(\mathcal{K}(T^7)) = 8 & \end{array}$$

It is somewhat interesting to row-reduce the powers of A , since the null spaces of these powers are the kernels of the powers of T . These are best done with software,

but here are two examples, first a mid-range power, then an extreme power.

$$\begin{array}{l}
 A^4 \xrightarrow{\text{RREF}} \\
 A^{50} \xrightarrow{\text{RREF}}
 \end{array}
 \begin{array}{c}
 \left[\begin{array}{cccccccccc}
 1 & 0 & 0 & 0 & -1 & -3 & 3 & -3 & 0 & 0 \\
 0 & 1 & 0 & 0 & 0 & -4 & 7 & -5 & 0 & -1 \\
 0 & 0 & 1 & 0 & 0 & -1 & 2 & -1 & 0 & -1 \\
 0 & 0 & 0 & 1 & 1 & -1 & 1 & -1 & 1 & -1 \\
 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0
 \end{array} \right] \\
 \\
 \left[\begin{array}{ccccccccccc}
 1 & 0 & -3 & \frac{2}{3} & -\frac{1}{3} & -\frac{2}{3} & -\frac{7}{3} & -\frac{2}{3} & \frac{2}{3} & \frac{7}{3} \\
 0 & 1 & -5 & 1 & 1 & 0 & -2 & -1 & 1 & 3 \\
 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0
 \end{array} \right]
 \end{array}$$

Once we get to the sixth power, the kernels do not change, and so because of the uniqueness of reduced row-echelon form, these do not change either.

3.1.2 Generalized Eigenspaces

These are the two main definitions of this section.

Definition 13. Suppose that $T: V \rightarrow V$ is a linear transformation. Suppose further that for $\mathbf{x} \neq \mathbf{0}$, $(T - \lambda I_V)^k(\mathbf{x}) = \mathbf{0}$ for some $k > 0$. Then \mathbf{x} is a **generalized eigenvector** of T with eigenvalue λ .

Definition 14. Suppose that $T: V \rightarrow V$ is a linear transformation. Define the **generalized eigenspace** of T for λ as

$$\mathcal{G}_T(\lambda) = \left\{ \mathbf{x} \mid (T - \lambda I_V)^k(\mathbf{x}) = \mathbf{0} \text{ for some } k \geq 0 \right\}$$

So the generalized eigenspace is composed of generalized eigenvectors, plus the zero vector. As the name implies, the generalized eigenspace is a subspace of V . But more topically, it is an invariant subspace of V relative to T .

Theorem 3.2. *Suppose that $T: V \rightarrow V$ is a linear transformation. Then the generalized eigenspace $\mathcal{G}_T(\lambda)$ is an invariant subspace of V relative to T .*

Proof. First we establish that $\mathcal{G}_T(\lambda)$ is a subspace of V . Note that $(T - \lambda I_V)^0(\mathbf{0}) = \mathbf{0}$, so by THEOREM LTTZZ we have $\mathbf{0} \in \mathcal{G}_T(\lambda)$.

Suppose that $\mathbf{x}, \mathbf{y} \in \mathcal{G}_T(\lambda)$. Then there are integers k, ℓ such that $(T - \lambda I_V)^k(\mathbf{x}) =$

$\mathbf{0}$ and $(T - \lambda I_V)^\ell(\mathbf{y}) = \mathbf{0}$. Set $m = k + \ell$,

$$\begin{aligned} (T - \lambda I_V)^m(\mathbf{x} + \mathbf{y}) &= (T - \lambda I_V)^m(\mathbf{x}) + (T - \lambda I_V)^m(\mathbf{y}) \\ &= (T - \lambda I_V)^{k+\ell}(\mathbf{x}) + (T - \lambda I_V)^{k+\ell}(\mathbf{y}) \\ &= (T - \lambda I_V)^\ell \left((T - \lambda I_V)^k(\mathbf{x}) \right) + \\ &\quad (T - \lambda I_V)^k \left((T - \lambda I_V)^\ell(\mathbf{y}) \right) \\ &= (T - \lambda I_V)^\ell(\mathbf{0}) + (T - \lambda I_V)^k(\mathbf{0}) \\ &= \mathbf{0} + \mathbf{0} = \mathbf{0} \end{aligned}$$

So $\mathbf{x} + \mathbf{y} \in \mathcal{G}_T(\lambda)$.

Suppose that $\mathbf{x} \in \mathcal{G}_T(\lambda)$ and $\alpha \in \mathbb{C}$. Then there is an integer k such that $(T - \lambda I_V)^k(\mathbf{x}) = \mathbf{0}$.

$$(T - \lambda I_V)^k(\alpha\mathbf{x}) = \alpha(T - \lambda I_V)^k(\mathbf{x}) = \alpha\mathbf{0} = \mathbf{0}$$

. So $\alpha\mathbf{x} \in \mathcal{G}_T(\lambda)$. By THEOREM TSS, $\mathcal{G}_T(\lambda)$ is a subspace of V .

Now we show that $\mathcal{G}_T(\lambda)$ is invariant relative to T . Suppose that $\mathbf{x} \in \mathcal{G}_T(\lambda)$. Then by Definition 14 there is an integer k such that $(T - \lambda I_V)^k(\mathbf{x}) = \mathbf{0}$. The following argument is due to Zoltan Toth.

$$\begin{aligned} (T - \lambda I_V)^k(T(\mathbf{x})) &= (T - \lambda I_V)^k(T(\mathbf{x})) - \lambda\mathbf{0} \\ &= (T - \lambda I_V)^k(T(\mathbf{x})) - \lambda(T - \lambda I_V)^k(\mathbf{x}) \\ &= (T - \lambda I_V)^k(T(\mathbf{x})) - (T - \lambda I_V)^k(\lambda\mathbf{x}) \\ &= (T - \lambda I_V)^k(T(\mathbf{x}) - \lambda\mathbf{x}) \\ &= (T - \lambda I_V)^k((T - \lambda I_V)(\mathbf{x})) \\ &= (T - \lambda I_V)^{k+1}(\mathbf{x}) \\ &= (T - \lambda I_V) \left((T - \lambda I_V)^k(\mathbf{x}) \right) \\ &= (T - \lambda I_V)(\mathbf{0}) = \mathbf{0} \end{aligned}$$

This qualifies $T(\mathbf{x})$ for membership in $\mathcal{G}_T(\lambda)$, so by Definition 14, $\mathcal{G}_T(\lambda)$ is invariant relative to T . \square

Before we compute some generalized eigenspaces, we state and prove one theorem that will make it much easier to create a generalized eigenspace, since it will allow us to use tools we already know well, and will remove some of the ambiguity of the clause “for some k ” in the definition.

Theorem 3.3. *Suppose that $T: V \rightarrow V$ is a linear transformation, $\dim(V) = n$, and λ is an eigenvalue of T . Then $\mathcal{G}_T(\lambda) = \mathcal{K}((T - \lambda I_V)^n)$.*

Proof. To establish the set equality, first suppose that $\mathbf{x} \in \mathcal{G}_T(\lambda)$. Then there is an integer k such that $(T - \lambda I_V)^k(\mathbf{x}) = \mathbf{0}$. This is equivalent to the statement that $\mathbf{x} \in \mathcal{K}((T - \lambda I_V)^k)$. No matter what the value of k is, $\langle\langle$ theorem about nested powers of kernels $\rangle\rangle$ gives

$$\mathbf{x} \in \mathcal{K}((T - \lambda I_V)^k) \subseteq \mathcal{K}((T - \lambda I_V)^n).$$

So, $\mathcal{G}_T(\lambda) \subseteq \mathcal{K}((T - \lambda I_V)^n)$.

For the opposite inclusion, suppose $\mathbf{y} \in \mathcal{K}((T - \lambda I_V)^n)$. Then $(T - \lambda I_V)^n(\mathbf{y}) = \mathbf{0}$, so $\mathbf{y} \in \mathcal{G}_T(\lambda)$ and thus $\mathcal{K}((T - \lambda I_V)^n) \subseteq \mathcal{G}_T(\lambda)$. So we have the desired equality of sets. \square

Theorem 3.3 allows us to compute generalized eigenspaces as a single kernel (or null space of a matrix representation) without considering all possible powers k . We can simply consider the case where $k = n$. It is worth noting that the “regular” eigenspace is a subspace of the generalized eigenspace since

$$\mathcal{E}_T(\lambda) = \mathcal{K}((T - \lambda I_V)^1) \subseteq \mathcal{K}((T - \lambda I_V)^n) = \mathcal{G}_T(\lambda)$$

where the subset inclusion is a consequence of Theorem 3.1.

Also, there is no such thing as a “generalized eigenvalue.” If λ is not an eigenvalue of T , then the kernel of $T - \lambda I_V$ is trivial and therefore subsequent powers of $T - \lambda I_V$ also have trivial kernels (Theorem 3.1 gives $m = 0$). So if we defined generalized eigenspaces for scalars that are not an eigenvalue, they would always be trivial. Alright, we know enough now to compute some generalized eigenspaces. We will record some information about algebraic and geometric multiplicities of eigenvalues (DEFINITION AME, DEFINITION GME) as we go, since these observations will be of interest in light of some future theorems.

Example 15. In order to gain some experience with generalized eigenspaces, we construct one and then also construct a matrix representation for the restriction to this invariant subspace.

Consider the linear transformation $T: \mathbb{C}^5 \rightarrow \mathbb{C}^5$ defined by $T(\mathbf{x}) = A\mathbf{x}$, where

$$A = \begin{bmatrix} -22 & -24 & -24 & -24 & -46 \\ 3 & 2 & 6 & 0 & 11 \\ -12 & -16 & -6 & -14 & -17 \\ 6 & 8 & 4 & 10 & 8 \\ 11 & 14 & 8 & 13 & 18 \end{bmatrix}$$

One of the eigenvalues of A is $\lambda = 2$, with geometric multiplicity $\gamma_T(2) = 1$, and algebraic multiplicity $\alpha_T(2) = 3$. We get the generalized eigenspace according to Theorem 3.3,

$$W = \mathcal{G}_T(2) = \mathcal{K}((T - 2I_{\mathbb{C}^5})^5) = \left\langle \left\{ \begin{bmatrix} -2 \\ 1 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ -1 \\ 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} -4 \\ 2 \\ 0 \\ 0 \\ 1 \end{bmatrix} \right\} \right\rangle = \langle \{\mathbf{w}_1, \mathbf{w}_2, \mathbf{w}_3\} \rangle$$

By Theorem 3.2, we know W is invariant relative to T , so we can employ DEFINITION LTR to form the restriction, $T|_W: W \rightarrow W$.

We will form the restriction of T to W , $T|_W$, since we will do this frequently in subsequent examples. For a basis of W we will use $C = \{\mathbf{w}_1, \mathbf{w}_2, \mathbf{w}_3\}$. Notice that $\dim(W) = 3$, so our matrix representation will be a square matrix of size 3.

Applying DEFINITION MR, we compute

$$\begin{aligned}\rho_C(T(\mathbf{w}_1)) &= \rho_C(A\mathbf{w}_1) \\ &= \rho_C\left(\begin{bmatrix} -4 \\ 2 \\ 2 \\ 0 \\ 0 \end{bmatrix}\right) = \rho_C\left(2\begin{bmatrix} -2 \\ 1 \\ 1 \\ 0 \\ 0 \end{bmatrix} + 0\begin{bmatrix} 0 \\ -1 \\ 0 \\ 1 \\ 0 \end{bmatrix} + 0\begin{bmatrix} -4 \\ 2 \\ 0 \\ 0 \\ 1 \end{bmatrix}\right) = \begin{bmatrix} 2 \\ 0 \\ 0 \end{bmatrix} \\ \rho_C(T(\mathbf{w}_2)) &= \rho_C(A\mathbf{w}_2) \\ &= \rho_C\left(\begin{bmatrix} 0 \\ -2 \\ 2 \\ 2 \\ -1 \end{bmatrix}\right) = \rho_C\left(2\begin{bmatrix} -2 \\ 1 \\ 1 \\ 0 \\ 0 \end{bmatrix} + 2\begin{bmatrix} 0 \\ -1 \\ 0 \\ 1 \\ 0 \end{bmatrix} + (-1)\begin{bmatrix} -4 \\ 2 \\ 0 \\ 0 \\ 1 \end{bmatrix}\right) = \begin{bmatrix} 2 \\ 2 \\ -1 \end{bmatrix} \\ \rho_C(T(\mathbf{w}_3)) &= \rho_C(A\mathbf{w}_3) \\ &= \rho_C\left(\begin{bmatrix} -6 \\ 3 \\ -1 \\ 0 \\ 2 \end{bmatrix}\right) = \rho_C\left((-1)\begin{bmatrix} -2 \\ 1 \\ 1 \\ 0 \\ 0 \end{bmatrix} + 0\begin{bmatrix} 0 \\ -1 \\ 0 \\ 1 \\ 0 \end{bmatrix} + 2\begin{bmatrix} -4 \\ 2 \\ 0 \\ 0 \\ 1 \end{bmatrix}\right) = \begin{bmatrix} -1 \\ 0 \\ 2 \end{bmatrix}\end{aligned}$$

So the matrix representation of $T|_W$ relative to C is

$$M_{C,C}^{T|_W} = \begin{bmatrix} 2 & 2 & -1 \\ 0 & 2 & 0 \\ 0 & -1 & 2 \end{bmatrix}$$

The question arises: how do we use a 3×3 matrix to compute with vectors from \mathbb{C}^5 ? To answer this question, consider the randomly chosen vector

$$\mathbf{w} = \begin{bmatrix} -4 \\ 4 \\ 4 \\ -2 \\ -1 \end{bmatrix}$$

First check that $\mathbf{w} \in \mathcal{G}_T(2)$. There are two ways to do this, first verify that

$$(T - 2I_{\mathbb{C}^5})^5(\mathbf{w}) = (A - 2I_5)^5\mathbf{w} = \mathbf{0}$$

meeting Definition 13 (with $k = 5$). Or, express \mathbf{w} as a linear combination of the basis C for W , to wit, $\mathbf{w} = 4\mathbf{w}_1 - 2\mathbf{w}_2 - \mathbf{w}_3$.

Now compute $T|_W(\mathbf{w})$ directly

$$T|_W(\mathbf{w}) = T(\mathbf{w}) = A\mathbf{w} = \begin{bmatrix} -10 \\ 9 \\ 5 \\ -4 \\ 0 \end{bmatrix}$$

It was necessary to verify that $\mathbf{w} \in \mathcal{G}_T(2)$. If we trust our work so far, then this output we just computed will also be an element of W , but it would be wise to check this anyway (using either of the methods we used for \mathbf{w}). We'll wait.

Now we will repeat this sample computation, but instead using the matrix representation of $T|_W$ relative to C .

$$\begin{aligned}
 T|_W(\mathbf{w}) &= \rho_C^{-1} \left(M_{C,C}^{T|_W} \rho_C(\mathbf{w}) \right) \\
 &= \rho_C^{-1} \left(M_{C,C}^{T|_W} \rho_C(4\mathbf{w}_1 - 2\mathbf{w}_2 - \mathbf{w}_3) \right) \\
 &= \rho_C^{-1} \left(\begin{bmatrix} 2 & 2 & -1 \\ 0 & 2 & 0 \\ 0 & -1 & 2 \end{bmatrix} \begin{bmatrix} 4 \\ -2 \\ -1 \end{bmatrix} \right) = \rho_C^{-1} \left(\begin{bmatrix} 5 \\ -4 \\ 0 \end{bmatrix} \right) \\
 &= 5 \begin{bmatrix} -2 \\ 1 \\ 1 \\ 0 \\ 0 \end{bmatrix} + (-4) \begin{bmatrix} 0 \\ -1 \\ 0 \\ 1 \\ 0 \end{bmatrix} + 0 \begin{bmatrix} -4 \\ 2 \\ 0 \\ 0 \\ 1 \end{bmatrix} = \begin{bmatrix} -10 \\ 9 \\ 5 \\ -4 \\ 0 \end{bmatrix}
 \end{aligned}$$

This matches the previous computation. Notice how the “action” of $T|_W$ is accomplished by a 3×3 matrix multiplying a column vector of size 3.

If you would like more practice with these sorts of computations, mimic the above using the other eigenvalue of T , which is $\lambda = -2$. The generalized eigenspace has dimension 2, so the matrix representation of the restriction to the generalized eigenspace will be a 2×2 matrix.

Our next two examples compute a complete set of generalized eigenspaces for a linear transformation.

Example 16. In Example 3 we presented two invariant subspaces of \mathbb{C}^4 . There was some mystery about just how these were constructed, but we can now reveal that they are generalized eigenspaces. Example 3 featured $T: \mathbb{C}^4 \rightarrow \mathbb{C}^4$ defined by $T(\mathbf{x}) = A\mathbf{x}$ with A given by

$$A = \begin{bmatrix} -8 & 6 & -15 & 9 \\ -8 & 14 & -10 & 18 \\ 1 & 1 & 3 & 0 \\ 3 & -8 & 2 & -11 \end{bmatrix}$$

A matrix representation of T relative to the standard basis (DEFINITION SUV) will equal A . So we can analyze A with the techniques of CHAPTER E. Doing so, we find two eigenvalues, $\lambda = 1, -2$, with multiplicities,

$$\alpha_T(1) = 2 \qquad \gamma_T(1) = 1 \qquad \alpha_T(-2) = 2 \qquad \gamma_T(-2) = 1$$

To apply Theorem 3.3 we subtract each eigenvalue from the diagonal entries of A ,

raise the result to the power $\dim(\mathbb{C}^4) = 4$, and compute a basis for the null space.

$$(A - (-2)I_4)^4 = \begin{bmatrix} 648 & -1215 & 729 & -1215 \\ -324 & 486 & -486 & 486 \\ -405 & 729 & -486 & 729 \\ 297 & -486 & 405 & -486 \end{bmatrix} \xrightarrow{\text{RREF}} \begin{bmatrix} 1 & 0 & 3 & 0 \\ 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

$$\mathcal{G}_T(-2) = \left\langle \left\{ \begin{bmatrix} -3 \\ -1 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ -1 \\ 0 \\ 1 \end{bmatrix} \right\} \right\rangle$$

$$(A - (1)I_4)^4 = \begin{bmatrix} 81 & -405 & -81 & -729 \\ -108 & -189 & -378 & -486 \\ -27 & 135 & 27 & 243 \\ 135 & 54 & 351 & 243 \end{bmatrix} \xrightarrow{\text{RREF}} \begin{bmatrix} 1 & 0 & 7/3 & 1 \\ 0 & 1 & 2/3 & 2 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

$$\mathcal{G}_T(1) = \left\langle \left\{ \begin{bmatrix} -7 \\ -2 \\ 3 \\ 0 \end{bmatrix}, \begin{bmatrix} -1 \\ -2 \\ 0 \\ 1 \end{bmatrix} \right\} \right\rangle$$

In Example 3 we concluded that these two invariant subspaces formed a direct sum of \mathbb{C}^4 , only at that time, they were called X and W . Now we can write

$$\mathbb{C}^4 = \mathcal{G}_T(1) \oplus \mathcal{G}_T(-2)$$

This is no accident. Notice that the dimension of each of these invariant subspaces is equal to the algebraic multiplicity of the associated eigenvalue. Not an accident either. (See the upcoming Theorem 3.4.)

Example 17. Define the linear transformation $S: \mathbb{C}^6 \rightarrow \mathbb{C}^6$ by $S(\mathbf{x}) = B\mathbf{x}$ where

$$\begin{bmatrix} 2 & -4 & 25 & -54 & 90 & -37 \\ 2 & -3 & 4 & -16 & 26 & -8 \\ 2 & -3 & 4 & -15 & 24 & -7 \\ 10 & -18 & 6 & -36 & 51 & -2 \\ 8 & -14 & 0 & -21 & 28 & 4 \\ 5 & -7 & -6 & -7 & 8 & 7 \end{bmatrix}$$

Then B will be the matrix representation of S relative to the standard basis and we can use the techniques of CHAPTER E applied to B in order to find the eigenvalues of S .

$$\alpha_S(3) = 2 \qquad \gamma_S(3) = 1 \qquad \alpha_S(-1) = 4 \qquad \gamma_S(-1) = 2$$

To find the generalized eigenspaces of S we need to subtract an eigenvalue from the diagonal elements of B , raise the result to the power $\dim(\mathbb{C}^6) = 6$ and compute

the null space. Here are the results for the two eigenvalues of S ,

$$\begin{aligned}
 (B - 3I_6)^6 &= \begin{bmatrix} 64000 & -152576 & -59904 & 26112 & -95744 & 133632 \\ 15872 & -39936 & -11776 & 8704 & -29184 & 36352 \\ 12032 & -30208 & -9984 & 6400 & -20736 & 26368 \\ -1536 & 11264 & -23040 & 17920 & -17920 & -1536 \\ -9728 & 27648 & -6656 & 9728 & -1536 & -17920 \\ -7936 & 17920 & 5888 & 1792 & 4352 & -14080 \end{bmatrix} \\
 &\xrightarrow{\text{RREF}} \begin{bmatrix} 1 & 0 & 0 & 0 & -4 & 5 \\ 0 & 1 & 0 & 0 & -1 & 1 \\ 0 & 0 & 1 & 0 & -1 & 1 \\ 0 & 0 & 0 & 1 & -2 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \\
 \mathcal{G}_S(3) &= \left\langle \left\{ \begin{bmatrix} 4 \\ 1 \\ 1 \\ 2 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} -5 \\ -1 \\ -1 \\ -1 \\ 0 \\ 1 \end{bmatrix} \right\} \right\rangle \\
 (B - (-1)I_6)^6 &= \begin{bmatrix} 6144 & -16384 & 18432 & -36864 & 57344 & -18432 \\ 4096 & -8192 & 4096 & -16384 & 24576 & -4096 \\ 4096 & -8192 & 4096 & -16384 & 24576 & -4096 \\ 18432 & -32768 & 6144 & -61440 & 90112 & -6144 \\ 14336 & -24576 & 2048 & -45056 & 65536 & -2048 \\ 10240 & -16384 & -2048 & -28672 & 40960 & 2048 \end{bmatrix} \\
 &\xrightarrow{\text{RREF}} \begin{bmatrix} 1 & 0 & -5 & 2 & -4 & 5 \\ 0 & 1 & -3 & 3 & -5 & 3 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \\
 \mathcal{G}_S(-1) &= \left\langle \left\{ \begin{bmatrix} 5 \\ 3 \\ 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} -2 \\ -3 \\ 0 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 4 \\ 5 \\ 0 \\ 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} -5 \\ -3 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} \right\} \right\rangle
 \end{aligned}$$

If we take the union of the two bases for these two invariant subspaces we obtain the set

$$C = \left\{ \begin{bmatrix} 4 \\ 1 \\ 1 \\ 2 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} -5 \\ -1 \\ -1 \\ -1 \\ 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 5 \\ 3 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} -2 \\ -3 \\ 0 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 4 \\ 5 \\ 0 \\ 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} -5 \\ -3 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} \right\}$$

You can check that this set is linearly independent (right now we have no guarantee this will happen). Once this is verified, we have a basis for \mathbb{C}^6 . This is enough for us to apply Theorem 1.1 and conclude that

$$\mathbb{C}^6 = \mathcal{G}_S(3) \oplus \mathcal{G}_S(-1)$$

This is no accident. Notice that the dimension of each of these invariant subspaces is equal to the algebraic multiplicity of the associated eigenvalue. Not an accident either. (See Theorem 3.4.)

Our principal interest in generalized eigenspaces is the following important theorem, which has been presaged by the two previous examples.

Theorem 3.4. *Suppose that $T: V \rightarrow V$ is a linear transformation with distinct eigenvalues $\lambda_1, \lambda_2, \lambda_3, \dots, \lambda_m$. Then*

$$V = \mathcal{G}_T(\lambda_1) \oplus \mathcal{G}_T(\lambda_2) \oplus \mathcal{G}_T(\lambda_3) \oplus \cdots \oplus \mathcal{G}_T(\lambda_m)$$

Proof. We will provide a complete proof soon. For now, we give an outline.

We now that a decomposition of the domain of a linear transformation into invariant subspaces will give a block diagonal matrix representation. But it cuts both ways. If there is a similarity transformation to a block diagonal matrix, then the columns of the nonsingular matrix used for similarity will be a basis that can be partitioned into bases of invariant subspaces that are a direct sum decomposition of the domain (THEOREM SCB). So we outline a sequence of similarity transformations that converts *any* square matrix to the appropriate block diagonal form.

1. Begin with the eigenvalues of the matrix, ordered so that equal eigenvalues are adjacent.
2. Determine the upper triangular matrix with these eigenvalues on the diagonal and similar to the original matrix as guaranteed by THEOREM UTMTR.
3. Suppose that the entry in row i and column j in the “upper half” (so $j > i$) has the value a . Suppose further that the diagonal entries (eigenvalues) λ_i and λ_j are different.

Define S to be the identity matrix, with the addition of the entry $\frac{a}{\lambda_j - \lambda_i}$ in row i and column j . Then a similarity transformation by S will place a zero in row i and column j . Here is where we begin to understand being careful about equal and different eigenvalues.

4. The similarity transformation of the previous step will change other entries of the matrix, *but only* in row i to the *right* of the entry of interest and column j *above* the entry of interest.
5. Begin in the bottom row, going only as far right as needed to get different eigenvalues, and “zero out” the rest of the row. Move up a row, work left to right, “zeroing out” as much of the row as possible. Continue moving up a row at a time, then move left to right in the row. The restriction to using different eigenvalues will cut a staircase pattern.
6. You should understand that the blocks left on the diagonal correspond to runs of equal eigenvalues on the diagonal. So each block has a size equal to the algebraic multiplicity of the eigenvalue.

□

Now, given any linear transformation, we can find a decomposition of the domain into a collection of invariant subspaces. And, as we have seen, such a decomposition will provide a basis for the domain so that a matrix representation relative to this basis will have a block diagonal form. Besides a decomposition into invariant subspaces, this proof has a bonus for us.

Corollary 3.5. *Suppose $T: V \rightarrow V$ is a linear transformation with eigenvalue λ . Then the dimension of the generalized eigenspace for λ is the algebraic multiplicity of λ , $\dim(\mathcal{G}_T(\lambda)) = \alpha_T(\lambda)$.*

Proof. Coming soon: as a consequence of proof, or by counting dimensions with inequality on geometric dimension. \square

We illustrate the use of this decomposition in building a block diagonal matrix representation.

Example 18. In Example 17 we computed the generalized eigenspaces of the linear transformation $S: \mathbb{C}^6 \rightarrow \mathbb{C}^6$ by $S(\mathbf{x}) = B\mathbf{x}$ where

$$B = \begin{bmatrix} 2 & -4 & 25 & -54 & 90 & -37 \\ 2 & -3 & 4 & -16 & 26 & -8 \\ 2 & -3 & 4 & -15 & 24 & -7 \\ 10 & -18 & 6 & -36 & 51 & -2 \\ 8 & -14 & 0 & -21 & 28 & 4 \\ 5 & -7 & -6 & -7 & 8 & 7 \end{bmatrix}$$

We also recognized that these generalized eigenspaces provided a vector space decomposition.

From these generalized eigenspaces, we found the basis

$$C = \{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3, \mathbf{v}_4, \mathbf{v}_5, \mathbf{v}_6\} \\ = \left\{ \begin{bmatrix} 4 \\ 1 \\ 1 \\ 2 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} -5 \\ -1 \\ -1 \\ -1 \\ 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 5 \\ 3 \\ 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} -2 \\ -3 \\ 0 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 4 \\ 5 \\ 0 \\ 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} -5 \\ -3 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} \right\}$$

of \mathbb{C}^6 where $\{\mathbf{v}_1, \mathbf{v}_2\}$ is a basis of $\mathcal{G}_S(3)$ and $\{\mathbf{v}_3, \mathbf{v}_4, \mathbf{v}_5, \mathbf{v}_6\}$ is a basis of $\mathcal{G}_S(-1)$

We can employ C in the construction of a matrix representation of S (DEFINITION

MR). Here are the computations,

$$\begin{aligned}
 \rho_C(S(\mathbf{v}_1)) &= \rho_C \begin{pmatrix} 11 \\ 3 \\ 3 \\ 7 \\ 4 \\ 1 \end{pmatrix} = \rho_C(4\mathbf{v}_1 + 1\mathbf{v}_2) = \begin{bmatrix} 4 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \\
 \rho_C(S(\mathbf{v}_2)) &= \rho_C \begin{pmatrix} -14 \\ -3 \\ -3 \\ -4 \\ -1 \\ 2 \end{pmatrix} = \rho_C((-1)\mathbf{v}_1 + 2\mathbf{v}_2) = \begin{bmatrix} -1 \\ 2 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \\
 \rho_C(S(\mathbf{v}_3)) &= \rho_C \begin{pmatrix} 23 \\ 5 \\ 5 \\ 2 \\ -2 \\ -2 \end{pmatrix} = \rho_C(5\mathbf{v}_3 + 2\mathbf{v}_4 + (-2)\mathbf{v}_5 + (-2)\mathbf{v}_6) = \begin{bmatrix} 0 \\ 0 \\ 5 \\ 2 \\ -2 \\ -2 \end{bmatrix} \\
 \rho_C(S(\mathbf{v}_4)) &= \rho_C \begin{pmatrix} -46 \\ -11 \\ -10 \\ -2 \\ 5 \\ 4 \end{pmatrix} = \rho_C((-10)\mathbf{v}_3 + (-2)\mathbf{v}_4 + 5\mathbf{v}_5 + 4\mathbf{v}_6) = \begin{bmatrix} 0 \\ 0 \\ -10 \\ -2 \\ 5 \\ 4 \end{bmatrix} \\
 \rho_C(S(\mathbf{v}_5)) &= \rho_C \begin{pmatrix} 78 \\ 19 \\ 17 \\ 1 \\ -10 \\ -7 \end{pmatrix} = \rho_C(17\mathbf{v}_3 + 1\mathbf{v}_4 + (-10)\mathbf{v}_5 + (-7)\mathbf{v}_6) = \begin{bmatrix} 0 \\ 0 \\ 17 \\ 1 \\ -10 \\ -7 \end{bmatrix} \\
 \rho_C(S(\mathbf{v}_6)) &= \rho_C \begin{pmatrix} -35 \\ -9 \\ -8 \\ 2 \\ 6 \\ 3 \end{pmatrix} = \rho_C((-8)\mathbf{v}_3 + 2\mathbf{v}_4 + 6\mathbf{v}_5 + 3\mathbf{v}_6) = \begin{bmatrix} 0 \\ 0 \\ -8 \\ 2 \\ 6 \\ 3 \end{bmatrix}
 \end{aligned}$$

These column vectors are the columns of the matrix representation, so we obtain

$$M_{C,C}^S = \begin{bmatrix} 4 & -1 & 0 & 0 & 0 & 0 \\ 1 & 2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 5 & -10 & 17 & -8 \\ 0 & 0 & 2 & -2 & 1 & 2 \\ 0 & 0 & -2 & 5 & -10 & 6 \\ 0 & 0 & -2 & 4 & -7 & 3 \end{bmatrix}$$

As before, the key feature of this representation is the 2×2 and 4×4 blocks on the diagonal. They arise from generalized eigenspaces and their sizes are equal to the algebraic multiplicities of the eigenvalues.

3.2 Nilpotent Linear Transformations

We will discover that nilpotent linear transformations are the essential obstacle in a non-diagonalizable linear transformation. So we will study them carefully, both as an object of inherent mathematical interest, but also as the object at the heart of the argument that leads to a pleasing canonical form for any linear transformation. Once we understand these linear transformations thoroughly, we will be able to easily analyze the structure of any linear transformation.

3.2.1 Nilpotent Linear Transformations

Definition 15. Suppose that $T: V \rightarrow V$ is a linear transformation such that there is an integer $p > 0$ such that $T^p(\mathbf{v}) = \mathbf{0}$ for every $\mathbf{v} \in V$. The smallest p for which this condition is met is called the *index* of T .

Of course, the linear transformation T defined by $T(\mathbf{v}) = \mathbf{0}$ will qualify as nilpotent of index 1. But are there others? Yes, of course.

Example 19. Recall that our definitions and theorems are being stated for linear transformations on abstract vector spaces, while our examples will work with square matrices (and use the same terms interchangeably). In this case, to demonstrate the existence of nontrivial nilpotent linear transformations, we desire a matrix such that some power of the matrix is the zero matrix. Consider powers of a 6×6 matrix A ,

$$A = \begin{bmatrix} -3 & 3 & -2 & 5 & 0 & -5 \\ -3 & 5 & -3 & 4 & 3 & -9 \\ -3 & 4 & -2 & 6 & -4 & -3 \\ -3 & 3 & -2 & 5 & 0 & -5 \\ -3 & 3 & -2 & 4 & 2 & -6 \\ -2 & 3 & -2 & 2 & 4 & -7 \end{bmatrix}$$

and compute powers of A ,

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

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

$$A^4 = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

Thus we can say that A is nilpotent of index 4.

Because it will presage some upcoming theorems, we will record some extra information about the eigenvalues and eigenvectors of A here. A has just one eigenvalue, $\lambda = 0$, with algebraic multiplicity 6 and geometric multiplicity 2. The eigenspace for this eigenvalue is

$$\mathcal{E}_A(0) = \left\langle \begin{bmatrix} 2 \\ 2 \\ 5 \\ 2 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} -1 \\ -1 \\ -5 \\ -1 \\ 0 \\ 1 \end{bmatrix} \right\rangle$$

If there were degrees of singularity, we might say this matrix was *very* singular, since zero is an eigenvalue with maximum algebraic multiplicity (THEOREM SMZE, THEOREM ME). Notice too that A is “far” from being diagonalizable (THEOREM DMFE).

With the existence of nontrivial nilpotent matrices settled, let’s look at another example.

Example 20. Consider the matrix

$$B = \begin{bmatrix} -1 & 1 & -1 & 4 & -3 & -1 \\ 1 & 1 & -1 & 2 & -3 & -1 \\ -9 & 10 & -5 & 9 & 5 & -15 \\ -1 & 1 & -1 & 4 & -3 & -1 \\ 1 & -1 & 0 & 2 & -4 & 2 \\ 4 & -3 & 1 & -1 & -5 & 5 \end{bmatrix}$$

and compute the second power of B ,

$$B^2 = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

So B is nilpotent of index 2.

Again, the only eigenvalue of B is zero, with algebraic multiplicity 6. The geometric multiplicity of the eigenvalue is 3, as seen in the eigenspace,

$$\mathcal{E}_B(0) = \left\langle \begin{bmatrix} 1 \\ 3 \\ 6 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ -4 \\ -7 \\ 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 2 \\ 1 \\ 0 \\ 0 \\ 1 \end{bmatrix} \right\rangle$$

Again, THEOREM DMFE tells us that B is far from being diagonalizable.

On a first encounter with the definition of a nilpotent matrix, you might wonder if such a thing was possible at all. That a high power of a nonzero object could be zero is so very different from our experience with scalars that it seems very unnatural. Hopefully the two previous examples were somewhat surprising. But we

have seen that matrix algebra does not always behave the way we expect (EXAMPLE MMNC), and we also now recognize matrix products not just as arithmetic, but as function composition (THEOREM MRCLT). With a couple examples completed, we turn to some general properties.

Theorem 3.6. *Suppose that $T: V \rightarrow V$ is a nilpotent linear transformation and λ is an eigenvalue of T . Then $\lambda = 0$.*

Proof. Let \mathbf{x} be an eigenvector of T for the eigenvalue λ , and suppose that T is nilpotent with index p . Then

$$\mathbf{0} = T^p(\mathbf{x}) = \lambda^p \mathbf{x}$$

Because \mathbf{x} is an eigenvector, it is nonzero, and therefore THEOREM SMEZV tells us that $\lambda^p = 0$ and so $\lambda = 0$. \square

Paraphrasing, all of the eigenvalues of a nilpotent linear transformation are zero. So in particular, the characteristic polynomial of a nilpotent linear transformation, T , on a vector space of dimension n , is simply $p_T(x) = x^n$.

The next theorem is not critical for what follows, but it will explain our interest in nilpotent linear transformations. More specifically, it is the first step in backing up the assertion that nilpotent linear transformations are the essential obstacle in a non-diagonalizable linear transformation. While it is not obvious from the statement of the theorem, it says that a nilpotent linear transformation is not diagonalizable, unless it is trivially so.

Theorem 3.7. *Suppose the linear transformation $T: V \rightarrow V$ is nilpotent. Then T is diagonalizable if and only if T is the zero linear transformation.*

Proof. (\Leftarrow) We start with the easy direction. Let $n = \dim(V)$. The linear transformation $Z: V \rightarrow V$ defined by $Z(\mathbf{v}) = \mathbf{0}$ for all $\mathbf{v} \in V$ is nilpotent of index $p = 1$ and a matrix representation relative to any basis of V is the $n \times n$ zero matrix, \mathcal{O} . Quite obviously, the zero matrix is a diagonal matrix (DEFINITION DIM) and hence Z is diagonalizable (DEFINITION DZM).

(\Rightarrow) Assume now that T is diagonalizable, so $\gamma_T(\lambda) = \alpha_T(\lambda)$ for every eigenvalue λ (THEOREM DMFE). By Theorem 3.6, T has only one eigenvalue (zero), which therefore must have algebraic multiplicity n (THEOREM NEM). So the geometric multiplicity of zero will be n as well, $\gamma_T(0) = n$.

Let B be a basis for the eigenspace $\mathcal{E}_T(0)$. Then B is a linearly independent subset of V of size n , and thus a basis of V . For any $\mathbf{x} \in B$ we have

$$T(\mathbf{x}) = 0\mathbf{x} = \mathbf{0}$$

So T is identically zero on a basis for B , and since the action of a linear transformation on a basis determines all of the values of the linear transformation (THEOREM LTDB), it is easy to see that $T(\mathbf{v}) = \mathbf{0}$ for every $\mathbf{v} \in V$. \square

So, other than one trivial case (the zero linear transformation), every nilpotent linear transformation is not diagonalizable. It remains to see what is so “essential” about this broad class of non-diagonalizable linear transformations.

3.2.2 Powers of Kernels of Nilpotent Linear Transformations

We return to our discussion of kernels of powers of linear transformations, now specializing to nilpotent linear transformations. We reprise Theorem 3.1, gaining just a little more precision in the conclusion.

Theorem 3.8. *Suppose $T: V \rightarrow V$ is a nilpotent linear transformation with index p and $\dim(V) = n$. Then $0 \leq p \leq n$ and*

$$\{\mathbf{0}\} = \mathcal{K}(T^0) \subsetneq \mathcal{K}(T^1) \subsetneq \mathcal{K}(T^2) \subsetneq \cdots \subsetneq \mathcal{K}(T^p) = \mathcal{K}(T^{p+1}) = \cdots = V$$

Proof. Since $T^p = 0$ it follows that $T^{p+j} = 0$ for all $j \geq 0$ and thus $\mathcal{K}(T^{p+j}) = V$ for $j \geq 0$. So the value of m guaranteed by THEOREM KPLT is at most p . The only remaining aspect of our conclusion that does not follow from Theorem 3.1 is that $m = p$. To see this, we must show that $\mathcal{K}(T^k) \subsetneq \mathcal{K}(T^{k+1})$ for $0 \leq k \leq p-1$. If $\mathcal{K}(T^k) = \mathcal{K}(T^{k+1})$ for some $k < p$, then $\mathcal{K}(T^k) = \mathcal{K}(T^p) = V$. This implies that $T^k = 0$, violating the fact that T has index p . So the smallest value of m is indeed p , and we learn that $p < n$. \square

The structure of the kernels of powers of nilpotent linear transformations will be crucial to what follows. But immediately we can see a practical benefit. Suppose we are confronted with the question of whether or not an $n \times n$ matrix, A , is nilpotent or not. If we don't quickly find a low power that equals the zero matrix, when do we stop trying higher and higher powers? Theorem 3.8 gives us the answer: if we don't see a zero matrix by the time we finish computing A^n , then it is not going to ever happen. We will now take a look at one example of Theorem 3.8 in action.

Example 21. We will recycle the nilpotent matrix A of index 4 from Example 19. We now know that would have only needed to look at the first 6 powers of A if the matrix had not been nilpotent and we wanted to discover that. We list bases for the null spaces of the powers of A . (Notice how we are using null spaces for matrices interchangeably with kernels of linear transformations, see THEOREM KNSI for justification.)

$$\begin{aligned} \mathcal{N}(A) &= \mathcal{N} \left(\begin{bmatrix} -3 & 3 & -2 & 5 & 0 & -5 \\ -3 & 5 & -3 & 4 & 3 & -9 \\ -3 & 4 & -2 & 6 & -4 & -3 \\ -3 & 3 & -2 & 5 & 0 & -5 \\ -3 & 3 & -2 & 4 & 2 & -6 \\ -2 & 3 & -2 & 2 & 4 & -7 \end{bmatrix} \right) = \left\langle \left\{ \begin{bmatrix} 2 \\ 2 \\ 5 \\ 2 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} -1 \\ -1 \\ -5 \\ -1 \\ 0 \\ 1 \end{bmatrix} \right\} \right\rangle \\ \mathcal{N}(A^2) &= \mathcal{N} \left(\begin{bmatrix} 1 & -2 & 1 & 0 & -3 & 4 \\ 0 & -2 & 1 & 1 & -3 & 4 \\ 3 & 0 & 0 & -3 & 0 & 0 \\ 1 & -2 & 1 & 0 & -3 & 4 \\ 0 & -2 & 1 & 1 & -3 & 4 \\ -1 & -2 & 1 & 2 & -3 & 4 \end{bmatrix} \right) = \left\langle \left\{ \begin{bmatrix} 0 \\ 1 \\ 2 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 2 \\ 1 \\ 0 \\ 2 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ -3 \\ 0 \\ 0 \\ 2 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 2 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} \right\} \right\rangle \\ \mathcal{N}(A^3) &= \mathcal{N} \left(\begin{bmatrix} 1 & 0 & 0 & -1 & 0 & 0 \\ 1 & 0 & 0 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & -1 & 0 & 0 \\ 1 & 0 & 0 & -1 & 0 & 0 \\ 1 & 0 & 0 & -1 & 0 & 0 \end{bmatrix} \right) = \left\langle \left\{ \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \\ 0 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} \right\} \right\rangle \\ \mathcal{N}(A^4) &= \mathcal{N} \left(\begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \right) = \left\langle \left\{ \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} \right\} \right\rangle \end{aligned}$$

With the exception of some convenience scaling of the basis vectors in $\mathcal{N}(A^2)$ these are exactly the basis vectors described in THEOREM BNS. We can see that the dimension of $\mathcal{N}(A)$ equals the geometric multiplicity of the zero eigenvalue. Why is this not an accident? We can see the dimensions of the kernels consistently increasing, and we can see that $\mathcal{N}(A^4) = \mathbb{C}^6$. But Theorem 3.8 says a little more. Each successive kernel should be a superset of the previous one. We ought to be able to begin with a basis of $\mathcal{N}(A)$ and *extend* it to a basis of $\mathcal{N}(A^2)$. Then we should be able to extend a basis of $\mathcal{N}(A^2)$ into a basis of $\mathcal{N}(A^3)$, all with repeated applications of THEOREM ELIS. Verify the following,

$$\begin{aligned} \mathcal{N}(A) &= \left\langle \left\{ \begin{bmatrix} 2 \\ 2 \\ 5 \\ 2 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} -1 \\ -1 \\ -5 \\ -1 \\ 0 \\ 1 \end{bmatrix} \right\} \right\rangle \\ \mathcal{N}(A^2) &= \left\langle \left\{ \begin{bmatrix} 2 \\ 2 \\ 5 \\ 2 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} -1 \\ -1 \\ -5 \\ -1 \\ 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 0 \\ -3 \\ 0 \\ 0 \\ 2 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 2 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} \right\} \right\rangle \\ \mathcal{N}(A^3) &= \left\langle \left\{ \begin{bmatrix} 2 \\ 2 \\ 5 \\ 2 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} -1 \\ -1 \\ -5 \\ -1 \\ 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 0 \\ -3 \\ 0 \\ 0 \\ 2 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 2 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} \right\} \right\rangle \\ \mathcal{N}(A^4) &= \left\langle \left\{ \begin{bmatrix} 2 \\ 2 \\ 5 \\ 2 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} -1 \\ -1 \\ -5 \\ -1 \\ 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 0 \\ -3 \\ 0 \\ 0 \\ 2 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 2 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \\ 0 \\ 0 \end{bmatrix} \right\} \right\rangle \end{aligned}$$

Do not be concerned at the moment about how these bases were constructed since we are not describing the applications of THEOREM ELIS here. Do verify carefully for each alleged basis that, (1) it is a superset of the basis for the previous kernel, (2) the basis vectors really are members of the kernel of the associated power of A , (3) the basis is a linearly independent set, (4) the size of the basis is equal to the size of the basis found previously for each kernel. With these verifications, you will know that we have successfully demonstrated what Theorem 3.8 guarantees.

3.2.3 Restrictions to Generalized Eigenspaces

We have seen that we can decompose the domain of a linear transformation into a direct sum of generalized eigenspaces (Theorem 3.4). And we know that we can then easily obtain a basis that leads to a block diagonal matrix representation. The blocks of this matrix representation are matrix representations of *restrictions* to the generalized eigenspaces (for example, Example 18). And the next theorem tells us that these restrictions, adjusted slightly, provide us with a broad class of nilpotent linear transformations.

Theorem 3.9. *Suppose $T: V \rightarrow V$ is a linear transformation with eigenvalue λ . Then the linear transformation $T|_{\mathcal{G}_T(\lambda)} - \lambda I_{\mathcal{G}_T(\lambda)}$ is nilpotent.*

Proof. Notice first that every subspace of V is invariant with respect to I_V , so $I_{\mathcal{G}_T(\lambda)} = I_V|_{\mathcal{G}_T(\lambda)}$. Let $n = \dim(V)$ and choose $\mathbf{v} \in \mathcal{G}_T(\lambda)$. Then with an application of Theorem 3.3,

$$(T|_{\mathcal{G}_T(\lambda)} - \lambda I_{\mathcal{G}_T(\lambda)})^n(\mathbf{v}) = (T - \lambda I_V)^n(\mathbf{v}) = \mathbf{0}$$

So by DEFINITION NLT, $T|_{\mathcal{G}_T(\lambda)} - \lambda I_{\mathcal{G}_T(\lambda)}$ is nilpotent. \square

The proof of Theorem 3.9 shows that the index of the linear transformation $T|_{\mathcal{G}_T(\lambda)} - \lambda I_{\mathcal{G}_T(\lambda)}$ is less than or equal to the dimension of V . In practice, it must be less than or equal to the dimension of the domain, $\mathcal{G}_T(\lambda)$. In any event, the exact value of this index will be of some interest, so we define it now. Notice that this is a property of the eigenvalue λ . In many ways it is similar to the algebraic and geometric multiplicities of an eigenvalue (DEFINITION AME, DEFINITION GME).

Definition 16. Suppose $T: V \rightarrow V$ is a linear transformation with eigenvalue λ . Then the **index** of λ , $\iota_T(\lambda)$, is the index of the nilpotent linear transformation $T|_{\mathcal{G}_T(\lambda)} - \lambda I_{\mathcal{G}_T(\lambda)}$.

Example 22. In Example 17 we computed the generalized eigenspaces of the linear transformation $S: \mathbb{C}^6 \rightarrow \mathbb{C}^6$ defined by $S(\mathbf{x}) = B\mathbf{x}$ where

$$B = \begin{bmatrix} 2 & -4 & 25 & -54 & 90 & -37 \\ 2 & -3 & 4 & -16 & 26 & -8 \\ 2 & -3 & 4 & -15 & 24 & -7 \\ 10 & -18 & 6 & -36 & 51 & -2 \\ 8 & -14 & 0 & -21 & 28 & 4 \\ 5 & -7 & -6 & -7 & 8 & 7 \end{bmatrix}$$

The generalized eigenspace $\mathcal{G}_S(3)$ has dimension 2, while $\mathcal{G}_S(-1)$ has dimension 4. We will investigate each thoroughly in turn, with the intent being to illustrate Theorem 3.9. Many of our computations will be repeats of those done in Example 18.

For $U = \mathcal{G}_S(3)$ we compute a matrix representation of $S|_U$ using the basis found in Example 17,

$$D = \{\mathbf{u}_1, \mathbf{u}_2\} = \left\{ \left(\begin{bmatrix} 4 \\ 1 \\ 1 \\ 2 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} -5 \\ -1 \\ -1 \\ -1 \\ 0 \\ 1 \end{bmatrix} \right) \right\}$$

Since D has size 2, we obtain a 2×2 matrix representation from

$$\begin{aligned} \rho_D(S|_U(\mathbf{u}_1)) &= \rho_D \left(\begin{bmatrix} 11 \\ 3 \\ 3 \\ 7 \\ 4 \\ 1 \end{bmatrix} \right) = \rho_D(4\mathbf{u}_1 + \mathbf{u}_2) = \begin{bmatrix} 4 \\ 1 \end{bmatrix} \\ \rho_D(S|_U(\mathbf{u}_2)) &= \rho_D \left(\begin{bmatrix} -14 \\ -3 \\ -3 \\ -4 \\ -1 \\ 2 \end{bmatrix} \right) = \rho_D((-1)\mathbf{u}_1 + 2\mathbf{u}_2) = \begin{bmatrix} -1 \\ 2 \end{bmatrix} \end{aligned}$$

Thus

$$M = M_{U,U}^{S|_U} = \begin{bmatrix} 4 & -1 \\ 1 & 2 \end{bmatrix}$$

Now we can illustrate Theorem 3.9 with powers of the matrix representation (rather than the restriction itself),

$$\begin{aligned} M - 3I_2 &= \begin{bmatrix} 1 & -1 \\ 1 & -1 \end{bmatrix} \\ (M - 3I_2)^2 &= \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix} \end{aligned}$$

So $M - 3I_2$ is a nilpotent matrix of index 2 (meaning that $S|_U - 3I_U$ is a nilpotent linear transformation of index 2) and according to Definition 16 we say $\iota_S(3) = 2$.

For $W = \mathcal{G}_S(-1)$ we compute a matrix representation of $S|_W$ using the basis found in Example 17,

$$E = \{\mathbf{w}_1, \mathbf{w}_2, \mathbf{w}_3, \mathbf{w}_4\} = \left\{ \begin{bmatrix} 5 \\ 3 \\ 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} -2 \\ -3 \\ 0 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 4 \\ 5 \\ 0 \\ 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} -5 \\ -3 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} \right\}$$

Since E has size 4, we obtain a 4×4 matrix representation (DEFINITION MR)

from

$$\begin{aligned} \rho_E(S|_W(\mathbf{w}_1)) &= \rho_E \begin{pmatrix} 23 \\ 5 \\ 5 \\ 2 \\ -2 \\ -2 \end{pmatrix} = \rho_E(5\mathbf{w}_1 + 2\mathbf{w}_2 + (-2)\mathbf{w}_3 + (-2)\mathbf{w}_4) = \begin{bmatrix} 5 \\ 2 \\ -2 \\ -2 \end{bmatrix} \\ \rho_E(S|_W(\mathbf{w}_2)) &= \rho_E \begin{pmatrix} -46 \\ -11 \\ -10 \\ -2 \\ 5 \\ 4 \end{pmatrix} = \rho_E((-10)\mathbf{w}_1 + (-2)\mathbf{w}_2 + 5\mathbf{w}_3 + 4\mathbf{w}_4) = \begin{bmatrix} -10 \\ -2 \\ 5 \\ 4 \end{bmatrix} \\ \rho_E(S|_W(\mathbf{w}_3)) &= \rho_E \begin{pmatrix} 78 \\ 19 \\ 17 \\ 1 \\ -10 \\ -7 \end{pmatrix} = \rho_E(17\mathbf{w}_1 + \mathbf{w}_2 + (-10)\mathbf{w}_3 + (-7)\mathbf{w}_4) = \begin{bmatrix} 17 \\ 1 \\ -10 \\ -7 \end{bmatrix} \\ \rho_E(S|_W(\mathbf{w}_4)) &= \rho_E \begin{pmatrix} -35 \\ -9 \\ -8 \\ 2 \\ 6 \\ 3 \end{pmatrix} = \rho_E((-8)\mathbf{w}_1 + 2\mathbf{w}_2 + 6\mathbf{w}_3 + 3\mathbf{w}_4) = \begin{bmatrix} -8 \\ 2 \\ 6 \\ 3 \end{bmatrix} \end{aligned}$$

Thus

$$N = M_{W,W}^{S|_W} = \begin{bmatrix} 5 & -10 & 17 & -8 \\ 2 & -2 & 1 & 2 \\ -2 & 5 & -10 & 6 \\ -2 & 4 & -7 & 3 \end{bmatrix}$$

Now we can illustrate Theorem 3.9 with powers of the matrix representation (rather than the restriction itself),

$$\begin{aligned} N - (-1)I_4 &= \begin{bmatrix} 6 & -10 & 17 & -8 \\ 2 & -1 & 1 & 2 \\ -2 & 5 & -9 & 6 \\ -2 & 4 & -7 & 4 \end{bmatrix} \\ (N - (-1)I_4)^2 &= \begin{bmatrix} -2 & 3 & -5 & 2 \\ 4 & -6 & 10 & -4 \\ 4 & -6 & 10 & -4 \\ 2 & -3 & 5 & -2 \end{bmatrix} \\ (N - (-1)I_4)^3 &= \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \end{aligned}$$

So $N - (-1)I_4$ is a nilpotent matrix of index 3 (meaning that $S|_W - (-1)I_W$ is a nilpotent linear transformation of index 3) and according to Definition 16 we say $\iota_S(-1) = 3$.

Notice that if we were to take the union of the two bases of the generalized eigenspaces, we would have a basis for \mathbb{C}^6 . Then a matrix representation of S relative

to this basis would be the same block diagonal matrix we found in Example 18, only we now understand each of these blocks as being very close to being a nilpotent matrix.

3.2.4 Jordan Blocks

We conclude this section about nilpotent linear transformations with an infinite family of nilpotent matrices and a doubly-infinite family of nearly nilpotent matrices.

Definition 17. Given the scalar $\lambda \in \mathbb{C}$, the Jordan block $J_n(\lambda)$ is the $n \times n$ matrix defined by

$$[J_n(\lambda)]_{ij} = \begin{cases} \lambda & i = j \\ 1 & j = i + 1 \\ 0 & \text{otherwise} \end{cases}$$

Example 23. A simple example of a Jordan block,

$$J_4(5) = \begin{bmatrix} 5 & 1 & 0 & 0 \\ 0 & 5 & 1 & 0 \\ 0 & 0 & 5 & 1 \\ 0 & 0 & 0 & 5 \end{bmatrix}$$

We will return to general Jordan blocks later, but in this section we are only interested in Jordan blocks where $\lambda = 0$. (But notice that $J_n(\lambda) - \lambda I_n = J_n(0)$.) Here is an example of why we are specializing in the $\lambda = 0$ case now.

Example 24. Consider

$$J_5(0) = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

and compute powers,

$$\begin{aligned} (J_5(0))^2 &= \begin{bmatrix} 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} \\ (J_5(0))^3 &= \begin{bmatrix} 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} \\ (J_5(0))^4 &= \begin{bmatrix} 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} \\ (J_5(0))^5 &= \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} \end{aligned}$$

So $J_5(0)$ is nilpotent of index 5. As before, we record some information about the eigenvalues and eigenvectors of this matrix. The only eigenvalue is zero, with algebraic multiplicity 5, the maximum possible (THEOREM ME). The geometric multiplicity of this eigenvalue is just 1, the minimum possible (THEOREM ME), as seen in the eigenspace,

$$\mathcal{E}_{J_5(0)}(0) = \left\langle \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \right\rangle$$

There should not be any real surprises in this example. We can watch the ones in the powers of $J_5(0)$ slowly march off to the upper-right hand corner of the powers. Or we can watch the columns of the identity matrix march right, falling off the edge as they go. In some vague way, the eigenvalues and eigenvectors of this matrix are equally extreme.

We can form combinations of Jordan blocks to build a variety of nilpotent matrices. Simply create a block diagonal matrix, where each block is a Jordan block.

Example 25. Consider the matrix

$$C = \begin{bmatrix} J_3(0) & \mathcal{O} & \mathcal{O} \\ \mathcal{O} & J_3(0) & \mathcal{O} \\ \mathcal{O} & \mathcal{O} & J_2(0) \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

and compute powers,

$$C^2 = \begin{bmatrix} 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

$$C^3 = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

So C is nilpotent of index 3. You should notice how block diagonal matrices behave in products (much like diagonal matrices) and that it was the largest Jordan block that determined the index of this combination. All eight eigenvalues are zero, and each of the three Jordan blocks contributes one eigenvector to a basis for the eigenspace, resulting in zero having a geometric multiplicity of 3.

Since nilpotent matrices only have zero as an eigenvalue (Theorem 3.6), the algebraic multiplicity will be the maximum possible. However, by creating block diagonal matrices with Jordan blocks on the diagonal you should be able to attain any desired geometric multiplicity for this lone eigenvalue. Likewise, the size of the largest Jordan block employed will determine the index of the matrix. So nilpotent matrices with various combinations of index, geometric multiplicity and algebraic multiplicity are easy to manufacture. The predictable properties of block diagonal matrices in matrix products and eigenvector computations, along with the next theorem, make this possible. You might find EXAMPLE NJB5 a useful companion to this proof.

Theorem 3.10. *The Jordan block $J_n(0)$ is nilpotent of index n .*

Proof. We need to establish a specific matrix is nilpotent of a specified index. The first column of $J_n(0)$ is the zero vector, and the remaining $n - 1$ columns are the standard unit vectors \mathbf{e}_i , $1 \leq i \leq n - 1$ (DEFINITION SUV), which are also the first $n - 1$ columns of the size n identity matrix I_n . As shorthand, write $J = J_n(0)$.

$$J = [\mathbf{0} \mid \mathbf{e}_1 \mid \mathbf{e}_2 \mid \mathbf{e}_3 \mid \dots \mid \mathbf{e}_{n-1}]$$

We will use the definition of matrix multiplication (DEFINITION MM), together with a proof by induction, to study the powers of J . Our claim is that

$$J^k = [\mathbf{0} \mid \mathbf{0} \mid \dots \mid \mathbf{0} \mid \mathbf{e}_1 \mid \mathbf{e}_2 \mid \dots \mid \mathbf{e}_{n-k}] \text{ for } 0 \leq k \leq n$$

For the base case, $k = 0$, and the definition of $J^0 = I_n$ establishes the claim.

For the induction step, first note that $J\mathbf{e}_1 = \mathbf{0}$ and $J\mathbf{e}_i = \mathbf{e}_{i-1}$ for $2 \leq i \leq n$.

Then, assuming the claim is true for k , we examine the $k + 1$ case,

$$\begin{aligned}
 J^{k+1} &= JJ^k \\
 &= J[\mathbf{0} | \mathbf{0} | \dots | \mathbf{0} | \mathbf{e}_1 | \mathbf{e}_2 | \dots | \mathbf{e}_{n-k}] \\
 &= [J\mathbf{0} | J\mathbf{0} | \dots | J\mathbf{0} | J\mathbf{e}_1 | J\mathbf{e}_2 | \dots | J\mathbf{e}_{n-k}] \\
 &= [\mathbf{0} | \mathbf{0} | \dots | \mathbf{0} | \mathbf{0} | \mathbf{e}_1 | \mathbf{e}_2 | \dots | \mathbf{e}_{n-k-1}] \\
 &= [\mathbf{0} | \mathbf{0} | \dots | \mathbf{0} | \mathbf{e}_1 | \mathbf{e}_2 | \dots | \mathbf{e}_{n-(k+1)}]
 \end{aligned}$$

This concludes the induction.

So J^k has a nonzero entry (a one) in row $n - k$ and column n , for $0 \leq k \leq n - 1$, and is therefore a nonzero matrix. However,

$$J^n = [\mathbf{0} | \mathbf{0} | \dots | \mathbf{0}] = \mathcal{O}$$

Thus, by 15, J is nilpotent of index n . □

3.3 Jordan Canonical Form

Nilpotent matrices and generalized eigenspaces are the essential ingredients for a canonical form applicable to any square matrix. In this section will progress from the specialized case of a nilpotent matrix to the totally general case of any square matrix.

3.3.1 Canonical Form for Nilpotent Linear Transformations

Our main purpose in this section is to find a basis so that a nilpotent linear transformation will have a pleasing, nearly-diagonal matrix representation. Of course, we will not have a definition for “pleasing,” nor for “nearly-diagonal.” But the short answer is that our preferred matrix representation will be built up from Jordan blocks, $J_k(0)$. Here’s the theorem. You will find Example $\langle\langle$ example-canonical-form-nilpotent-size-6-index-4 $\rangle\rangle$, just following, helpful as you study this proof, since it uses the same notation, and is large enough to (barely) illustrate the full generality of the theorem.

Theorem 3.11. *Suppose that $T: V \rightarrow V$ is a nilpotent linear transformation of index p . Then there is a basis for V so that the matrix representation, $M_{B,B}^T$, is block diagonal with each block being a Jordan block, $J_k(0)$. The size of the largest block is the index p , and the total number of blocks is the nullity of T , $n(T)$.*

Proof. The proof is constructive, as we will explicitly manufacture the basis, and so can be used in practice. As we begin, the basis vectors will not be in the proper order, but we will rearrange them at the end of the proof. For convenience, define $n_i = n(T^i)$, so for example, $n_0 = 0$, $n_1 = n(T)$ and $n_p = n(T^p) = \dim(V)$. Define $s_i = n_i - n_{i-1}$, for $1 \leq i \leq p$, so we can think of s_i as “how much bigger” $\mathcal{K}(T^i)$ is than $\mathcal{K}(T^{i-1})$. In particular, Theorem 3.8 implies that $s_i > 0$ for $1 \leq i \leq p$.

We build a set of vectors $\mathbf{z}_{i,j}$, $1 \leq i \leq p$, $1 \leq j \leq s_i$. Each $\mathbf{z}_{i,j}$ will be an element of $\mathcal{K}(T^i)$ and not an element of $\mathcal{K}(T^{i-1})$. In total, we will obtain a linearly independent set of $\sum_{i=1}^p s_i = \sum_{i=1}^p n_i - n_{i-1} = n_p - n_0 = \dim(V)$ vectors that form a basis of V . We construct this set in pieces, starting at the “wrong” end. Our procedure will build a series of subspaces, Z_i , each lying in between $\mathcal{K}(T^{i-1})$ and $\mathcal{K}(T^i)$, having bases $\mathbf{z}_{i,j}$, $1 \leq j \leq s_i$, and which together equal V as a direct sum. Now would be a good time to review the results on direct sums collected in Section 1.2.

We build the subspace Z_p first (this is what we meant by “starting at the wrong end”). $\mathcal{K}(T^{p-1})$ is a proper subspace of $\mathcal{K}(T^p) = V$ (Theorem 3.8). Theorem 1.2 says that there is a subspace of V that will pair with the subspace $\mathcal{K}(T^{p-1})$ to form a direct sum of V . Call this subspace Z_p , and choose vectors $\mathbf{z}_{p,j}$, $1 \leq j \leq s_p$ as a basis of Z_p , which we will denote as B_p . Note that we have a fair amount of freedom in how to choose these first basis vectors. Several observations will be useful in the next step. First $V = \mathcal{K}(T^{p-1}) \oplus Z_p$. The basis $B_p = \{\mathbf{z}_{p,1}, \mathbf{z}_{p,2}, \mathbf{z}_{p,3}, \dots, \mathbf{z}_{p,s_p}\}$ is linearly independent. For $1 \leq j \leq s_p$, $\mathbf{z}_{p,j} \in \mathcal{K}(T^p) = V$. Since the two subspaces of a direct sum have no nonzero vectors in common, for $1 \leq j \leq s_p$, $\mathbf{z}_{p,j} \notin \mathcal{K}(T^{p-1})$. That was comparably easy.

If obtaining Z_p was easy, getting Z_{p-1} will be harder. We will repeat the next step $p-1$ times, and so will do it carefully the first time. Eventually, Z_{p-1} will have dimension s_{p-1} . However, obtaining the first s_p vectors of a basis for Z_{p-1} are straightforward. Define $\mathbf{z}_{p-1,j} = T(\mathbf{z}_{p,j})$, $1 \leq j \leq s_p$. Notice that we have no choice in creating these vectors, they are a consequence of our choices for $\mathbf{z}_{p,j}$. In retrospect (i.e. on a second reading of this proof), you will recognize this as the key step in realizing a matrix representation of a nilpotent linear transformation with Jordan blocks. We need to know that this set of vectors is linearly independent. We consider a relation of linear dependence on $\mathbf{z}_{p-1,j}$, $1 \leq j \leq s_p$, and massage it,

$$\begin{aligned} \mathbf{0} &= a_1 \mathbf{z}_{p-1,1} + a_2 \mathbf{z}_{p-1,2} + a_3 \mathbf{z}_{p-1,3} + \dots + a_{s_p} \mathbf{z}_{p-1,s_p} \\ &= a_1 T(\mathbf{z}_{p,1}) + a_2 T(\mathbf{z}_{p,2}) + a_3 T(\mathbf{z}_{p,3}) + \dots + a_{s_p} T(\mathbf{z}_{p,s_p}) \\ &= T(a_1 \mathbf{z}_{p,1} + a_2 \mathbf{z}_{p,2} + a_3 \mathbf{z}_{p,3} + \dots + a_{s_p} \mathbf{z}_{p,s_p}) \end{aligned}$$

Define

$$\mathbf{x} = a_1 \mathbf{z}_{p,1} + a_2 \mathbf{z}_{p,2} + a_3 \mathbf{z}_{p,3} + \dots + a_{s_p} \mathbf{z}_{p,s_p}$$

The statement just above means that $\mathbf{x} \in \mathcal{K}(T) \subseteq \mathcal{K}(T^{p-1})$ (Theorem 3.8). As defined, \mathbf{x} is a linear combination of the basis vectors B_p , and therefore $\mathbf{x} \in Z_p$. Thus $\mathbf{x} \in \mathcal{K}(T^{p-1}) \cap Z_p$. Because $V = \mathcal{K}(T^{p-1}) \oplus Z_p$, Theorem 1.4 tells us that $\mathbf{x} = \mathbf{0}$. Now we recognize the definition of \mathbf{x} as a relation of linear dependence on the linearly independent set B_p , and therefore conclude that $a_1 = a_2 = \dots = a_{s_p} = 0$. This establishes the linear independence of $\mathbf{z}_{p-1,j}$, $1 \leq j \leq s_p$.

We also need to know where the vectors $\mathbf{z}_{p-1,j}$, $1 \leq j \leq s_p$ live. First we demonstrate that they are members of $\mathcal{K}(T^{p-1})$.

$$T^{p-1}(\mathbf{z}_{p-1,j}) = T^{p-1}(T(\mathbf{z}_{p,j})) = T^p(\mathbf{z}_{p,j}) = \mathbf{0}$$

So $\mathbf{z}_{p-1,j} \in \mathcal{K}(T^{p-1})$, $1 \leq j \leq s_p$.

Moreover, these vectors are not elements of $\mathcal{K}(T^{p-2})$. Suppose to the contrary that $\mathbf{z}_{p-1,j} \in \mathcal{K}(T^{p-2})$. Then

$$\mathbf{0} = T^{p-2}(\mathbf{z}_{p-1,j}) = T^{p-2}(T(\mathbf{z}_{p,j})) = T^{p-1}(\mathbf{z}_{p,j})$$

which contradicts the earlier statement that $\mathbf{z}_{p,j} \notin \mathcal{K}(T^{p-1})$. So $\mathbf{z}_{p-1,j} \notin \mathcal{K}(T^{p-2})$, $1 \leq j \leq s_p$.

Now choose any basis $C_{p-2} = \{\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3, \dots, \mathbf{u}_{n_{p-2}}\}$ for $\mathcal{K}(T^{p-2})$. We want to extend this basis by adding in the $\mathbf{z}_{p-1,j}$ to span a subspace of $\mathcal{K}(T^{p-1})$. But first we establish that this set is linearly independent. Let a_k , $1 \leq k \leq n_{p-2}$ and b_j , $1 \leq j \leq s_p$ be the scalars in a relation of linear dependence,

$$\mathbf{0} = a_1 \mathbf{u}_1 + a_2 \mathbf{u}_2 + \dots + a_{n_{p-2}} \mathbf{u}_{n_{p-2}} + b_1 \mathbf{z}_{p-1,1} + b_2 \mathbf{z}_{p-1,2} + \dots + b_{s_p} \mathbf{z}_{p-1,s_p}$$

Then,

$$\begin{aligned}
\mathbf{0} &= T^{p-2}(\mathbf{0}) \\
&= T^{p-2}(a_1\mathbf{u}_1 + a_2\mathbf{u}_2 + \cdots + a_{n_{p-2}}\mathbf{u}_{n_{p-2}} + b_1\mathbf{z}_{p-1,1} + b_2\mathbf{z}_{p-1,2} + \cdots + b_{s_p}\mathbf{z}_{p-1,s_p}) \\
&= a_1T^{p-2}(\mathbf{u}_1) + a_2T^{p-2}(\mathbf{u}_2) + \cdots + a_{n_{p-2}}T^{p-2}(\mathbf{u}_{n_{p-2}}) + \\
&\quad b_1T^{p-2}(\mathbf{z}_{p-1,1}) + b_2T^{p-2}(\mathbf{z}_{p-1,2}) + \cdots + b_{s_p}T^{p-2}(\mathbf{z}_{p-1,s_p}) \\
&= a_1\mathbf{0} + a_2\mathbf{0} + \cdots + a_{n_{p-2}}\mathbf{0} + \\
&\quad b_1T^{p-2}(\mathbf{z}_{p-1,1}) + b_2T^{p-2}(\mathbf{z}_{p-1,2}) + \cdots + b_{s_p}T^{p-2}(\mathbf{z}_{p-1,s_p}) \\
&= b_1T^{p-2}(\mathbf{z}_{p-1,1}) + b_2T^{p-2}(\mathbf{z}_{p-1,2}) + \cdots + b_{s_p}T^{p-2}(\mathbf{z}_{p-1,s_p}) \\
&= b_1T^{p-2}(T(\mathbf{z}_{p,1})) + b_2T^{p-2}(T(\mathbf{z}_{p,2})) + \cdots + b_{s_p}T^{p-2}(T(\mathbf{z}_{p,s_p})) \\
&= b_1T^{p-1}(\mathbf{z}_{p,1}) + b_2T^{p-1}(\mathbf{z}_{p,2}) + \cdots + b_{s_p}T^{p-1}(\mathbf{z}_{p,s_p}) \\
&= T^{p-1}(b_1\mathbf{z}_{p,1} + b_2\mathbf{z}_{p,2} + \cdots + b_{s_p}\mathbf{z}_{p,s_p})
\end{aligned}$$

Define

$$\mathbf{y} = b_1\mathbf{z}_{p,1} + b_2\mathbf{z}_{p,2} + \cdots + b_{s_p}\mathbf{z}_{p,s_p}$$

The statement just above means that $\mathbf{y} \in \mathcal{K}(T^{p-1})$. As defined, \mathbf{y} is a linear combination of the basis vectors B_p , and therefore $\mathbf{y} \in Z_p$. Thus $\mathbf{y} \in \mathcal{K}(T^{p-1}) \cap Z_p$. Because $V = \mathcal{K}(T^{p-1}) \oplus Z_p$, Theorem 1.4 tells us that $\mathbf{y} = \mathbf{0}$. Now we recognize the definition of \mathbf{y} as a relation of linear dependence on the linearly independent set B_p , and therefore $b_1 = b_2 = \cdots = b_{s_p} = 0$. Return to the full relation of linear dependence with both sets of scalars (the a_i and b_j). Now that we know that $b_j = 0$ for $1 \leq j \leq s_p$, this relation of linear dependence simplifies to a relation of linear dependence on just the basis C_{p-1} . Therefore, $a_i = 0$, $1 \leq i \leq n_{p-1}$ and we have the desired linear independence.

Define a new subspace of $\mathcal{K}(T^{p-1})$ by

$$Q_{p-1} = \langle \{ \mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3, \dots, \mathbf{u}_{n_{p-1}}, \mathbf{z}_{p-1,1}, \mathbf{z}_{p-1,2}, \mathbf{z}_{p-1,3}, \dots, \mathbf{z}_{p-1,s_p} \} \rangle$$

By Theorem 1.2 there exists a subspace of $\mathcal{K}(T^{p-1})$ which will pair with Q_{p-1} to form a direct sum. Call this subspace R_{p-1} , so by definition, $\mathcal{K}(T^{p-1}) = Q_{p-1} \oplus R_{p-1}$. We are interested in the dimension of R_{p-1} . Note first, that since the spanning set of Q_{p-1} is linearly independent, $\dim(Q_{p-1}) = n_{p-2} + s_p$. Then

$$\begin{aligned}
\dim(R_{p-1}) &= \dim(\mathcal{K}(T^{p-1})) - \dim(Q_{p-1}) \\
&= n_{p-1} - (n_{p-2} + s_p) = (n_{p-1} - n_{p-2}) - s_p = s_{p-1} - s_p
\end{aligned}$$

Notice that if $s_{p-1} = s_p$, then R_{p-1} is trivial. Now choose a basis of R_{p-1} , and denote these $s_{p-1} - s_p$ vectors as $\mathbf{z}_{p-1,s_p+1}, \mathbf{z}_{p-1,s_p+2}, \mathbf{z}_{p-1,s_p+3}, \dots, \mathbf{z}_{p-1,s_{p-1}}$. This is another occasion to notice that we have some freedom in this choice.

We now have $\mathcal{K}(T^{p-1}) = Q_{p-1} \oplus R_{p-1}$, and we have bases for each of the two subspaces. The union of these two bases will therefore be a linearly independent set in $\mathcal{K}(T^{p-1})$ with size

$$\begin{aligned}
(n_{p-2} + s_p) + (s_{p-1} - s_p) &= n_{p-2} + s_{p-1} = n_{p-2} + n_{p-1} - n_{p-2} \\
&= n_{p-1} = \dim(\mathcal{K}(T^{p-1}))
\end{aligned}$$

So with the proper size the following set is a basis of $\mathcal{K}(T^{p-1})$,

$$\{ \mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_{n_{p-2}}, \mathbf{z}_{p-1,1}, \mathbf{z}_{p-1,2}, \dots, \mathbf{z}_{p-1,s_p}, \mathbf{z}_{p-1,s_p+1}, \mathbf{z}_{p-1,s_p+2}, \dots, \mathbf{z}_{p-1,s_{p-1}} \}$$

We built up this basis in three parts, we will now split it in half. Define the subspace Z_{p-1} by

$$Z_{p-1} = \langle B_{p-1} \rangle = \langle \{ \mathbf{z}_{p-1,1}, \mathbf{z}_{p-1,2}, \dots, \mathbf{z}_{p-1,s_{p-1}} \} \rangle$$

where we have implicitly denoted the basis as B_{p-1} . Then 1.1 allows us to split up the basis for $\mathcal{K}(T^{p-1})$ as $C_{p-1} \cup B_{p-1}$ and write

$$\mathcal{K}(T^{p-1}) = \mathcal{K}(T^{p-2}) \oplus Z_{p-1} >$$

Whew! This is a good place to recap what we have achieved. The vectors $\mathbf{z}_{i,j}$ form bases for the subspaces Z_i and right now

$$V = \mathcal{K}(T^{p-1}) \oplus Z_p = \mathcal{K}(T^{p-2}) \oplus Z_{p-1} \oplus Z_p$$

The key feature of this decomposition of V is that the first s_p vectors in the basis for Z_{p-1} are outputs of the linear transformation T using the basis vectors of Z_p as inputs.

Now we want to further decompose $\mathcal{K}(T^{p-2})$, into $\mathcal{K}(T^{p-3})$ and Z_{p-2} . The procedure is the same as above, so we will only sketch the key steps. Checking the details proceeds in the same manner as above. Technically, we could have set up the preceding as the induction step in a proof by induction (PROOF TECHNIQUE I), but this probably would make the proof harder to understand.

Apply T to each element of B_{p-1} , to create vectors $\mathbf{z}_{p-2,j}$, $1 \leq j \leq s_{p-1}$. These vectors form a linearly independent set, and each is an element of $\mathcal{K}(T^{p-2})$, but not an element of $\mathcal{K}(T^{p-3})$. Grab a basis C_{p-3} of $\mathcal{K}(T^{p-3})$ and add on the newly-created vectors $\mathbf{z}_{p-2,j}$, $1 \leq j \leq s_{p-1}$. This expanded set is linearly independent, and we can define a subspace Q_{p-2} with this set as its basis. Theorem 1.2 gives us a subspace R_{p-2} such that $\mathcal{K}(T^{p-2}) = Q_{p-2} \oplus R_{p-2}$. Vectors $\mathbf{z}_{p-2,j}$, $s_{p-1} + 1 \leq j \leq s_{p-2}$ are chosen as a basis for R_{p-2} once the relevant dimensions have been verified. The union of C_{p-3} and $\mathbf{z}_{p-2,j}$, $1 \leq j \leq s_{p-2}$ then form a basis of $\mathcal{K}(T^{p-2})$, which can be split into two parts to yield the decomposition

$$\mathcal{K}(T^{p-2}) = \mathcal{K}(T^{p-3}) \oplus Z_{p-2}$$

Here Z_{p-2} is the subspace of $\mathcal{K}(T^{p-2})$ with basis $B_{p-2} = \{ \mathbf{z}_{p-2,j} \mid 1 \leq j \leq s_{p-2} \}$. Finally,

$$V = \mathcal{K}(T^{p-1}) \oplus Z_p = \mathcal{K}(T^{p-2}) \oplus Z_{p-1} \oplus Z_p = \mathcal{K}(T^{p-3}) \oplus Z_{p-2} \oplus Z_{p-1} \oplus Z_p$$

Again, the key feature of this decomposition is that the first vectors in the basis of Z_{p-2} are outputs of T using vectors from the basis Z_{p-1} as inputs (and in turn, some of these inputs are outputs of T derived from inputs in Z_p).

Now assume we repeat this procedure until we decompose $\mathcal{K}(T^2)$ into subspaces $\mathcal{K}(T)$ and Z_2 . Finally, decompose $\mathcal{K}(T)$ into subspaces $\mathcal{K}(T^0) = \mathcal{K}(I_n) = \{ \mathbf{0} \}$ and Z_1 , so that we recognize the vectors $\mathbf{z}_{i,j}$, $1 \leq j \leq s_1 = n_1$ as elements of $\mathcal{K}(T)$. The set

$$B = B_1 \cup B_2 \cup B_3 \cup \dots \cup B_p = \{ \mathbf{z}_{i,j} \mid 1 \leq i \leq p, 1 \leq j \leq s_i \}$$

is linearly independent by Theorem 1.5 and has size

$$\sum_{i=1}^p s_i = \sum_{i=1}^p n_i - n_{i-1} = n_p - n_0 = \dim(V)$$

So B is a basis of V .

We desire a matrix representation of T relative to B , but first we will reorder the elements of B . The following display lists the elements of B in the desired order,

when read across the rows left-to-right in the usual way. Notice that we established the existence of these vectors column-by-column, and beginning on the right.

$\mathbf{z}_{1,1}$	$\mathbf{z}_{2,1}$	$\mathbf{z}_{3,1}$	\cdots	$\mathbf{z}_{p,1}$
$\mathbf{z}_{1,2}$	$\mathbf{z}_{2,2}$	$\mathbf{z}_{3,2}$	\cdots	$\mathbf{z}_{p,2}$
	\vdots		\vdots	
\mathbf{z}_{1,s_p}	\mathbf{z}_{2,s_p}	\mathbf{z}_{3,s_p}	\cdots	\mathbf{z}_{p,s_p}
\mathbf{z}_{1,s_p+1}	\mathbf{z}_{2,s_p+1}	\mathbf{z}_{3,s_p+1}	\cdots	
	\vdots		\vdots	
\mathbf{z}_{1,s_3}	\mathbf{z}_{2,s_3}	\mathbf{z}_{3,s_3}		
	\vdots			
\mathbf{z}_{1,s_2}	\mathbf{z}_{2,s_2}			
	\vdots			
\mathbf{z}_{1,s_1}				

It is difficult to layout this table with the notation we have been using, but it would not be especially useful to invent some notation to overcome the difficulty. (One approach would be to define something like the inverse of the nonincreasing function, $i \rightarrow s_i$.) Do notice that there are $s_1 = n_1$ rows, and p columns. Column i is the basis B_i . The vectors in the first column are elements of $\mathcal{K}(T)$. Each row is the same length, or shorter, than the one above it. If we apply T to any vector in the table, other than those in the first column, the output is the preceding vector in the row.

Now contemplate the matrix representation of T relative to B as we read across the rows of the table above. In the first row, $T(\mathbf{z}_{1,1}) = \mathbf{0}$, so the first column of the representation is the zero column. Next, $T(\mathbf{z}_{2,1}) = \mathbf{z}_{1,1}$, so the second column of the representation is a vector with a single one in the first entry, and zeros elsewhere. Next, $T(\mathbf{z}_{3,1}) = \mathbf{z}_{2,1}$, so column 3 of the representation is a zero, then a one, then all zeros. Continuing in this vein, we obtain the first p columns of the representation, which is the Jordan block $J_p(0)$ followed by rows of zeros.

When we apply T to the basis vectors of the second row, what happens? Applying T to the first vector, the result is the zero vector, so the representation gets a zero column. Applying T to the second vector in the row, the output is simply the first vector in that row, making the next column of the representation all zeros plus a lone one, sitting just above the diagonal. Continuing, we create a Jordan block, sitting on the diagonal of the matrix representation. It is not possible in general to state the size of this block, but since the second row is no longer than the first, it cannot have size larger than p .

Since there are as many rows as the dimension of $\mathcal{K}(T)$, the representation contains as many Jordan blocks as the nullity of T , $n(T)$. Each successive block is smaller than the preceding one, with the first, and largest, having size p . The blocks are Jordan blocks since the basis vectors $\mathbf{z}_{i,j}$ were often defined as the result of applying T to other elements of the basis already determined, and then we rearranged the basis into an order that placed outputs of T just before their inputs, excepting the start of each row, which was an element of $\mathcal{K}(T)$. \square

The proof of Theorem 3.11 is constructive, so we can use it to create bases of nilpotent linear transformations with pleasing matrix representations. Recall that

Theorem 3.7 told us that nilpotent linear transformations are almost never diagonalizable, so this is progress. As we have hinted before, with a nice representation of nilpotent matrices, it will not be difficult to build up representations of other non-diagonalizable matrices. Here is the promised example which illustrates the previous theorem. It is a useful companion to your study of the proof of Theorem 3.11.

Example 26. The 6×6 matrix, A , of Example 19 is nilpotent. If we define the linear transformation $T: \mathbb{C}^6 \rightarrow \mathbb{C}^6$ by $T(\mathbf{x}) = A\mathbf{x}$, then T is nilpotent and we can seek a basis of \mathbb{C}^6 that yields a matrix representation with Jordan blocks on the diagonal. Since T has index 4 and nullity 2, from Theorem 3.11 we can expect the largest Jordan block to be $J_4(0)$, and there will be just 2 blocks. This only leaves enough room for the second block to have size 2.

To determine nullities, we will recycle the bases for the null spaces of the powers of A from Example 21, rather than recomputing them. We will also use the same notation used in the proof of Theorem 3.11.

To begin, $s_4 = n_4 - n_3 = 6 - 5 = 1$, so we need one vector of $\mathcal{K}(T^4) = \mathbb{C}^6$, that is not in $\mathcal{K}(T^3)$, to be a basis for Z_4 . We have a lot of latitude in this choice, and we have not described any sure-fire method for constructing a vector *outside* of a subspace. Looking at the basis for $\mathcal{K}(T^3)$ we see that if a vector is in this subspace, and has a nonzero value in the first entry, then it must also have a nonzero value in the fourth entry. So the vector

$$\mathbf{z}_{4,1} = \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}$$

will not be an element of $\mathcal{K}(T^3)$. (Notice that many other choices could be made here, so our basis will not be unique.) This completes the determination of $Z_p = Z_4$.

Next, $s_3 = n_3 - n_2 = 5 - 4 = 1$, so we again need just a single basis vector for Z_3 . We start by evaluating T with each basis vector of Z_4 ,

$$\mathbf{z}_{3,1} = T(\mathbf{z}_{4,1}) = A\mathbf{z}_{4,1} = \begin{bmatrix} -3 \\ -3 \\ -3 \\ -3 \\ -3 \\ -2 \end{bmatrix}$$

Since $s_3 = s_4$, the subspace R_3 is trivial, and there is nothing left to do, $\mathbf{z}_{3,1}$ is the lone basis vector of Z_3 .

Now $s_2 = n_2 - n_1 = 4 - 2 = 2$, so the construction of Z_2 will not be as simple as the construction of Z_3 . We first apply T to the basis vector of Z_2 ,

$$\mathbf{z}_{2,1} = T(\mathbf{z}_{3,1}) = A\mathbf{z}_{3,1} = \begin{bmatrix} 1 \\ 0 \\ 3 \\ 1 \\ 0 \\ -1 \end{bmatrix}$$

The two basis vectors of $\mathcal{K}(T^1)$, together with $\mathbf{z}_{2,1}$, form a basis for Q_2 . Because $\dim(\mathcal{K}(T^2)) - \dim(Q_2) = 4 - 3 = 1$ we need only find a single basis vector for R_2 .

This vector must be an element of $\mathcal{K}(T^2)$, but not an element of Q_2 . Again, there is a variety of vectors that fit this description, and we have no precise algorithm for finding them. Since they are plentiful, they are not too hard to find. We add up the four basis vectors of $\mathcal{K}(T^2)$, ensuring an element of $\mathcal{K}(T^2)$. Then we check to see if the vector is a linear combination of three vectors: the two basis vectors of $\mathcal{K}(T^1)$ and $\mathbf{z}_{2,1}$. Having passed the tests, we have chosen

$$\mathbf{z}_{2,2} = \begin{bmatrix} 2 \\ 1 \\ 2 \\ 2 \\ 2 \\ 1 \end{bmatrix}$$

Thus, $Z_2 = \langle \{\mathbf{z}_{2,1}, \mathbf{z}_{2,2}\} \rangle$.

Lastly, $s_1 = n_1 - n_0 = 2 - 0 = 2$. Since $s_2 = s_1$, we again have a trivial R_1 and need only complete our basis by evaluating the basis vectors of Z_2 with T ,

$$\mathbf{z}_{1,1} = T(\mathbf{z}_{2,1}) = A\mathbf{z}_{2,1} = \begin{bmatrix} 1 \\ 1 \\ 0 \\ 1 \\ 1 \\ 1 \end{bmatrix} \quad \mathbf{z}_{1,2} = T(\mathbf{z}_{2,2}) = A\mathbf{z}_{2,2} = \begin{bmatrix} -2 \\ -2 \\ -5 \\ -2 \\ -1 \\ 0 \end{bmatrix}$$

Now we reorder these vectors as the desired basis,

$$B = \{\mathbf{z}_{1,1}, \mathbf{z}_{2,1}, \mathbf{z}_{3,1}, \mathbf{z}_{4,1}, \mathbf{z}_{1,2}, \mathbf{z}_{2,2}\}$$

We now apply DEFINITION MR to build a matrix representation of T relative to

B ,

$$\begin{aligned} \rho_B(T(\mathbf{z}_{1,1})) = \rho_B(A\mathbf{z}_{1,1}) = \rho_B(\mathbf{0}) &= \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \\ \rho_B(T(\mathbf{z}_{2,1})) = \rho_B(A\mathbf{z}_{2,1}) = \rho_B(\mathbf{z}_{1,1}) &= \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \\ \rho_B(T(\mathbf{z}_{3,1})) = \rho_B(A\mathbf{z}_{3,1}) = \rho_B(\mathbf{z}_{2,1}) &= \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \\ \rho_B(T(\mathbf{z}_{4,1})) = \rho_B(A\mathbf{z}_{4,1}) = \rho_B(\mathbf{z}_{3,1}) &= \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} \\ \rho_B(T(\mathbf{z}_{1,2})) = \rho_B(A\mathbf{z}_{1,2}) = \rho_B(\mathbf{0}) &= \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \\ \rho_B(T(\mathbf{z}_{2,2})) = \rho_B(A\mathbf{z}_{2,2}) = \rho_B(\mathbf{z}_{1,2}) &= \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 1 \\ 0 \end{bmatrix} \end{aligned}$$

Installing these vectors as the columns of the matrix representation we have

$$M_{B,B}^T = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

which is a block diagonal matrix with Jordan blocks $J_4(0)$ and $J_2(0)$.

If we construct the matrix S having the vectors of B as columns, then THEOREM SCB tells us that a similarity transformation by S relates the original matrix representation of T with the matrix representation consisting of Jordan blocks, in other words, $S^{-1}AS = M_{B,B}^T$.

Notice that constructing interesting examples of matrix representations requires domains with dimensions bigger than just two or three. Going forward our examples will get even bigger.

3.3.2 Restrictions to Generalized Eigenspaces

We now know how to make canonical matrix representations of a what seems to be a narrow class of linear transformations—the nilpotent ones (Theorem 3.11). However, since the restriction of any linear transformation to one of its generalized eigenspace is only a small adjustment away from being a nilpotent linear transformation (Theorem 3.9) we can extend the utility of our previous representation easily.

Theorem 3.12. *Suppose that $T: V \rightarrow V$ is a linear transformation with eigenvalue λ . Then there is a basis of the the generalized eigenspace $\mathcal{G}_T(\lambda)$ such that the restriction $T|_{\mathcal{G}_T(\lambda)}: \mathcal{G}_T(\lambda) \rightarrow \mathcal{G}_T(\lambda)$ has a matrix representation that is block diagonal where each block is a Jordan block of the form $J_k(\lambda)$ (with varying values of k).*

Proof. Theorem 3.9 tells us that $T|_{\mathcal{G}_T(\lambda)} - \lambda I_{\mathcal{G}_T(\lambda)}$ is a nilpotent linear transformation. Theorem 3.11 tells us that a nilpotent linear transformation has a basis for its domain that yields a matrix representation that is block diagonal where the blocks are Jordan blocks of the form $J_k(0)$. So let B be a basis of $\mathcal{G}_T(\lambda)$ that yields such a matrix representation for $T|_{\mathcal{G}_T(\lambda)} - \lambda I_{\mathcal{G}_T(\lambda)}$.

We can write

$$T|_{\mathcal{G}_T(\lambda)} = (T|_{\mathcal{G}_T(\lambda)} - \lambda I_{\mathcal{G}_T(\lambda)}) + \lambda I_{\mathcal{G}_T(\lambda)}$$

Then the matrix representation of $\lambda I_{\mathcal{G}_T(\lambda)}$ relative to the basis B is then simply the diagonal matrix λI_m , where $m = \dim(\mathcal{G}_T(\lambda))$. So we have the rather unwieldy expression,

$$M_{B,B}^{T|_{\mathcal{G}_T(\lambda)}} = M_{B,B}^{(T|_{\mathcal{G}_T(\lambda)} - \lambda I_{\mathcal{G}_T(\lambda)}) + \lambda I_{\mathcal{G}_T(\lambda)}} = M_{B,B}^{T|_{\mathcal{G}_T(\lambda)} - \lambda I_{\mathcal{G}_T(\lambda)}} + \lambda M_{B,B}^{I_{\mathcal{G}_T(\lambda)}}$$

The first representation in the final expression has Jordan blocks with zero in every diagonal entry by Theorem 3.11, while the second representation has λ in every diagonal entry of the matrix. The result of adding the two representations is to convert the Jordan blocks from the form $J_k(0)$ to the form $J_k(\lambda)$. \square

Of course, Theorem 3.11 provides some extra information on the sizes of the Jordan blocks in a representation and we could carry over this information to Theorem 3.9, but we will save this description and incorporate it into our final major result in the next section.

3.3.3 Jordan Canonical Form

Begin with any linear transformation that has an identical domain and codomain. Build a block diagonal representation from a direct sum decomposition into (invariant) generalized eigenspaces. For each generalized eigenspace, further refine the block into a sequence of Jordan blocks (with common diagonal elements) from the restriction to the generalized eigenspace, which is very nearly nilpotent. Then you have Jordan canonical form. Other than cosmetic reorderings, it is a unique representative of the equivalence class of similar matrices.

We remove the ambiguity from trivial reorderings of eigenvalues and Jordan blocks with a careful definition.

Definition 18. A square matrix is in **Jordan canonical form** if it meets the following requirements:

1. The matrix is block diagonal.
2. Each block is a Jordan block.
3. If $\rho < \lambda$ then the block $J_k(\rho)$ occupies rows with indices greater than the indices of the rows occupied by $J_\ell(\lambda)$.
4. If $\rho = \lambda$ and $\ell < k$, then the block $J_\ell(\lambda)$ occupies rows with indices greater than the indices of the rows occupied by $J_k(\lambda)$.

Theorem 3.13. *Suppose $T: V \rightarrow V$ is a linear transformation. Then there is a basis B for V such that the matrix representation of T with the following properties:*

1. *The matrix representation is in Jordan canonical form.*
2. *If $J_k(\lambda)$ is one of the Jordan blocks, then λ is an eigenvalue of T .*
3. *For each eigenvalue λ , the largest block of the form $J_k(\lambda)$ has size equal to the index of λ , $\nu_T(\lambda)$.*
4. *For each eigenvalue λ , the number of blocks of the form $J_k(\lambda)$ is the geometric multiplicity of λ , $\gamma_T(\lambda)$.*
5. *For each eigenvalue λ , the number of rows occupied by blocks of the form $J_k(\lambda)$ is the algebraic multiplicity of λ , $\alpha_T(\lambda)$.*

Proof. This theorem is really just the consequence of applying to T , consecutively, Theorem 3.4, Theorem 3.12 and Theorem 3.11.

Theorem 3.4 gives us a decomposition of V into generalized eigenspaces, one for each distinct eigenvalue. Since these generalized eigenspaces are invariant relative to T , this provides a block diagonal matrix representation where each block is the matrix representation of the restriction of T to the generalized eigenspace.

Restricting T to a generalized eigenspace results in a “nearly nilpotent” linear transformation, as stated more precisely in Theorem ???. We unravel Theorem ??? in the proof of Theorem 3.12 so that we can apply Theorem 3.11 about representations of nilpotent linear transformation.

We know the dimension of a generalized eigenspace is the algebraic multiplicity of the eigenvalue (Corollary 3.5), so the blocks associated with the generalized eigenspaces are square with a size equal to the algebraic multiplicity. In refining the basis for a block associated with a single eigenvalue, and producing Jordan blocks, the results of Theorem 3.11 apply. The total number of blocks will be the nullity of $T|_{\mathcal{G}_T(\lambda)} - \lambda I_{\mathcal{G}_T(\lambda)}$, which is the geometric multiplicity of λ as an eigenvalue of T (DEFINITION GME). The largest of the Jordan blocks will have size equal to the index of the nilpotent linear transformation $T|_{\mathcal{G}_T(\lambda)} - \lambda I_{\mathcal{G}_T(\lambda)}$, which is exactly the definition of the index of the eigenvalue λ (Definition 16). \square

Before we do some examples of this result, notice how close Jordan canonical form is to a diagonal matrix. Or, equivalently, notice how close we have come to diagonalizing a matrix (DEFINITION DZM). We have a matrix representation which has diagonal entries that are the eigenvalues of a matrix. Each occurs on the diagonal as many times as the algebraic multiplicity. However, when the geometric multiplicity is strictly less than the algebraic multiplicity, we have some entries in the representation just above the diagonal (the “superdiagonal”). Furthermore, we

have some idea how often this happens if we know the geometric multiplicity and the index of the eigenvalue.

We now recognize just how plain a diagonalizable linear transformation really is. For each eigenvalue, the generalized eigenspace is just the regular eigenspace, and it decomposes into a direct sum of one-dimensional subspaces, each spanned by a different eigenvector chosen from a basis of eigenvectors for the eigenspace.

Some authors create matrix representations of nilpotent linear transformations where the Jordan block has the ones just below the diagonal (the “subdiagonal”). No matter, it is really the same, just different. We have also defined Jordan canonical form to place blocks for the larger eigenvalues earlier, and for blocks with the same eigenvalue, we place the larger sized blocks earlier. This is fairly standard, but there is no reason we could not order the blocks differently. It would be the same, just different. The reason for choosing *some* ordering is to be assured that there is just *one* canonical matrix representation for each linear transformation.

Example 27. Suppose that $T: \mathbb{C}^{10} \rightarrow \mathbb{C}^{10}$ is the linear transformation defined by $T(\mathbf{x}) = A\mathbf{x}$ where

$$A = \begin{bmatrix} -6 & 9 & -7 & -5 & 5 & 12 & -22 & 14 & 8 & 21 \\ -3 & 5 & -3 & -1 & 2 & 7 & -12 & 9 & 1 & 12 \\ 8 & -9 & 8 & 6 & 0 & -14 & 25 & -13 & -4 & -26 \\ -7 & 9 & -7 & -5 & 0 & 13 & -23 & 13 & 2 & 24 \\ 0 & -1 & 0 & -1 & -3 & -2 & 3 & -4 & -2 & -3 \\ 3 & 2 & 1 & 2 & 9 & -1 & 1 & 5 & 5 & -5 \\ -1 & 3 & -3 & -2 & 4 & 3 & -6 & 4 & 4 & 3 \\ 3 & -4 & 3 & 2 & 1 & -5 & 9 & -5 & 1 & -9 \\ 0 & 2 & 0 & 0 & 2 & 2 & -4 & 4 & 2 & 4 \\ -4 & 4 & -5 & -4 & -1 & 6 & -11 & 4 & 1 & 10 \end{bmatrix}$$

We will find a basis for \mathbb{C}^{10} that will yield a matrix representation of T in Jordan canonical form. First we find the eigenvalues, and their multiplicities, with the techniques of CHAPTER E.

$$\begin{array}{lll} \lambda = 2 & \alpha_T(2) = 2 & \gamma_T(2) = 2 \\ \lambda = 0 & \alpha_T(0) = 3 & \gamma_T(-1) = 2 \\ \lambda = -1 & \alpha_T(-1) = 5 & \gamma_T(-1) = 2 \end{array}$$

For each eigenvalue, we can compute a generalized eigenspace. By Theorem 3.4 we know that \mathbb{C}^{10} will decompose into a direct sum of these invariant eigenspaces, and we can restrict T to each part of this decomposition. At this stage we know that the Jordan canonical form will be block diagonal with blocks of size 2, 3 and 5, since the dimensions of the generalized eigenspaces are equal to the algebraic multiplicities of the eigenvalues (Theorem 3.5). The geometric multiplicities tell us how many Jordan blocks occupy each of the three larger blocks, but we will discuss this as we analyze each eigenvalue. We do not yet know the index of each eigenvalue (though we can easily infer it for $\lambda = 2$) and even if we did have this information, it only determines the size of the largest Jordan block (per eigenvalue). We will press ahead, considering each eigenvalue one at a time.

The eigenvalue $\lambda = 2$ has “full” geometric multiplicity, and is not an impediment to diagonalizing T . We will treat it in full generality anyway. First we compute the generalized eigenspace. Since Theorem 3.3 says that $\mathcal{G}_T(2) = \mathcal{K}\left((T - 2I_{\mathbb{C}^{10}})^{10}\right)$ we

can compute this generalized eigenspace as a null space derived from the matrix A ,

$$(A - 2I_{10})^{10} \xrightarrow{\text{RREF}} \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -2 & -1 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & -1 & -1 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & 2 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & -1 & -2 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & -2 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

$$\mathcal{G}_T(2) = \mathcal{K}\left((A - 2I_{10})^{10}\right) = \left\langle \left\{ \begin{bmatrix} 2 \\ 1 \\ -1 \\ 1 \\ -1 \\ 2 \\ 1 \\ 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \\ -2 \\ 2 \\ 0 \\ -1 \\ 0 \\ -1 \\ 0 \\ 1 \end{bmatrix} \right\} \right\rangle$$

The restriction of T to $\mathcal{G}_T(2)$ relative to the two basis vectors above has a matrix representation that is a 2×2 diagonal matrix with the eigenvalue $\lambda = 2$ as the diagonal entries. So these two vectors will be the first two vectors in our basis for \mathbb{C}^{10} ,

$$\mathbf{v}_1 = \begin{bmatrix} 2 \\ 1 \\ -1 \\ 1 \\ -1 \\ 2 \\ 1 \\ 0 \\ 1 \\ 0 \end{bmatrix} \quad \mathbf{v}_2 = \begin{bmatrix} 1 \\ 1 \\ -2 \\ 2 \\ 0 \\ -1 \\ 0 \\ -1 \\ 0 \\ 1 \end{bmatrix}$$

Notice that it was not strictly necessary to compute the 10-th power of $A - 2I_{10}$. With $\alpha_T(2) = \gamma_T(2)$ the null space of the matrix $A - 2I_{10}$ contains *all* of the generalized eigenvectors of T for the eigenvalue $\lambda = 2$. But there was no harm in computing the 10-th power either. This discussion is equivalent to the observation that the linear transformation $T|_{\mathcal{G}_T(2)}: \mathcal{G}_T(2) \rightarrow \mathcal{G}_T(2)$ is nilpotent of index 1. In other words, $\iota_T(2) = 1$.

The eigenvalue $\lambda = 0$ will not be quite as simple, since the geometric multiplicity is strictly less than the algebraic multiplicity. As before, we first compute the generalized eigenspace. Since Theorem 3.3 says that $\mathcal{G}_T(0) = \mathcal{K}\left((T - 0I_{\mathbb{C}^{10}})^{10}\right)$ we

can compute this generalized eigenspace as a null space derived from the matrix A ,

$$(A - 0I_{10})^{10} \xrightarrow{\text{RREF}} \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 & -1 \\ 0 & 1 & 0 & 0 & 0 & 0 & -1 & 0 & -1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & 2 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & -2 & -1 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & -1 & 0 & -1 & 2 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

$$\mathcal{G}_T(0) = \mathcal{K}\left((A - 0I_{10})^{10}\right) = \left\langle \left\{ \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 1 \\ 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \\ -1 \\ 2 \\ -1 \\ 1 \\ 0 \\ -1 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \\ -2 \\ 1 \\ 0 \\ -2 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} \right\} \right\rangle = \langle F \rangle$$

So $\dim(\mathcal{G}_T(0)) = 3 = \alpha_T(0)$, as expected. We will use these three basis vectors for the generalized eigenspace to construct a matrix representation of $T|_{\mathcal{G}_T(0)}$, where F is being defined implicitly as the basis of $\mathcal{G}_T(0)$. We construct this representation

as usual, applying DEFINITION MR,

$$\begin{aligned}
 \rho_F \left(T|_{\mathcal{G}_T(0)} \left(\begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \\ 1 \\ 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} \right) \right) &= \rho_F \left(\begin{bmatrix} -1 \\ 0 \\ 2 \\ -1 \\ 0 \\ 2 \\ 0 \\ 0 \\ 0 \\ -1 \end{bmatrix} \right) = \rho_F \left((-1) \begin{bmatrix} 1 \\ 0 \\ -2 \\ 1 \\ 0 \\ -2 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} \right) = \begin{bmatrix} 0 \\ 0 \\ -1 \end{bmatrix} \\
 \rho_F \left(T|_{\mathcal{G}_T(0)} \left(\begin{bmatrix} 1 \\ 1 \\ -1 \\ 2 \\ -1 \\ 1 \\ 1 \\ 0 \\ -1 \\ 1 \\ 0 \end{bmatrix} \right) \right) &= \rho_F \left(\begin{bmatrix} 1 \\ 0 \\ -2 \\ 1 \\ 0 \\ -2 \\ 0 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} \right) = \rho_F \left((1) \begin{bmatrix} 1 \\ 0 \\ -2 \\ 1 \\ 0 \\ -2 \\ 0 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} \right) = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \\
 \rho_F \left(T|_{\mathcal{G}_T(0)} \left(\begin{bmatrix} 1 \\ 0 \\ -2 \\ 1 \\ 0 \\ -2 \\ 0 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} \right) \right) &= \rho_F \left(\begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \right) = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}
 \end{aligned}$$

So we have the matrix representation

$$M = M_{F,F}^{T|_{\mathcal{G}_T(0)}} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ -1 & 1 & 0 \end{bmatrix}$$

By Theorem 3.9 we can obtain a nilpotent matrix from this matrix representation by subtracting the eigenvalue from the diagonal elements, and then we can apply Theorem 3.11 to $M - (0)I_3$. First check that $(M - (0)I_3)^2 = \mathcal{O}$, so we know that the index of $M - (0)I_3$ as a nilpotent matrix, and that therefore $\lambda = 0$ is an eigenvalue of T with index 2, $\iota_T(0) = 2$. To determine a basis of \mathbb{C}^3 that converts $M - (0)I_3$ to canonical form, we need the null spaces of the powers of $M - (0)I_3$. For convenience, set $N = M - (0)I_3$.

$$\begin{aligned}
 \mathcal{N}(N^1) &= \left\langle \left\{ \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \right\} \right\rangle \\
 \mathcal{N}(N^2) &= \left\langle \left\{ \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \right\} \right\rangle = \mathbb{C}^3
 \end{aligned}$$

Then we choose a vector from $\mathcal{N}(N^2)$ that is not an element of $\mathcal{N}(N^1)$. Any vector with unequal first two entries will fit the bill, say

$$\mathbf{z}_{2,1} = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}$$

where we are employing the notation in Theorem 3.11. The next step is to multiply this vector by N to get part of the basis for $\mathcal{N}(N^1)$,

$$\mathbf{z}_{1,1} = N\mathbf{z}_{2,1} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ -1 & 1 & 0 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ -1 \end{bmatrix}$$

We need a vector to pair with $\mathbf{z}_{1,1}$ that will make a basis for the two-dimensional subspace $\mathcal{N}(N^1)$. Examining the basis for $\mathcal{N}(N^1)$ we see that a vector with its first two entries equal will do the job.

$$\mathbf{z}_{1,2} = \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}$$

Reordering, we find the basis,

$$C = \{\mathbf{z}_{1,1}, \mathbf{z}_{2,1}, \mathbf{z}_{1,2}\} = \left\{ \begin{bmatrix} 0 \\ 0 \\ -1 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} \right\}$$

From this basis, we can get a matrix representation of N (when viewed as a linear transformation) relative to the basis C for \mathbb{C}^3 ,

$$\begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} = \begin{bmatrix} J_2(0) & \mathcal{O} \\ \mathcal{O} & J_1(0) \end{bmatrix}$$

Now we add back the eigenvalue $\lambda = 0$ to the representation of N to obtain a representation for M . Of course, with an eigenvalue of zero, the change is not apparent, so we will not display the same matrix again. This is the second block of the Jordan canonical form for T . However, the three vectors in C will not suffice as basis vectors for the domain of T —they have the wrong size! The vectors in C are vectors in the domain of a linear transformation defined by the matrix M . But M was a matrix representation of $T|_{\mathcal{G}_T(0)} - 0I_{\mathcal{G}_T(0)}$ relative to the basis F for $\mathcal{G}_T(0)$. We need to “uncoordinatize” each of the basis vectors in C to produce a linear combination of vectors in F that will be an element of the generalized eigenspace $\mathcal{G}_T(0)$. These will be the next three vectors of our final answer, a basis for \mathbb{C}^{10} that

has a pleasing matrix representation.

$$\begin{aligned}
 \mathbf{v}_3 = \rho_F^{-1} \left(\begin{bmatrix} 0 \\ 0 \\ -1 \end{bmatrix} \right) &= 0 \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 1 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} + 0 \begin{bmatrix} 1 \\ 1 \\ -1 \\ 2 \\ -1 \\ 1 \\ 0 \\ -1 \\ 1 \\ 0 \end{bmatrix} + (-1) \begin{bmatrix} 1 \\ 0 \\ -2 \\ 1 \\ 0 \\ -2 \\ 0 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} = \begin{bmatrix} -1 \\ 0 \\ 2 \\ -1 \\ 0 \\ 2 \\ 0 \\ 0 \\ 0 \\ 0 \\ -1 \end{bmatrix} \\
 \mathbf{v}_4 = \rho_F^{-1} \left(\begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} \right) &= 1 \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 1 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} + 0 \begin{bmatrix} 1 \\ 1 \\ -1 \\ 2 \\ -1 \\ 1 \\ 0 \\ -1 \\ 1 \\ 0 \end{bmatrix} + 0 \begin{bmatrix} 1 \\ 0 \\ -2 \\ 1 \\ 0 \\ -2 \\ 0 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} = \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 1 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \\
 \mathbf{v}_5 = \rho_F^{-1} \left(\begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} \right) &= 1 \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 1 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} + 1 \begin{bmatrix} 1 \\ 1 \\ -1 \\ 2 \\ -1 \\ 1 \\ 0 \\ -1 \\ 1 \\ 0 \end{bmatrix} + 0 \begin{bmatrix} 1 \\ 0 \\ -2 \\ 1 \\ 0 \\ -2 \\ 0 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} = \begin{bmatrix} 1 \\ 2 \\ -1 \\ 2 \\ -1 \\ 2 \\ 1 \\ -1 \\ 1 \\ 0 \\ 0 \end{bmatrix}
 \end{aligned}$$

Five down, five to go. Basis vectors, that is. $\lambda = -1$ is the smallest eigenvalue, but it will require the most computation. First we compute the generalized eigenspace. Since Theorem 3.3 says that $\mathcal{G}_T(-1) = \mathcal{K}\left((T - (-1)I_{\mathbb{C}^{10}})^{10}\right)$ we can

compute this generalized eigenspace as a null space derived from the matrix A ,

$$(A - (-1)I_{10})^{10} \xrightarrow{\text{RREF}} \begin{bmatrix} 1 & 0 & 1 & 0 & 1 & 0 & -1 & 1 & 0 & 1 \\ 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 0 & 1 & 0 & 0 & -2 \\ 0 & 0 & 0 & 0 & 0 & 1 & -2 & 1 & 0 & 2 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

$$\mathcal{G}_T(-1) = \mathcal{K}\left((A - (-1)I_{10})^{10}\right) = \left\langle \left\{ \begin{bmatrix} -1 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} -1 \\ -1 \\ 0 \\ -1 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \\ 0 \\ -1 \\ 0 \\ 2 \\ 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} -1 \\ -1 \\ 0 \\ 0 \\ 0 \\ -1 \\ 0 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} -1 \\ 0 \\ 0 \\ 2 \\ 0 \\ -2 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} \right\} \right\rangle = \langle F \rangle$$

So $\dim(\mathcal{G}_T(-1)) = 5 = \alpha_T(-1)$, as expected. We will use these five basis vectors for the generalized eigenspace to construct a matrix representation of $T|_{\mathcal{G}_T(-1)}$, where F is being recycled and defined now implicitly as the basis of $\mathcal{G}_T(-1)$.

We construct this representation as usual, applying DEFINITION MR,

$$\rho_F \left(T|_{\mathcal{G}_T(-1)} \left(\begin{bmatrix} -1 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \right) \right) = \rho_F \left(\begin{bmatrix} -1 \\ 0 \\ 0 \\ 0 \\ 0 \\ -2 \\ -2 \\ 0 \\ 0 \\ -1 \end{bmatrix} \right) \quad (3.1)$$

$$= \rho_F \left(0 \begin{bmatrix} -1 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} + 0 \begin{bmatrix} -1 \\ -1 \\ 0 \\ -1 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} + (-2) \begin{bmatrix} 1 \\ 0 \\ 0 \\ -1 \\ 0 \\ 2 \\ 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} + 0 \begin{bmatrix} -1 \\ -1 \\ 0 \\ 0 \\ 0 \\ -1 \\ 0 \\ 1 \\ 0 \\ 0 \end{bmatrix} + (-1) \begin{bmatrix} -1 \\ 0 \\ 0 \\ 2 \\ 0 \\ -2 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} \right) = \begin{bmatrix} 0 \\ 0 \\ -2 \\ 0 \\ -1 \end{bmatrix}$$

$$\rho_F \left(T|_{\mathcal{G}_T(-1)} \left(\begin{bmatrix} -1 \\ -1 \\ 0 \\ -1 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \right) \right) = \rho_F \left(\begin{bmatrix} 7 \\ 1 \\ -5 \\ 3 \\ -1 \\ 2 \\ 4 \\ 0 \\ 0 \\ 3 \end{bmatrix} \right) \quad (3.2)$$

$$= \rho_F \left((-5) \begin{bmatrix} -1 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} + (-1) \begin{bmatrix} -1 \\ -1 \\ 0 \\ -1 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} + 4 \begin{bmatrix} 1 \\ 0 \\ 0 \\ -1 \\ 0 \\ 2 \\ 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} + 0 \begin{bmatrix} -1 \\ -1 \\ 0 \\ 0 \\ 0 \\ -1 \\ 0 \\ 1 \\ 0 \\ 0 \end{bmatrix} + 3 \begin{bmatrix} -1 \\ 0 \\ 0 \\ 2 \\ 0 \\ -2 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} \right) = \begin{bmatrix} -5 \\ -1 \\ 4 \\ 0 \\ 3 \end{bmatrix} \quad (3.3)$$

$$\rho_F \left(T|_{\mathcal{G}_T(-1)} \left(\begin{bmatrix} 1 \\ 0 \\ 0 \\ -1 \\ 0 \\ 2 \\ 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} \right) \right) = \rho_F \left(\begin{bmatrix} 1 \\ 0 \\ -1 \\ 1 \\ 0 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} \right) \quad (3.4)$$

So we have the matrix representation of the restriction of T (again recycling and redefining the matrix M)

$$M = M_{F,F}^{T|_{g_T(-1)}} = \begin{bmatrix} 0 & -5 & -1 & 2 & 6 \\ 0 & -1 & 0 & -1 & -1 \\ -2 & 4 & 1 & -1 & -6 \\ 0 & 0 & 0 & 1 & 2 \\ -1 & 3 & 1 & -2 & -6 \end{bmatrix}$$

Theorem 3.9 says we can obtain a nilpotent matrix from this matrix representation by subtracting the eigenvalue from the diagonal elements, and then we can apply Theorem 3.11 to $M - (-1)I_5$. First check that $(M - (-1)I_5)^3 = \mathcal{O}$, so we know that the index of $M - (-1)I_5$ as a nilpotent matrix, and that therefore $\lambda = -1$ is an eigenvalue of T with index 3, $\iota_T(-1) = 3$. To determine a basis of \mathbb{C}^5 that converts $M - (-1)I_5$ to canonical form, we need the null spaces of the powers of $M - (-1)I_5$. Again, for convenience, set $N = M - (-1)I_5$.

$$\begin{aligned} \mathcal{N}(N^1) &= \left\langle \left\{ \begin{bmatrix} 1 \\ 0 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} -3 \\ 1 \\ 0 \\ -2 \\ 2 \end{bmatrix} \right\} \right\rangle \\ \mathcal{N}(N^2) &= \left\langle \left\{ \begin{bmatrix} 3 \\ 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} -3 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} \right\} \right\rangle \\ \mathcal{N}(N^3) &= \left\langle \left\{ \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} \right\} \right\rangle = \mathbb{C}^5 \end{aligned}$$

Then we choose a vector from $\mathcal{N}(N^3)$ that is not an element of $\mathcal{N}(N^2)$. The sum of the four basis vectors for $\mathcal{N}(N^2)$ sum to a vector with all five entries equal to 1. We will adjust with the first entry to create a vector not in $\mathcal{N}(N^2)$,

$$\mathbf{z}_{3,1} = \begin{bmatrix} 0 \\ 1 \\ 1 \\ 1 \\ 1 \end{bmatrix}$$

where we are employing the notation in Theorem 3.11. The next step is to multiply this vector by N to get a portion of the basis for $\mathcal{N}(N^2)$,

$$\mathbf{z}_{2,1} = N\mathbf{z}_{3,1} = \begin{bmatrix} 1 & -5 & -1 & 2 & 6 \\ 0 & 0 & 0 & -1 & -1 \\ -2 & 4 & 2 & -1 & -6 \\ 0 & 0 & 0 & 2 & 2 \\ -1 & 3 & 1 & -2 & -5 \end{bmatrix} \begin{bmatrix} 0 \\ 1 \\ 1 \\ 1 \\ 1 \end{bmatrix} = \begin{bmatrix} 2 \\ -2 \\ -1 \\ 4 \\ -3 \end{bmatrix}$$

We have a basis for the two-dimensional subspace $\mathcal{N}(N^1)$ and we can add to that the vector $\mathbf{z}_{2,1}$ and we have three of four basis vectors for $\mathcal{N}(N^2)$. These three

vectors span the subspace we call Q_2 . We need a fourth vector outside of Q_2 to complete a basis of the four-dimensional subspace $\mathcal{N}(N^2)$. Check that the vector

$$\mathbf{z}_{2,2} = \begin{bmatrix} 3 \\ 1 \\ 3 \\ 1 \\ 1 \end{bmatrix}$$

is an element of $\mathcal{N}(N^2)$ that lies outside of the subspace Q_2 . This vector was constructed by getting a nice basis for Q_2 and forming a linear combination of this basis that specifies three of the five entries of the result. Of the remaining two entries, one was changed to move the vector outside of Q_2 and this was followed by a change to the remaining entry to place the vector into $\mathcal{N}(N^2)$. The vector $\mathbf{z}_{2,2}$ is the lone basis vector for the subspace we call R_2 .

The remaining two basis vectors are easy to come by. They are the result of applying N to each of the two most recently determined basis vectors,

$$\mathbf{z}_{1,1} = N\mathbf{z}_{2,1} = \begin{bmatrix} 3 \\ -1 \\ 0 \\ 2 \\ -2 \end{bmatrix} \quad \mathbf{z}_{1,2} = N\mathbf{z}_{2,2} = \begin{bmatrix} 3 \\ -2 \\ -3 \\ 4 \\ -4 \end{bmatrix}$$

Now we reorder these basis vectors, to arrive at the basis

$$C = \{\mathbf{z}_{1,1}, \mathbf{z}_{2,1}, \mathbf{z}_{3,1}, \mathbf{z}_{1,2}, \mathbf{z}_{2,2}\} = \left\{ \begin{bmatrix} 3 \\ -1 \\ 0 \\ 2 \\ -2 \end{bmatrix}, \begin{bmatrix} 2 \\ -2 \\ -1 \\ 4 \\ -3 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 1 \\ 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 3 \\ -2 \\ -3 \\ 4 \\ -4 \end{bmatrix}, \begin{bmatrix} 3 \\ 1 \\ 3 \\ 1 \\ 1 \end{bmatrix} \right\}$$

A matrix representation of N relative to C is

$$\begin{bmatrix} 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} = \begin{bmatrix} J_3(0) & \mathcal{O} \\ \mathcal{O} & J_2(0) \end{bmatrix}$$

To obtain a matrix representation of M , we add back in the matrix $(-1)I_5$, placing the eigenvalue back along the diagonal, and slightly modifying the Jordan blocks,

$$\begin{bmatrix} -1 & 1 & 0 & 0 & 0 \\ 0 & -1 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 & 0 \\ 0 & 0 & 0 & -1 & 1 \\ 0 & 0 & 0 & 0 & -1 \end{bmatrix} = \begin{bmatrix} J_3(-1) & \mathcal{O} \\ \mathcal{O} & J_2(-1) \end{bmatrix}$$

The basis C yields a pleasant matrix representation for the *restriction* of the linear transformation $T - (-1)I$ to the generalized eigenspace $\mathcal{G}_T(-1)$. However, we must remember that these vectors in \mathbb{C}^5 are representations of vectors in \mathbb{C}^{10} relative to the basis F . Each needs to be “un-coordinatized” before joining our final basis.

Here we go,

$$\begin{aligned}
 \mathbf{v}_6 &= \rho_F^{-1} \begin{pmatrix} 3 \\ -1 \\ 0 \\ 2 \\ -2 \end{pmatrix} = 3 \begin{bmatrix} -1 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} + (-1) \begin{bmatrix} -1 \\ -1 \\ 0 \\ -1 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} + 0 \begin{bmatrix} 1 \\ 0 \\ 0 \\ -1 \\ 0 \\ 2 \\ 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} + 2 \begin{bmatrix} -1 \\ -1 \\ 0 \\ 0 \\ -1 \\ 0 \\ 0 \\ 1 \\ 0 \\ 0 \end{bmatrix} + (-2) \begin{bmatrix} -1 \\ 0 \\ 0 \\ 2 \\ 0 \\ -2 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} = \begin{bmatrix} -2 \\ -1 \\ 3 \\ -3 \\ -1 \\ 2 \\ 0 \\ 2 \\ 0 \\ -2 \end{bmatrix} \\
 \mathbf{v}_7 &= \rho_F^{-1} \begin{pmatrix} 2 \\ -2 \\ -1 \\ 4 \\ -3 \end{pmatrix} = 2 \begin{bmatrix} -1 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} + (-2) \begin{bmatrix} -1 \\ -1 \\ 0 \\ -1 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} + (-1) \begin{bmatrix} 1 \\ 0 \\ 0 \\ -1 \\ 0 \\ 2 \\ 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} + 4 \begin{bmatrix} -1 \\ -1 \\ 0 \\ 0 \\ 0 \\ -1 \\ 0 \\ 1 \\ 0 \\ 0 \end{bmatrix} + (-3) \begin{bmatrix} -1 \\ 0 \\ 0 \\ 2 \\ 0 \\ -2 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} = \begin{bmatrix} -2 \\ -2 \\ 2 \\ -3 \\ -2 \\ 0 \\ -1 \\ 4 \\ 0 \\ -3 \end{bmatrix} \\
 \mathbf{v}_8 &= \rho_F^{-1} \begin{pmatrix} 0 \\ 1 \\ 1 \\ 1 \\ 1 \end{pmatrix} = 0 \begin{bmatrix} -1 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} + 1 \begin{bmatrix} -1 \\ -1 \\ 0 \\ -1 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} + 1 \begin{bmatrix} 1 \\ 0 \\ 0 \\ -1 \\ 0 \\ 2 \\ 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} + 1 \begin{bmatrix} -1 \\ -1 \\ 0 \\ 0 \\ 0 \\ -1 \\ 0 \\ 1 \\ 0 \\ 0 \end{bmatrix} + 1 \begin{bmatrix} -1 \\ 0 \\ 0 \\ 2 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} = \begin{bmatrix} -2 \\ -2 \\ 0 \\ 0 \\ 1 \\ -1 \\ 1 \\ 1 \\ 0 \\ 1 \end{bmatrix} \\
 \mathbf{v}_9 &= \rho_F^{-1} \begin{pmatrix} 3 \\ -2 \\ -3 \\ 4 \\ -4 \end{pmatrix} = 3 \begin{bmatrix} -1 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} + (-2) \begin{bmatrix} -1 \\ -1 \\ 0 \\ -1 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} + (-3) \begin{bmatrix} 1 \\ 0 \\ 0 \\ -1 \\ 0 \\ 2 \\ 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} + 4 \begin{bmatrix} -1 \\ -1 \\ 0 \\ 0 \\ 0 \\ -1 \\ 0 \\ 1 \\ 0 \\ 0 \end{bmatrix} + (-4) \begin{bmatrix} -1 \\ 0 \\ 0 \\ 2 \\ 0 \\ -2 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} = \begin{bmatrix} -4 \\ -2 \\ 3 \\ -3 \\ -2 \\ -2 \\ -3 \\ 4 \\ 0 \\ -4 \end{bmatrix} \\
 \mathbf{v}_{10} &= \rho_F^{-1} \begin{pmatrix} 3 \\ 1 \\ 3 \\ 1 \\ 1 \end{pmatrix} = 3 \begin{bmatrix} -1 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} + 1 \begin{bmatrix} -1 \\ -1 \\ 0 \\ -1 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} + 3 \begin{bmatrix} 1 \\ 0 \\ 0 \\ -1 \\ 0 \\ 2 \\ 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} + 1 \begin{bmatrix} -1 \\ -1 \\ 0 \\ 0 \\ 0 \\ -1 \\ 0 \\ 1 \\ 0 \\ 0 \end{bmatrix} + 1 \begin{bmatrix} -1 \\ 0 \\ 0 \\ 2 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} = \begin{bmatrix} -3 \\ -2 \\ 3 \\ -2 \\ 1 \\ 3 \\ 3 \\ 1 \\ 0 \\ 1 \end{bmatrix}
 \end{aligned}$$

To summarize, we list the entire basis $B = \{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3, \dots, \mathbf{v}_{10}\}$,

$$\begin{array}{ccccc}
 \mathbf{v}_1 & \begin{bmatrix} 2 \\ 1 \\ -1 \\ 1 \\ -1 \\ 2 \\ 1 \\ 0 \\ 1 \\ 0 \end{bmatrix} & \mathbf{v}_2 & \begin{bmatrix} 1 \\ 1 \\ -2 \\ 2 \\ 0 \\ -1 \\ 0 \\ -1 \\ 0 \\ 1 \end{bmatrix} & \mathbf{v}_3 & \begin{bmatrix} -1 \\ 0 \\ 2 \\ -1 \\ 0 \\ 2 \\ 0 \\ 0 \\ 0 \\ -1 \end{bmatrix} & \mathbf{v}_4 & \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \\ 1 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} & \mathbf{v}_5 & \begin{bmatrix} 1 \\ 2 \\ -1 \\ 2 \\ -1 \\ 2 \\ 1 \\ -1 \\ 1 \\ 0 \end{bmatrix} \\
 \mathbf{v}_6 & \begin{bmatrix} -2 \\ -1 \\ 3 \\ -3 \\ -1 \\ 2 \\ 0 \\ 2 \\ 0 \\ -2 \end{bmatrix} & \mathbf{v}_7 & \begin{bmatrix} -2 \\ -2 \\ 2 \\ -3 \\ -2 \\ 0 \\ -1 \\ 4 \\ 0 \\ -3 \end{bmatrix} & \mathbf{v}_8 & \begin{bmatrix} -2 \\ -2 \\ 0 \\ 0 \\ 1 \\ -1 \\ 1 \\ 1 \\ 0 \\ 1 \end{bmatrix} & \mathbf{v}_9 & \begin{bmatrix} -4 \\ -2 \\ 3 \\ -3 \\ -2 \\ -2 \\ -3 \\ 4 \\ 0 \\ -4 \end{bmatrix} & \mathbf{v}_{10} & \begin{bmatrix} -3 \\ -2 \\ 3 \\ -2 \\ 1 \\ 3 \\ 3 \\ 1 \\ 0 \\ 1 \end{bmatrix}
 \end{array}$$

The resulting matrix representation is

$$M_{B,B}^T = \begin{bmatrix} 2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & -1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & -1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 \end{bmatrix}$$

If you are not inclined to check all of these computations, here are a few that should convince you of the amazing properties of the basis B . Compute the matrix-vector products $A\mathbf{v}_i$, $1 \leq i \leq 10$. In each case the result will be a vector of the form $\lambda\mathbf{v}_i + \delta\mathbf{v}_{i-1}$, where λ is one of the eigenvalues (you should be able to predict ahead of time *which* one) and $\delta \in \{0, 1\}$.

Alternatively, if we can write inputs to the linear transformation T as linear combinations of the vectors in B , then the “action” of T is reduced to a matrix-vector product with the exceedingly simple matrix that is the Jordan canonical form. Wow!

Chapter 4

Applications

4.1 Least Squares

Solving a linear system of equations is a fundamental use of the tools of linear algebra. You know from introductory linear algebra that a linear system may have no solution. In an applied situation there could be many reasons for this, and it begs the question: what to do next?

We often construct mathematical models of practical situations, frequently in an effort to measure various parameters of the model. Suppose we think that interest rates R , as measured by the rate on one-year government bonds, are a linear function of construction activity C , as measured by the number of permits issued for new construction in the last 30 days. So the “hotter” the construction market, the greater the demand for loans, so the cost of money (the interest rate) is greater. With a good model, we might be able to predict interest rates by examining public records for changes in the number of construction permits issued.

So we have a mathematical model

$$I = aR + b$$

where we do not know a or b , the parameters of the model. But we would like to *know* the values of these parameters. Or at least have good *estimates*, since we understand that our model is an extremely simple representation of a much more complicated situation. So we collect data by obtaining monthly records of the interest rate and construction permits for the past sixty years. Now we have 720 pairs of permits issued and interest rates. We can substitute each pair into our model and we get a linear equation in the parameters a and b . Two such equations would be likely to have a unique solution, however if we consider all 720 equations there is unlikely to be a solution. Why do we say that? Imagine the linear system we would obtain from all these equations. The 720×2 coefficient matrix will have the values of C in the first column, and the second column will be all 1's. The vector of constants will be the values of R . If you form the augmented matrix of this system and row-reduce, you will be virtually guaranteed to get a pivot column in the third column and thus have no solution (THEOREM RCLS).

The lack of a solution for our model can be expressed more clearly by saying that vector of constants (the values of R) is not in the column space of the coefficient matrix. Or the vector of C values is not a linear combination of the vector of R values and the vector of all ones. If it were, then the scalars in the linear combination would be our solution for a and b . We will temporarily leave our economic forecasting example to be more general, but we will return to the example soon.

Consider the general problem of solving $A\mathbf{x} = \mathbf{b}$ when there is no solution. We *desire* an \mathbf{x} so that $A\mathbf{x}$ equals \mathbf{b} , but we will *settle* for an \mathbf{x} so that $A\mathbf{x}$ is *very close* to \mathbf{b} . What do we mean by close? One measure would be that we want the difference $A\mathbf{x} - \mathbf{b}$ to be a short vector, as short as possible. In other words we want to minimize the norm, $\|A\mathbf{x} - \mathbf{b}\|$. We will return to this minimization idea shortly, but let's ask informally what *direction* this short vector might have?

We have a subspace, the column space of A , represented by all possibilities for $A\mathbf{x}$ as we course over all values of \mathbf{x} . We have a vector \mathbf{b} that is not in the subspace. Which element in the subspace is closest to \mathbf{b} ? Consider the situation in two dimensions, where the subspace is a line through the origin, and the point lies off the line. From the point off the line, which of the infinitely many points on the line is closest, and what direction would you take to get there? In three dimensions, consider a subspace that is a plane containing the origin, and the point lies off the plane. From the point off the plane, which of the infinitely many points on the plane is closest, and what direction would you take to get there? You should have answered “in a perpendicular direction” and “in the direction of the normal vector to the plane”. Move orthogonally to the subspace. More exactly, in the direction of an element of the orthogonal complement of the subspace.

More carefully, $A\mathbf{x} - \mathbf{b} \in (\mathcal{C}(A))^\perp = \mathcal{N}(A^*)$, so

$$\mathbf{0} = A^*(A\mathbf{x} - \mathbf{b}) = A^*A\mathbf{x} - A^*\mathbf{b} \quad \Rightarrow \quad A^*A\mathbf{x} = A^*\mathbf{b}.$$

This linear system is called the **normal equations**, due to the role of orthogonality in its derivation. Several good things have happened. Primarily, A^*A is a square matrix, is positive semi-definite, and if A has full column rank, then A^*A is nonsingular. With a nonsingular coefficient matrix, we have the unique solution

$$\mathbf{x} = (A^*A)^{-1} A^*\mathbf{b}.$$

We consider two very simple examples of using the normal equations, in cases where our geometric intuition is useful.

Example 28. Consider the extremely trivial system of equations in one variable, $A\mathbf{x} = \mathbf{b}$ with

$$A = \begin{bmatrix} 1 \\ 2 \end{bmatrix} \quad \mathbf{x} = [x_1] \quad \mathbf{b} = \begin{bmatrix} 5 \\ 15 \end{bmatrix}$$

Quite clearly this system has no solution. Forming the normal equations we get the silly system

$$[5] [x_1] = [35]$$

which has the unique solution $x_1 = 7$. What is more interesting is that

$$A [7] = \begin{bmatrix} 7 \\ 14 \end{bmatrix}.$$

This is a vector in $\mathcal{C}(A)$ and the difference

$$\mathbf{r} = \begin{bmatrix} 5 \\ 15 \end{bmatrix} - \begin{bmatrix} 7 \\ 14 \end{bmatrix} = \begin{bmatrix} -2 \\ 1 \end{bmatrix}$$

is orthogonal to the columns of A , in other words \mathbf{r} is in the orthogonal complement of the column space of A . Geometrically, the point $(5, 15)$ is not on the line $y = 2x$, but $(7, 14)$ is. The line has slope 2, while the line segment joining $(5, 15)$ to $(7, 14)$ has slope $-\frac{1}{2}$, making it perpendicular to the line. So $(7, 14)$ is the point on the line closest to $(5, 15)$ and leads to the solution $x_1 = 7$ of the resulting system.

Let's do it again, staying simple, but not trivially so.

Example 29. Consider the simple system of equations in two variables, $A\mathbf{x} = \mathbf{b}$ with

$$A = \begin{bmatrix} 1 & -1 \\ 2 & 3 \\ 3 & 0 \end{bmatrix} \quad \mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \quad \mathbf{b} = \begin{bmatrix} 1 \\ 5 \\ -2 \end{bmatrix}$$

If you try to solve this system, say by row-reducing an augmented matrix, you will discover that there is no solution. The vector \mathbf{b} is not in the column space of A . The normal equations give the system

$$\begin{bmatrix} 14 & 5 \\ 5 & 10 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} 5 \\ 14 \end{bmatrix}$$

which has the unique solution $x_1 = -\frac{4}{23}$, $x_2 = \frac{171}{115}$. What is more interesting is that

$$A \begin{bmatrix} -\frac{4}{23} \\ \frac{171}{115} \end{bmatrix} = \begin{bmatrix} -\frac{191}{115} \\ \frac{473}{115} \\ -\frac{12}{23} \end{bmatrix}.$$

This is a vector in $\mathcal{C}(A)$ and the difference

$$\mathbf{r} = \begin{bmatrix} 1 \\ 5 \\ -2 \end{bmatrix} - \begin{bmatrix} -\frac{191}{115} \\ \frac{473}{115} \\ -\frac{12}{23} \end{bmatrix} = \begin{bmatrix} \frac{306}{115} \\ \frac{102}{115} \\ -\frac{34}{23} \end{bmatrix}$$

is orthogonal to both of the columns of A , in other words \mathbf{r} is in the orthogonal complement of the column space of A . Geometrically, the point $(1, 5, -2)$ is not on the plane spanned by the columns of A , which has equation $9x + 3y - 5z = 0$, but $(-\frac{191}{115}, \frac{473}{115}, -\frac{12}{23})$ is on the plane. The plane has a normal vector $9\vec{i} + 3\vec{j} - 5\vec{k}$, while the vector joining $(1, 5, -2)$ to $(-\frac{191}{115}, \frac{473}{115}, -\frac{12}{23})$ is $\frac{306}{115}\vec{i} + \frac{102}{115}\vec{j} - \frac{34}{23}\vec{k}$, which is a scalar multiple of the normal vector to the plane. So $(-\frac{191}{115}, \frac{473}{115}, -\frac{12}{23})$ is the point on the plane closest to $(1, 5, -2)$ and leads to the solution $x_1 = -\frac{4}{23}$, $x_2 = \frac{171}{115}$ of the resulting system.

What does a solution to the normal equations look like for our economic forecasting model? As discussed above A is a 720×2 matrix, where the first column has the numbers of construction permits, c_i for $1 \leq i \leq 720$, and the second column is all ones. The vector \mathbf{b} contains the 720 values of the interest rate, r_i , $1 \leq i \leq 720$. So

$$A^*A = \begin{bmatrix} \sum c_i^2 & \sum c_i \\ \sum c_i & 720 \end{bmatrix}$$

$$A^*\mathbf{b} = \begin{bmatrix} \sum c_i r_i \\ \sum r_i \end{bmatrix}$$

Then,

$$\begin{aligned} \begin{bmatrix} a \\ b \end{bmatrix} &= \mathbf{x} = (A^*A)^{-1} A^*\mathbf{b} \\ &= \begin{bmatrix} \sum c_i^2 & \sum c_i \\ \sum c_i & 720 \end{bmatrix}^{-1} \begin{bmatrix} \sum c_i r_i \\ \sum r_i \end{bmatrix} \\ &= \frac{1}{720 \sum c_i^2 - (\sum c_i)^2} \begin{bmatrix} 720 & -\sum c_i \\ -\sum c_i & \sum c_i^2 \end{bmatrix} \begin{bmatrix} \sum c_i r_i \\ \sum r_i \end{bmatrix} \\ &= \frac{1}{720 \sum c_i^2 - (\sum c_i)^2} \begin{bmatrix} 720 \sum c_i r_i - \sum c_i \sum r_i \\ -\sum c_i \sum c_i r_i + \sum c_i^2 \sum r_i \end{bmatrix} \end{aligned}$$

The expressions above can be cleaned up some, but the point is to see that we have an expression for each of a and b , that depends solely on the 720 pairs of data points (c_i, r_i) , $1 \leq i \leq 720$. These expressions may look familiar as the most basic case of **linear regression**. Exercise 21 asks you to derive these expressions, and in the more typical order. With estimates in hand, you can now consult the number of construction permits issued and form a prediction of interest rates.

Exercise 21. Suppose we have n pairs of a dependent variable y that varies linearly according to the independent variable x , (x_i, y_i) , $1 \leq i \leq n$. Model the data by the linear equation $y = a + bx$. Solve the normal equations to find expressions that estimate the parameters a and b .

Exercise 22. Find the least-squares estimate obtained from data modeled by the linear equation $y = bx$ (a situation where we know there is no y -intercept).

Exercise 23. Suppose you have data modeled by a single quantity, z , that depends on two independent variables, x and y , linearly according to the model $z = a + bx + cy$. So your data points are triples (x_i, y_i, z_i) , $1 \leq i \leq n$. Can you solve the normal equations to obtain expressions estimating a, b, c ? This might not be a very instructive exercise, but perhaps determine A^*A and $A^*\mathbf{b}$ before letting the matrix inverse dissuade you. What does the geometric picture of the data and your resulting estimates look like?

So far we have not stated any theorems, to say nothing of proving anything. Moving in an orthogonal direction feels like a good idea, but is it really best? Here is a theorem that suggests it is best according to one natural criteria.

Theorem 4.1. *Suppose that A is an $m \times n$ matrix and $\mathbf{b} \in \mathbb{C}^m$. Then $r(\mathbf{x}) = \|\mathbf{Ax} - \mathbf{b}\|$ is minimized by a solution, $\hat{\mathbf{x}}$, to the normal equations, $A^*\mathbf{Ax} = A^*\mathbf{b}$.*

Proof. For any vector \mathbf{x} we can write

$$\begin{aligned} \mathbf{Ax} - \mathbf{b} &= \mathbf{Ax} - A\hat{\mathbf{x}} + A\hat{\mathbf{x}} - \mathbf{b} \\ &= A(\mathbf{x} - \hat{\mathbf{x}}) + (A\hat{\mathbf{x}} - \mathbf{b}) \end{aligned}$$

Clearly the first term is a vector in the column space of A . The second vector, by virtue of $\hat{\mathbf{x}}$ being a solution to the normal equations, is an element of $\mathcal{N}(A^*)$, the orthogonal complement of the column space (Theorem ((result on orthogonal complement of column space))). So this is the promised decomposition of $\mathbf{Ax} - \mathbf{b}$ into the element of a subspace (the column space of A here) and the orthogonal complement of the subspace. Since these two vectors are orthogonal, we can apply the generalized version of the Pythagorean Theorem ((Pythagorean theorem, perhaps into FCLA)))

$$\begin{aligned} \|\mathbf{Ax} - \mathbf{b}\|^2 &= \|A(\mathbf{x} - \hat{\mathbf{x}}) + (A\hat{\mathbf{x}} - \mathbf{b})\|^2 \\ &= \|A(\mathbf{x} - \hat{\mathbf{x}})\|^2 + \|(A\hat{\mathbf{x}} - \mathbf{b})\|^2 \\ &\geq \|(A\hat{\mathbf{x}} - \mathbf{b})\|^2 \end{aligned}$$

This inequality establishes that $\hat{\mathbf{x}}$ minimizes $r(\mathbf{x})$. □

There are other measures of distance than the norm (we are using what would be called the **2-norm**, in order to differentiate it from other norms). But “Euclidean distance”, as we have used here, is often the most tractable and the results are often the most useful. In statistics, it is assumed that deviations in data, deviations from what the model would predict, are due to quantities that vary according

to probability distributions. With this assumption, estimators, such as our least-squares estimates, become complicated functions of these random quantities. With the 2-norm, and reasonable assumptions about the probability distributions for the data, the resulting least-squares estimators have probability distributions that are well-understood, thus making it possible to understand their properties and behavior.

4.1.1 Computing Least-Squares Solutions

In Exercise 23 we suggest formulating general expressions for least-squares estimates of three parameters of a model. This begins to get a bit absurd when you invert a matrix larger than size 2. Instead it makes sense to formulate from the data the matrix A^*A and the vector $A^*\mathbf{b}$, and then proceed to a numerical solution. As A^*A is a Hermitian positive definite matrix ($\langle\langle$ theorem/lemma on positive definite $\rangle\rangle$), it can be decomposed into a Cholesky factorization twice as fast as we can decompose an arbitrary matrix into an LU factorization ($\langle\langle$ cost to form Cholesky $\rangle\rangle$). The Cholesky factorization allows a round of forward-solving followed by a round of back-solving to find a solution to the normal equations.

But it gets better. Suppose A has rank n and we have a thin QR decomposition of A , $A = QR$ where Q has orthogonal columns and R is upper triangular with positive diagonal entries. Notice in particular that R is invertible. Then

$$\begin{aligned} Q^*\mathbf{b} &= (R^*)^{-1} R^* Q^*\mathbf{b} \\ &= (R^*)^{-1} A^*\mathbf{b} \\ &= (R^*)^{-1} A^*A\mathbf{x} \\ &= (R^*)^{-1} R^* Q^* A\mathbf{x} \\ &= Q^* Q R\mathbf{x} \\ &= R\mathbf{x} \end{aligned}$$

This system has an upper triangular coefficient matrix, so one round of back-solving will provide a solution for \mathbf{x} . And the product $Q^*\mathbf{b}$ will be reasonably well-behaved due to the orthogonal columns of Q . Notice that in the case of a square matrix A , the matrix Q will be invertible, and the same system that in general leads to the least-squares solution will also provide the exact solution in this special case.

Exercise 24. Compute the least-squares solution to the system $A\mathbf{x} = \mathbf{b}$. Compute the residual vector associated with your solution.

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

Solution. We compute the pieces of the normal equations

$$A^*A = \begin{bmatrix} 5 & 7 \\ 7 & 10 \end{bmatrix} \qquad A^*\mathbf{b} = \begin{bmatrix} 10 \\ 14 \end{bmatrix}$$

So the solution is

$$\mathbf{x} = \begin{bmatrix} 5 & 7 \\ 7 & 10 \end{bmatrix}^{-1} \begin{bmatrix} 10 \\ 14 \end{bmatrix} = \begin{bmatrix} 10 & -7 \\ -7 & 5 \end{bmatrix} \begin{bmatrix} 10 \\ 14 \end{bmatrix} = \begin{bmatrix} 2 \\ 0 \end{bmatrix}$$

The residual is the difference

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

Notice that residual is indeed orthogonal to the column space of A .

4.2 Curve Fitting

4.2.1 Interpolating Polynomials

Given two points in the plane, there is a unique line through them. Given three points in the plane, and not in a line, there is a unique parabola through them. Given four points in the plane, there is a unique polynomial, of degree 3 or less, passing through them. And so on. We can prove this result, and give a procedure for finding the polynomial with the help of Vandermonde matrices ($\langle\langle$ section on vandermonde matrices $\rangle\rangle$).

Theorem 4.2. *Suppose $\{(x_i, y_i) \mid 1 \leq i \leq n+1\}$ is a set of $n+1$ points in the plane where the x -coordinates are all different. Then there is a unique polynomial of degree n or less, $p(x)$, such that $p(x_i) = y_i$, $1 \leq i \leq n+1$.*

Proof. Write $p(x) = a_0 + a_1x + a_2x^2 + \cdots + a_nx^n$. To meet the conclusion of the theorem, we desire,

$$y_i = p(x_i) = a_0 + a_1x_i + a_2x_i^2 + \cdots + a_nx_i^n \quad 1 \leq i \leq n+1$$

This is a system of $n+1$ linear equations in the $n+1$ variables $a_0, a_1, a_2, \dots, a_n$. The vector of constants in this system is the vector containing the y -coordinates of the points. More importantly, the coefficient matrix is a Vandermonde matrix ($\langle\langle$ definition-vandermonde $\rangle\rangle$) built from the x -coordinates $x_1, x_2, x_3, \dots, x_{n+1}$. Since we have required that these scalars all be different, $\langle\langle$ theorem-NVM $\rangle\rangle$ tells us that the coefficient matrix is nonsingular and THEOREM NMUS says the solution for the coefficients of the polynomial exists, and the solution is unique. As a practical matter, THEOREM SNCM provides an expression for the solution. \square

Example 30. Suppose we have the following 5 points in the plane and we wish to pass a degree 4 polynomial through them.

i	1	2	3	4	5
x_i	-3	-1	2	3	6
y_i	276	16	31	144	2319

Table 4.1: Points on a polynomial

The required system of equations has a coefficient matrix that is the Vandermonde matrix where row i is successive powers of x_i

$$A = \begin{bmatrix} 1 & -3 & 9 & -27 & 81 \\ 1 & -1 & 1 & -1 & 1 \\ 1 & 2 & 4 & 8 & 16 \\ 1 & 3 & 9 & 27 & 81 \\ 1 & 6 & 36 & 216 & 1296 \end{bmatrix}$$

THEOREM NMUS provides a solution as

$$\begin{aligned} \begin{bmatrix} a_0 \\ a_1 \\ a_2 \\ a_3 \\ a_4 \end{bmatrix} &= A^{-1} \begin{bmatrix} 276 \\ 16 \\ 31 \\ 144 \\ 2319 \end{bmatrix} \\ &= \begin{bmatrix} -\frac{1}{15} & \frac{9}{143} & \frac{9}{13} & -\frac{1}{2} & \frac{1}{42} \\ 0 & -\frac{3}{7} & \frac{4}{1} & -\frac{1}{3} & \frac{84}{11} \\ \frac{5}{108} & -\frac{1}{56} & -\frac{4}{1} & \frac{17}{72} & -\frac{1}{756} \\ -\frac{1}{54} & \frac{1}{21} & -\frac{1}{12} & \frac{18}{1} & -\frac{1}{756} \\ \frac{1}{540} & -\frac{1}{168} & \frac{1}{60} & -\frac{1}{72} & \frac{1}{756} \end{bmatrix} \begin{bmatrix} 276 \\ 16 \\ 31 \\ 144 \\ 2319 \end{bmatrix} = \begin{bmatrix} 3 \\ -4 \\ 5 \\ -2 \\ 2 \end{bmatrix} \end{aligned}$$

So the polynomial is $p(x) = 3 - 4x + 5x^2 - 2x^3 + 2x^4$.

The unique polynomial passing through a set of points is known as the **interpolating polynomial** and it has many uses. Unfortunately, when confronted with data from an experiment the situation may not be so simple or clear cut. Read on.

4.2.2 Fitting Curves Through Data

To construct an interpolating polynomial, we presumed we wanted a curve passing through a set of points *exactly*. Sometimes we have a similar, but distinctly different situation. We have a set of data points x_i , $1 \leq i \leq n$, where the x_i are m -tuples. We have a model or a physical law which suggests that each m -tuple satisfies some linear equation with k unknown parameters. We wish to estimate the parameters. If we can formulate a linear system with the parameters as the variables, then we can use a least-squares estimate (Section 4.1). We illustrate with two examples.

Example 31. Suppose we believe the twelve data points below are related by a degree three polynomial, $y = p(x) = a_0 + a_1x + a_2x^2 + a_3x^3$. We have four unknown parameters, the coefficients of the polynomial. For each point we can create a 5 tuple, $(1, x_i, x_i^2, x_i^3, y_i)$, with entries that are related by the linear equation $a_0 + a_1x_i + a_2x_i^2 + a_3x_i^3 = y_i$. So we have 12 linear equations in 4 variables. The coefficient matrix A has 12 rows and 4 columns, similar in spirit to a Vandemonde matrix (Section ??), though not even square. The vector of constants is the 12 values of y_i .

Here are the relevant pieces of the system, the normal equations, and the solution.

$$A = \begin{bmatrix} 1 & 0.142 & 0.020 & 0.003 \\ 1 & 0.646 & 0.417 & 0.269 \\ 1 & 0.954 & 0.909 & 0.867 \\ 1 & 2.958 & 8.751 & 25.886 \\ 1 & 2.975 & 8.851 & 26.332 \\ 1 & 3.167 & 10.032 & 31.775 \\ 1 & 3.413 & 11.649 & 39.757 \\ 1 & 4.302 & 18.504 & 79.595 \\ 1 & 5.552 & 30.830 & 171.180 \\ 1 & 6.576 & 43.247 & 284.403 \\ 1 & 7.958 & 63.325 & 503.917 \\ 1 & 8.028 & 64.444 & 517.341 \end{bmatrix} \quad \mathbf{b} = \begin{bmatrix} -10.867 \\ 10.120 \\ 8.172 \\ 11.693 \\ 18.931 \\ 16.215 \\ 3.863 \\ -7.971 \\ -24.108 \\ -31.217 \\ 0.719 \\ 9.550 \end{bmatrix}$$

x_i	y_i
0.142	-10.867
0.645	10.120
0.953	8.1728
2.958	11.693
2.975	18.931
3.167	16.215
3.413	3.863
4.301	-7.971
5.552	-24.108
6.576	-31.217
7.957	0.719
8.027	9.550

Table 4.2: Points on a polynomial

$$A^*A = \begin{bmatrix} 12.000 & 46.671 & 260.978 & 1681.324 \\ 46.671 & 260.978 & 1681.324 & 11718.472 \\ 260.978 & 1681.324 & 11718.472 & 85542.108 \\ 1681.324 & 11718.472 & 85542.108 & 642050.755 \end{bmatrix}$$

$$A^*\mathbf{b} = \begin{bmatrix} 5.102 \\ -122.81 \\ -1090.783 \\ -6856.475 \end{bmatrix} \quad \mathbf{x} = \begin{bmatrix} -17.726 \\ 47.157 \\ -16.122 \\ 1.323 \end{bmatrix}$$

So the polynomial obtained from a least-squares fit is

$$\hat{p}(x) = 1.323x^3 - 16.122x^2 + 47.157x - 17.726$$

With other models, it may be necessary to rearrange the equation to “linearize” it. For example, if the relationship between x and y is exponential and is given by $y = ae^{bx}$ then applying the logarithm to both sides would yield $\log(y) = \log(a) + bx$. Then by using pairs $(x_i, \log(y_i))$, a least-squares solution would provide estimates of $\log(a)$ and b , which could be easily converted to estimates of a and b .

Chapter 5

Topics

In this chapter we collect a variety of useful topics that are independent of much of the other material, though not exclusively so.

5.1 Vandermonde Matrices

Alexandre-Théophile Vandermonde was a French mathematician in the 1700's who was among the first to write about basic properties of the determinant (such as the effect of swapping two rows). However, the determinant that bears his name (??) does not appear in any of his four published mathematical papers.

Definition 19. A square matrix of size n , A , is a **Vandermonde matrix** if there are scalars, $x_1, x_2, x_3, \dots, x_n$ such that $[A]_{ij} = x_i^{j-1}$, $1 \leq i \leq n$, $1 \leq j \leq n$.

Example 32. The matrix

$$A = \begin{bmatrix} 1 & 2 & 4 & 8 \\ 1 & -3 & 9 & -27 \\ 1 & 1 & 1 & 1 \\ 1 & 4 & 16 & 64 \end{bmatrix}$$

is a Vandermonde matrix since it meets the definition with $x_1 = 2$, $x_2 = -3$, $x_3 = 1$, $x_4 = 4$.

Vandermonde matrices are not very interesting as numerical matrices, but instead appear more often in proofs and applications where the scalars x_i are carried as symbols. Two such applications are in the sections on secret-sharing ([section-SAS](#)) and curve-fitting ([4.2](#)). Principally, we would like to know when Vandermonde matrices are nonsingular, and the most convenient way to check this is by determining when the determinant is nonzero ([THEOREM SMZD](#)). As a bonus, the determinant of a Vandermonde matrix has an especially pleasing formula.

Theorem 5.1. *Suppose that A is a Vandermonde matrix of size n built with the scalars $x_1, x_2, x_3, \dots, x_n$. Then*

$$\det(A) = \prod_{1 \leq i < j \leq n} (x_j - x_i).$$

Proof. The proof is by induction ([PROOF TECHNIQUE I](#)) on n , the size of the matrix. An empty product for a 1×1 matrix might make a good base case, but we'll start

at $n = 2$ instead. For a 2×2 Vandermonde matrix, we have

$$\det(A) = \begin{vmatrix} 1 & x_1 \\ 1 & x_2 \end{vmatrix} = x_2 - x_1 = \prod_{1 \leq i < j \leq 2} (x_j - x_i)$$

For the induction step we will perform row operations on A to obtain the determinant of A as a multiple of the determinant of an $(n-1) \times (n-1)$ Vandermonde matrix. The notation in this theorem tends to obscure your intuition about the changes effected by various row and column manipulations. Construct a 4×4 Vandermonde matrix with four symbols as the scalars (x_1, x_2, x_3, x_4) , or perhaps a, b, c, d and play along with the example as you study the proof.

First we convert most of the first column to zeros. Subtract row n from each of the other $n-1$ rows to form a matrix B . By THEOREM DRCMA, B has the same determinant as A . The entries of B , in the first $n-1$ rows, i.e. for $1 \leq i \leq n-1$, $1 \leq j \leq n-1$, are

$$[B]_{ij} = x_i^{j-1} - x_n^{j-1} = (x_i - x_n) \sum_{k=0}^{j-2} x_i^{j-2-k} x_n^k$$

As the elements of row i , $1 \leq i \leq n-1$, have the common factor $(x_i - x_n)$, we form the new matrix C that differs from B by the removal of this factor from each of the first $n-1$ rows. This will change the determinant, as we will track carefully in a moment. We also have a first column with zeros in each location, except row n , so we can use it for a column expansion computation of the determinant. We now know,

$$\begin{aligned} \det(A) &= \det(B) \\ &= (x_1 - x_n)(x_2 - x_n) \cdots (x_{n-1} - x_n) \det(C) \\ &= (x_1 - x_n)(x_2 - x_n) \cdots (x_{n-1} - x_n)(1)(-1)^{n+1} \det(C(n-1|1)) \\ &= (x_1 - x_n)(x_2 - x_n) \cdots (x_{n-1} - x_n)(-1)^{n-1} \det(C(n-1|1)) \\ &= (x_n - x_1)(x_n - x_2) \cdots (x_n - x_{n-1}) \det(C(n-1|1)) \end{aligned}$$

For convenience, denote $D = C(n-1|1)$. Entries of this matrix are similar to those of B , but the factors used to build C are gone, and since the first column is gone, there is a slight re-indexing relative to the columns. For $1 \leq i \leq n-1$, $1 \leq j \leq n-1$,

$$[D]_{ij} = \sum_{k=0}^{j-1} x_i^{j-1-k} x_n^k$$

We will perform many column operations on the matrix D , always of the type where we multiply a column by a scalar and add the result to another column. As such, THEOREM DRCM insures that the determinant will remain constant. We will work column by column, left to right, to convert D into a Vandermonde matrix with scalars $x_1, x_2, x_3, \dots, x_{n-1}$. More precisely, we will build a sequence of matrices $D = D_1, D_2, \dots, D_{n-1}$, where each is obtainable from the previous by a sequence of determinant-preserving column operations and the first ℓ columns of D_ℓ are the first ℓ columns of a Vandermonde matrix with scalars $x_1, x_2, x_3, \dots, x_{n-1}$. We could establish this claim by induction (PROOF TECHNIQUE I) on ℓ if we were to expand the claim to specify the exact values of the final $n-1-\ell$ columns as well. Since the claim is that matrices with certain properties exist, we will instead establish the claim by constructing the desired matrices one-by-one procedurally. The extension to an inductive proof should be clear, but not especially illuminating.

Set $D_1 = D$ to begin, and note that the entries of the first column of D_1 are, for $1 \leq i \leq n-1$,

$$[D_1]_{i1} = \sum_{k=0}^{1-1} x_i^{1-1-k} x_n^k = 1 = x_i^{1-1}$$

So the first column of D_1 has the properties we desire. We will use this column of all 1's to remove the highest power of x_n from each of the remaining columns and so build D_2 . Precisely, perform the $n-2$ column operations where column 1 is multiplied by x_n^{j-1} and subtracted from column j , for $2 \leq j \leq n-1$. Call the result D_2 , and examine its entries in columns 2 through $n-1$. For $1 \leq i \leq n-1$, $2 \leq j \leq n-1$,

$$\begin{aligned} [D_2]_{ij} &= -x_n^{j-1} [D_1]_{i1} + [D_1]_{ij} \\ &= -x_n^{j-1}(1) + \sum_{k=0}^{j-1} x_i^{j-1-k} x_n^k \\ &= -x_n^{j-1} + x_i^{j-1-(j-1)} x_n^{j-1} + \sum_{k=0}^{j-2} x_i^{j-1-k} x_n^k \\ &= \sum_{k=0}^{j-2} x_i^{j-1-k} x_n^k \end{aligned}$$

In particular, we examine column 2 of D_2 . For $1 \leq i \leq n-1$,

$$[D_2]_{i2} = \sum_{k=0}^{2-2} x_i^{2-1-k} x_n^k = x_i^1 = x_i^{2-1}$$

Now, form D_3 . Perform the $n-3$ column operations where column 2 of D_2 is multiplied by x_n^{j-2} and subtracted from column j , for $3 \leq j \leq n-1$. The result is D_3 , whose entries we now compute. For $1 \leq i \leq n-1$,

$$\begin{aligned} [D_3]_{ij} &= -x_n^{j-2} [D_2]_{i2} + [D_2]_{ij} \\ &= -x_n^{j-2} x_i^1 + \sum_{k=0}^{j-2} x_i^{j-1-k} x_n^k \\ &= -x_n^{j-2} x_i^1 + x_i^{j-1-(j-2)} x_n^{j-2} + \sum_{k=0}^{j-3} x_i^{j-1-k} x_n^k \\ &= \sum_{k=0}^{j-3} x_i^{j-1-k} x_n^k \end{aligned}$$

Specifically, we examine column 3 of D_3 . For $1 \leq i \leq n-1$,

$$\begin{aligned} [D_3]_{i3} &= \sum_{k=0}^{3-3} x_i^{3-1-k} x_n^k \\ &= x_i^2 = x_i^{3-1} \end{aligned}$$

We could continue this procedure $n-4$ more times, eventually totaling $\frac{1}{2}(n^2 - 3n + 2)$ column operations, and arriving at D_{n-1} , the Vandermonde matrix of size $n-1$ built from the scalars $x_1, x_2, x_3, \dots, x_{n-1}$. Informally, we chop off the last term of every sum, until a single term is left in a column, and it is of the right form for

the Vandermonde matrix. This desired column is then used in the next iteration to chop off some more final terms for columns to the right. Now we can apply our induction hypothesis to the determinant of D_{n-1} and arrive at an expression for $\det A$,

$$\begin{aligned}
 \det(A) &= \det(C) \\
 &= \prod_{k=1}^{n-1} (x_n - x_k) \det(D) \\
 &= \prod_{k=1}^{n-1} (x_n - x_k) \det(D_{n-1}) \\
 &= \prod_{k=1}^{n-1} (x_n - x_k) \prod_{1 \leq i < j \leq n-1} (x_j - x_i) \\
 &= \prod_{1 \leq i < j \leq n} (x_j - x_i)
 \end{aligned}$$

which is the desired result. \square

Before we had 5.1 we could see that if two of the scalar values were equal, then the Vandermonde matrix would have two equal rows and hence be singular (THEOREM DERC, THEOREM SMZD). But with this expression for the determinant, we can establish the converse.

Theorem 5.2. *A Vandermonde matrix of size n with scalars $x_1, x_2, x_3, \dots, x_n$ is nonsingular if and only if the scalars are all different.*

Proof. Let A denote the Vandermonde matrix with scalars $x_1, x_2, x_3, \dots, x_n$. By THEOREM SMZD, A is nonsingular if and only if the determinant of A is nonzero. The determinant is given by 5.1, and this product is nonzero if and only if each term of the product is nonzero. This condition translates to $x_i - x_j \neq 0$ whenever $i \neq j$. In other words, the matrix is nonsingular if and only if the scalars are all different. \square

Appendix A

GNU Free Documentation License

Version 1.3, 3 November 2008

Copyright ©2000, 2001, 2002, 2007, 2008 [Free Software Foundation, Inc.](#)

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

0. PREAMBLE The purpose of this License is to make a manual, textbook, or other functional and useful document “free” in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of “copyleft”, which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

1. APPLICABILITY AND DEFINITIONS This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The “Document”, below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as “you”. You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A “Modified Version” of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A “Secondary Section” is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document’s overall subject (or to related matters) and contains

nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The “Invariant Sections” are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The “Cover Texts” are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A “Transparent” copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not “Transparent” is called “Opaque”.

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The “Title Page” means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, “Title Page” means the text near the most prominent appearance of the work’s title, preceding the beginning of the body of the text.

The “publisher” means any person or entity that distributes copies of the Document to the public.

A section “Entitled XYZ” means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as “Acknowledgements”, “Dedications”, “Endorsements”, or “History”.) To “Preserve the Title” of such a section when you modify the Document means that it remains a section “Entitled XYZ” according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

2. VERBATIM COPYING You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. COPYING IN QUANTITY If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

4. MODIFICATIONS You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

1. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
2. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at

least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.

3. State on the Title page the name of the publisher of the Modified Version, as the publisher.
4. Preserve all the copyright notices of the Document.
5. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
6. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
7. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
8. Include an unaltered copy of this License.
9. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
10. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
11. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
12. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
13. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
14. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
15. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled “Endorsements”, provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

5. COMBINING DOCUMENTS You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled “History” in the various original documents, forming one section Entitled “History”; likewise combine any sections Entitled “Acknowledgements”, and any sections Entitled “Dedications”. You must delete all sections Entitled “Endorsements”.

6. COLLECTIONS OF DOCUMENTS You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

7. AGGREGATION WITH INDEPENDENT WORKS A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an “aggregate” if the copyright resulting from the compilation is not used to limit the legal rights of the compilation’s users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document's Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

8. TRANSLATION Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled "Acknowledgements", "Dedications", or "History", the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

9. TERMINATION You may not copy, modify, sublicense, or distribute the Document except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, or distribute it is void, and will automatically terminate your rights under this License.

However, if you cease all violation of this License, then your license from a particular copyright holder is reinstated (a) provisionally, unless and until the copyright holder explicitly and finally terminates your license, and (b) permanently, if the copyright holder fails to notify you of the violation by some reasonable means prior to 60 days after the cessation.

Moreover, your license from a particular copyright holder is reinstated permanently if the copyright holder notifies you of the violation by some reasonable means, this is the first time you have received notice of violation of this License (for any work) from that copyright holder, and you cure the violation prior to 30 days after your receipt of the notice.

Termination of your rights under this section does not terminate the licenses of parties who have received copies or rights from you under this License. If your rights have been terminated and not permanently reinstated, receipt of a copy of some or all of the same material does not give you any rights to use it.

10. FUTURE REVISIONS OF THIS LICENSE The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See [Copyleft](#).

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation. If the Document specifies that a proxy can decide which future versions of this License can be used, that proxy's public

statement of acceptance of a version permanently authorizes you to choose that version for the Document.

11. RELICENSING “Massive Multiauthor Collaboration Site” (or “MMC Site”) means any World Wide Web server that publishes copyrightable works and also provides prominent facilities for anybody to edit those works. A public wiki that anybody can edit is an example of such a server. A “Massive Multiauthor Collaboration” (or “MMC”) contained in the site means any set of copyrightable works thus published on the MMC site.

“CC-BY-SA” means the Creative Commons Attribution-Share Alike 3.0 license published by Creative Commons Corporation, a not-for-profit corporation with a principal place of business in San Francisco, California, as well as future copyleft versions of that license published by that same organization.

“Incorporate” means to publish or republish a Document, in whole or in part, as part of another Document.

An MMC is “eligible for relicensing” if it is licensed under this License, and if all works that were first published under this License somewhere other than this MMC, and subsequently incorporated in whole or in part into the MMC, (1) had no cover texts or invariant sections, and (2) were thus incorporated prior to November 1, 2008.

The operator of an MMC Site may republish an MMC contained in the site under CC-BY-SA on the same site at any time before August 1, 2009, provided the MMC is eligible for relicensing.

ADDENDUM: How to use this License for your documents To use this License in a document you have written, include a copy of the License in the document and put the following copyright and license notices just after the title page:

Copyright ©YEAR YOUR NAME

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.3 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is included in the section entitled “GNU Free Documentation License”.

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the “with . . . Texts.” line with this:

with the Invariant Sections being LIST THEIR TITLES, with the Front-Cover Texts being LIST, and with the Back-Cover Texts being LIST.

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.

