DIAGONALIZATION OF MATRICES

A MAJOR QUALIFYING PROJECT REPORT

SUBMITTED TO THE FACULTY OF
WORCESTER POLYTECHNIC INSTITUTE
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE
DEGREE OF BACHELOR OF SCIENCE

WRITTEN BY

THEOPHILUS MARKS

APPROVED BY: PADRAIG Ó CATHÁIN

MAY 5, 2021

Abstract

This thesis aims to study criteria for diagonalizability over finite fields. First we review basic and major results in linear algebra, including the Jordan Canonical Form, the Cayley-Hamilton theorem, and the spectral theorem; the last of which is an important criterion for for matrix diagonalizability over fields of characteristic 0, but fails in finite fields. We then calculate the number of diagonalizable and non-diagonalizable 2×2 matrices and 2×2 symmetric matrices with repeated and distinct eigenvalues over finite fields. Finally, we look at results by Brouwer, Gow, and Sheekey for enumerating symmetric nilpotent matrices over finite fields.

Summary

(1) In Chapter 1 we review basic linear algebra including eigenvectors and eigenvalues, diagonalizability, generalized eigenvectors and eigenspaces, and the Cayley-Hamilton theorem (we provide an alternative proof in the appendix). We explore the Jordan canonical form and the rational canonical form, and prove some necessary and sufficient conditions for diagonalizability related to the minimal polynomial and the JCF.

(2) In Chapter 2 we review inner products and inner product spaces, adjoints, and isometries. We prove the spectral theorem, which tells us that symmetric, orthogonal, and self-adjoint matrices are all diagonalizable over fields of characteristic zero. In particular, a nilpotent symmetric matrix in characteristic zero must be the zero matrix. We will explain later why the spectral theorem does not hold over finite fields.

(3) In Chapter 3 we start with a review of group theory and the orbit-stabilizer theorem, which we use to count 2×2 matrices over finite fields. We do this computation in four cases deriving from the specification of a matrix as being diagonalizable or non-diagonalizable, and having repeated or distinct eigenvalues. We see that the number of diagonalizable matrices depends on $q \mod 4$. We then repeat the same enumeration for 2×2 symmetric matrices. These counts turn out to be related to cyclotomic numbers, which are the number of solutions to the equation $x^2 + 1 = y^2$ in $\mathbb{F}_q$. We conclude the chapter with some observations on 2×2 symmetric nilpotent matrices.

(4) In Chapter 4 we review bilinear and quadratic forms, symmetric and alternating forms, and isotropic vectors. We see that isotropic vectors explain why the spectral theorem fails over finite fields, since it requires the absence of isotropic vectors, and it is known as a consequence of the Chevalley-Warning Theorem that a quadratic form in positive characteristic has an isotropic vector in dimension ≥ 3. Finally, we review the fitting decomposition of a matrix and give a theorem of Hall enumerating the nilpotent matrices in a vector space over a finite field. We then explore the enumeration of symmetric nilpotent matrices using methods developed by Brouwer, Gow, and Sheekey. They give an algorithm involving combinatorics on Young diagrams to evaluate the number of symmetric nilpotent matrices with a given Jordan Canonical Form in any dimension and for any finite field.

(5) In Chapter 5 we state our conclusions and give several possible directions for future research.

Acknowledgements

Special thanks to Padraig Ó Catháin and Georgina Quinn for helping write this MQP.

CONTENTS

1. Structure of a Linear Operator

In this chapter we prove two important theorems in linear algebra: the Jordan Canonical Form Theorem and the Cayley-Hamilton Theorem. The Jordan Canonical Form gives a necessary and sufficient condition for diagonalizability over algebraically closed fields, while the Cayley-Hamilton theorem gives a result which holds for arbitrary fields. We also discuss the Rational Canonical Form, which can tell us if a matrix is diagonalizable. Later in the thesis, we will consider diagonalizability over finite fields.
The material presented here is standard, and proofs may be found in [Axl15] and [FIS97]. We assume familiarity with linear algebra up to the definition of eigenvalues and eigenvectors.

1.1. Eigenvectors and eigenvalues.

Definition 1.1. *Let $M : V \to V$ be a linear transformation. A non-zero vector $v \in V$ is an eigenvector of M if $M(v) = \lambda v$ for some scalar λ.*

If λ is a scalar for which there exists a non-zero vector solution to the equation $M(v) = \lambda v$ then λ is an eigenvalue of M. The set of all eigenvectors corresponding to λ is the eigenspace associated to λ.

We will now give a characterisation of the values λ which are eigenvalues of a linear operator T.

Proposition 1.1. *Let V and W be vector spaces and let $T : W \to V$ be linear. If V is finite-dimensional then*

$$\textit{nullity}(T) + \textit{rank}(T) = \textit{dim}(T).$$

Proof. [FIS97]. □

Proposition 1.2. *The scalar λ is an eigenvalue of T if and only if $T - \lambda I$ is non-invertible.*

Proof. Suppose that $T - \lambda I$ is not invertible, then it is not a bijection $V \to V$, and by 1.1 there exists a non-zero vector v such that $(T - \lambda I)v = \mathbf{0}$. But this is equivalent to $Tv = \lambda v$, so v is an eigenvector of T with eigenvalue λ.

In the other direction, if λ is an eigenvalue, there exists $v \in V$ such that $Tv = \lambda v$. Rearranging the equation shows that $v \in N(T - \lambda I)$. □

Definition 1.2. *Let V be a vector space and let $T : V \to V$ be linear. A subspace W of V is called **T-invariant** if $T(x) \in W$ for every $x \in W$, that is, $T(W) \subseteq W$. If W is T-invariant, we define the **restriction of T on W** to be the function $T_W : W \to W$ defined by $T_W(x) = T(x)$ for all $x \in W$.*

Proposition 1.3. *The set of all eigenvectors of $T : V \to V$ with eigenvalue λ form a T-invariant subspace of V.*

Proof. Recall that U is a subspace of V if and only if it is closed under vector addition and scalar multiplication. Suppose that $T(v_1) = \lambda v_1$ and that $T(v_2) = \lambda v_2$ (for the same value of λ!). Then

$$T(v_1 + v_2) = T(v_1) + T(v_2) = \lambda v_1 + \lambda v_2 = \lambda(v_1 + v_2)\,.$$

Similarly,

$$T(\alpha v_1) = \alpha T(v_1) = \alpha\lambda v_1 = \lambda(\alpha v_1)\,.$$

□

The next proof shows that the eigenspaces of T are disjoint.

Proposition 1.4. *Suppose that $\lambda_1, \ldots, \lambda_m$ are distinct eigenvalues of T with eigenvectors $v_1, \ldots, v_m$. Then $\{v_1, \ldots, v_m\}$ is a linearly independent set of vectors.*

Proof. This is a proof by contradiction. Suppose the claim to be false, and let k be the smallest integer such that $\{v_1, \ldots, v_k\}$ is linearly dependent. Then there exist scalars $\alpha_1, \ldots, \alpha_{k-1}$ (not all zero) such that

$$v_k = \alpha_1 v_1 + \alpha_2 v_2 + \ldots + \alpha_{k-1} v_{k-1}\,. \tag{1}$$

Apply T to both sides of Equation 1, to get

$$\lambda_k v_k = \alpha_1\lambda_1 v_1 + \alpha_2\lambda_2 v_2 + \ldots + \alpha_{k-1}\lambda_{k-1} v_{k-1}\,.$$

Multiply both sides of Equation 1 by λ_k to get

$$\lambda_k v_k = \alpha_1\lambda_k v_1 + \alpha_2\lambda_k v_2 + \ldots + \alpha_{k-1}\lambda_k v_{k-1}\,,$$

and subtract one from the other to get

$$\mathbf{0} = \alpha_1(\lambda_1 - \lambda_k)v_1 + \alpha_2(\lambda_2 - \lambda_k)v_2 + \ldots + \alpha_{k-1}(\lambda_{k-1} - \lambda_k)v_{k-1}\,.$$

Since the λ_i are distinct, the coefficients on the right hand side are not identically zero, and we have produced a linear dependence among the set $\{v_1, \ldots, v_{k-1}\}$ which contradicts the choice of k. □

Corollary 1.1. *Let $T : V \to V$ where dim$(V) = n$. Then T has at most n distinct eigenvalues.*

Proof. By Proposition 1.4, eigenvectors with distinct eigenvalues are linearly independent, and a linearly independent set has size at most dim(T). □

Corollary 1.1 is a statement about *eigenvalues* not *eigenvectors*. A moment's consideration of Proposition 1.3 shows that as soon as T has an eigenspace of dimension 2, it is possible for T to have **many** distinct eigenvectors.

Theorem 1.1. *[Axl15] Suppose that V is a finite dimensional vector space over an algebraically closed field k and $T : V \to V$ is a linear transformation. Then T has an eigenvector.*

Proof. Let n be the dimension of V and consider the set $\{v, T(v), T^2(v), \ldots, T^n(v)\}$ which contains $n + 1$ vectors and so must be linearly dependent. Suppose that

$$\alpha_0 v + \alpha_1 T(v) + \ldots + \alpha_n T^n(v) = \mathbf{0}\,,$$

and set $p(x) = \alpha_n x^n + \alpha_{n-1}x^{n-1} + \ldots + \alpha_0$. Since k is algebraically closed,

$$p(x) = (x - \gamma_0)\,(x - \gamma_1)\cdots(x - \gamma_n)\ ,$$

for not necessarily distinct scalars $\gamma_0, \ldots, \gamma_n$. Then

$$p(T)v = (T - \gamma_0 I)\,(T - \gamma_1 I) \cdots (T - \gamma_n I)\,v = \mathbf{0}\,.$$

If all of the terms $T - \gamma_k I$ were invertible, then the product would be invertible too. Hence there exists a scalar γ_k for which $N(T - \gamma_k I)$ is non-zero. By Proposition 1.2, the γ_k is an eigenvalue of T. □

The algebraically closed condition in Theorem 1.1 is (really) necessary, as can be seen already with the matrix

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

over any field which is not of characteristic 2 and which does not contain a square root of -1.

1.2. **Upper Triangular Matrices and Invariant Subspaces.**

Definition 1.3. *A matrix M is **upper triangular** if all entries below the diagonal are zero.*

Proposition 1.5. *Let V be a finite dimensional vector space over an algebraically closed field. Let $T : V \to V$ be a linear transformation. Let $B = \{v_1, \ldots, v_n\}$ be a basis for V. The following are equivalent.*

(1) *$Tv_k \in \langle v_1, \ldots, v_k \rangle$ for $1 \leq k \leq n$.*
(2) *$\langle v_1, \ldots, v_k \rangle$ is T-invariant for $1 \leq k \leq n$.*
(3) *The matrix of T with respect to B is upper triangular.*

Proof. **1 implies 2:** Immediate from the definition of a T-invariant subspace.
2 implies 3: This follows from a careful application of the rules for matrices. Recall that the k^{th} column of $[T]_B$ is $[Tv_k]_B$, which is a column vector expressing v_k as a linear combination of vectors from B. Since $v_k \in \langle v_1, \ldots, v_k \rangle$ this column vector has non-zero entries only for $i \leq k$. These are the entries above the diagonal. Hence $[T]_B$ is upper triangular as required.
3 implies 1: This follows directly from the assumption that T is upper-triangular. □

Definition 1.4. *The configuration of subspaces considered in Proposition 1.5 is often called a **chain of invariant subspaces** for T.*

We will prove that any operator on a finite dimensional vector space admits a chain of invariant subspaces. A one dimensional invariant subspace is an eigenvector, this will be the base case for our induction. The inductive step will require the following construction.

Proposition 1.6. *Suppose that U is a T-invariant subspace of V. Then T/U is a vector space, on which T acts by*

$$T(v + U) = T(v) + U\,.$$

The image of a T invariant subspace of V is a T-invariant subspace of V/U, and the preimage of a T-invariant subspace of V/U is a T-invariant subspace of V.

Proof. The cosets of U in V form a vector space. It suffices to check that the action of T on V/U is linear.

$$\begin{aligned} T(v_1 + v_2 + U) &= T(v_1 + v_2) + U = T(v_1) + T(v_2) + U \\ &= (T(v_1) + U) + (T(v_2) + U) = T(v_1 + U) + T(v_2 + U) \end{aligned}$$

The proof for scalar multiplication is similar.

Suppose that W is a T-invariant subspace of V. By definition, $W/U = \{w+U \mid w \in W\}$. Since W is T-invariant, $T(w + U) = T(w) + U \in W/U$ and W/U is T-invariant. Similarly, suppose that X is a T-invariant subspace of V/U. Define the preimage of X to be $X^+ = \{x \in V \mid x + U \in X\}$. Then for any $x, y \in X^+$,

$$(x + U) + (y + U) = (x + y) + U\,,$$

so $x+y \in X^+$, and similarly for scalar multiplication, and closure under the action of T. □

Now we can prove our result.

Theorem 1.2 (Cauchy)**.** *Let V be a finite dimensional vector space over an algebraically closed field, and $T : V \to V$ a linear transformation. There exists a chain $U_1 \leq U_2 \leq \cdots \leq U_n = V$ of T-invariant subspaces of V. Equivalently, there exists a basis of V with respect to which T is upper triangular.*

Proof. The proof will be by induction on the dimension of V. The result holds trivially if $\dim(V) = 1$. Suppose that the desired result holds in dimension $\leq k$, and let V be of dimension $k + 1$.

By hypothesis, T has an eigenvalue λ with corresponding eigenvector v, and $U = \langle v \rangle$ is a T-invariant subspace of V. The dimension of V/U is k, so by induction there exists a basis $B = \{v_1 + U, \ldots, v_k + U\}$ of V/U with respect to which T is upper triangular. By Proposition 1.5, $T(v_i + U) \in \langle v_1 + U, \ldots, v_i + U \rangle$ for any $1 \leq i \leq k$.

Since $U = \langle v \rangle$, it follows that $T(v) \in \langle v \rangle$ and $T(v_i) \in \langle v, v_1, \ldots, v_i \rangle$ for each $1 \leq i \leq k$. By Proposition 1.5, the matrix of T with respect to the ordered basis $\{v, v_1, \ldots, v_k\}$ is upper triangular, as required. □

Proposition 1.7. *Suppose $T \in \mathcal{L}(V)$ and $p(x) \in \mathbb{F}[x]$. Then $N(p(T))$ and $R(p(T))$ are T-invariant.*

Proof. Suppose $v \in N(p(T))$. Then

$$(p(T))(Tv) = T(p(T)v) = T(0) = 0$$

thus $Tv \in N(p(T))$ and hence $N(p(T))$ is T-invariant.
Now suppose $v \in R(p(T))$. Let $u \in V$ be such that $p(T)u = v$. Then

$$p(T)(Tu) = T(p(T)u) = Tv$$

thus $Tv \in R(p(T))$ and hence $R(p(T))$ is T-invariant. □

1.3. **diagonalization: a basis of eigenvectors.**

So far, we have seen in Theorem 1.1 that a matrix defined over an algebraically closed field has an eigenvector. Cauchy's theorem on upper triangular matrices, Theorem 1.2, is an inductive version of this: for any matrix M there exists a basis $\{v_1, v_2, \ldots, v_n\}$ of the underlying space and a sequence of subspaces $U_i = \{v_1, \ldots, v_i\}$ such that $v_{i+1} + U_i$ is an eigenvector of V/U_i.

A particularly nice special case of Cauchy's theorem arises when every v_i is *itself an eigenvector*.

Proposition 1.8. *The set $B = \{v_1, \ldots, v_n\}$ is a basis for V consisting of eigenvectors of M if and only if $[M]_B$ is a diagonal matrix.*

Proof. Suppose that B is a basis of eigenvectors. By definition, the i^{th} column of M is $[Mv_i]_B$. By hypothesis, this vector is non-zero only in the i^{th} entry, so $[M]_B$ is diagonal.

In the other direction, if $[M]_B$ is diagonal, then the standard basis vectors *with respect to* B are eigenvectors of M. The result follows. □

Proposition 1.9. *Suppose that M, N are square matrices and X is an invertible matrix such that $M = XNX^{-1}$. Then for any positive integer n,*

$$M^n = XN^nX^{-1}\,.$$

Proof. We will proceed by induction. The claim is trivial for $n = 1$, let us verify it for $n = 2$:

$$\begin{aligned} M^2 &= \left(XNX^{-1}\right)\left(XNX^{-1}\right) \\ &= XN\left(X^{-1}X\right)NX^{-1} \\ &= XN^2X^{-1} \end{aligned}$$

Suppose that the induction hypothesis holds for $n = k$. Then

$$M^kM = \left(XN^kX^{-1}\right)\left(XNX^{-1}\right) = XN^{k+1}X^{-1}\,.$$

□

In this thesis, we will prove several necessary and sufficient criteria for a matrix to be diagonalizable. For now, we give only one.

Theorem 1.3. *Suppose that M is an $n \times n$ matrix with n distinct eigenvalues. Then M is diagonalizable.*

Proof. By Proposition 1.4, the eigenvectors corresponding to distinct eigenvalues are linearly independent, so there exists a basis of eigenvectors and Proposition 1.8 applies. □

Diagonalization is a useful property because diagonal matrices have several known properties which simplify computations involving them. For example, a power of a diagonal matrix can be computed by raising its diagonal entries to that power, and the determinant of a diagonal matrix is equal to the product of its diagonal entries.

1.4. **Generalised Eigenvectors and Generalised Eigenspaces.**

Recall that $v \in V$ is an *eigenvector* of T with eigenvalue λ if and only if $(T - \lambda I)v = \mathbf{0}$. Here we will explore some obstructions to diagonalization.

Definition 1.5. *A vector $v \in V$ is a **generalised eigenvector** of T with eigenvalue λ if and only if there exists a positive integer m such that*

$$(T - \lambda I)^m v = \boldsymbol{0}\,.$$

*We call the least m which satisfies the above equation the **height** of the generalised eigenvector.*

As an example, consider the matrix

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

which has $[1,0,0]^\top$ as an eigenvector, $[0,1,0]^\top$ as a generalised eigenvector of height 2 and $[0,0,1]^\top$ as a generalised eigenvector of height 3. Expressions like this can always be evaluated as polynomials in λ and the (finitely many) vectors $(T - \lambda I)^m v$ of lower height obtained from v.

We begin by showing that the generalised eigenvectors form a subspace of V. The result as proved is a little non-constructive as there is not yet an explicit bound on the height: we provide this in Proposition 1.11.

Proposition 1.10. *Suppose that V has dimension n and let $T : V \to V$ be a linear operator. The generalised eigenvectors with eigenvalue λ form a T-invariant subspace of V, denoted G_λ.*

Proof. Suppose that v_1, v_2 are generalised eigenvectors of T with eigenvalue λ. Say $(T - \lambda I)^{m_1} v_1 = (T - \lambda I)^{m_2} v_2 = \mathbf{0}$. Let $m = \max\{m_1, m_2\}$, then

$$(T - \lambda I)^m (v_1 + v_2) = (T - \lambda I)^m v_1 + (T - \lambda I)^m v_2 = \mathbf{0}\,,$$

and

$$(T - \lambda I)^{m_1} (\alpha v_1) = \alpha \mathbf{0} = \mathbf{0}\,.$$

So the generalised eigenvectors form a subspace G_λ of V, as required.

To see that G_λ is T-invariant, observe that $T(T - \lambda I)^m = (T - \lambda I)^m T$ for any non-negative integer m. Suppose that $v \in G_\lambda$ and $(T - \lambda I)^m v = \mathbf{0}$. Then

$$(T - \lambda I)^m (Tv) = T(T - \lambda I)^m v = T\mathbf{0} = \mathbf{0}$$

so $Tv \in G_\lambda$ and so G_λ is T-invariant. □

Next, we begin to explore the structure of a G_λ.

Definition 1.6. *Let $v \in G_\lambda$ be a generalised eigenvector. The **cycle** of v is*

$$\langle (T - \lambda I)^m v \mid m \in \mathbb{N} \rangle\,.$$

*We say that a cycle is **maximal** in G_λ if it is not contained in a larger cycle.*

Proposition 1.11. *Suppose that $v \in V$ is a generalised eigenvector of T with eigenvalue λ and height m. Then the vectors $(T-\lambda I)^j v$ for $j = 0, 1, \ldots, m-1$ are linearly independent.*

Proof. We write $M = (T-\lambda I)$ for notational convenience. Suppose that there exists a linear dependence between the vectors $M^j v$, say

$$M^{m-1}v + \alpha_{m-2}M^{m-2}v + \ldots + \alpha_1 Mv + \alpha_0 v = \mathbf{0}\,.$$

We multiply this equation by M^{m-1} on both sides:

$$M^{2m-2}v + \alpha_{m-1}M^{2m-3} + \ldots + \alpha_1 M^m v + \alpha_0 M^{m-1}v = \mathbf{0}\,,$$

but $M^{m+k}v = M^k\,(M^m v) = \mathbf{0}$, since v is a generalised eigenvector of height m. So this equation is equivalent to $\alpha_0 T^{m-1}v = \mathbf{0}$, and we conclude that $\alpha_0 = 0$.

Next, multiply the equation

$$M^{m-1}v + \alpha_{m-2}M^{m-2}v + \ldots + \alpha_1 Mv = \mathbf{0}$$

by M^{m-2} to show that $\alpha_1 = 0$. Proceeding in this way, one finds that all coefficients are 0 and so the set

$$\{M^{m-1}v, M^{m-2}v, \ldots, Mv, v\}$$

is linearly independent, as required. □

We have shown that a generalised eigenvector of height m generates a cycle which is a subspace of dimension m contained in G_λ. One consequence of Proposition 1.11 is that when $\dim(V) = n$, there can be no generalised eigenvectors of height greater than n. Note that there may be multiple linearly independent generalised eigenvectors of the same height: this happens already for eigenvectors, where there may be multiple linearly independent eigenvectors with the same eigenvalue.

Definition 1.7. *A* direct sum *of a matrix A of size $m \times n$ and matrix B of size $p \times q$ is a matrix of size $(m+p) \times (n+q)$:*

$$A \oplus B = \begin{bmatrix} A & 0 \\ 0 & B \end{bmatrix}$$

Proposition 1.12. *Let $T : V \to V$ be a linear operator on a vector space of dimension n. Then $V = N(T-\lambda I)^n \oplus R(T-\lambda I)^n$.*

Proof. The Rank-Nullity theorem shows that the dimensions of $\mathrm{N}(T-\lambda I)^n$ and $\mathrm{R}(T-\lambda I)^n$ sum to n. Suppose that $v \in \mathrm{N}(T-\lambda I)^n \cap \mathrm{R}(T-\lambda I)^n$. Then there exists $u \in V$ such that $(T-\lambda I)^n u = v$ and $(T-\lambda I)^{2n} u = \mathbf{0}$. Thus v is a generalised eigenvector of T of height greater than n, which contradicts Proposition 1.11. □

Note that raising $(T-\lambda I)$ to a sufficiently high power is necessary in Proposition 1.12. The range and nullspace of an arbitrary matrix are not disjoint.

Definition 1.8. *An operator is called **nilpotent** if some power of it equals 0.*

Example 1.1. *The operator $T \in \mathcal{L}(\mathbb{F}^3)$ defined by*

$$T(x_1, x_2, x_3) = (x_2, x_3, 0)$$

is nilpotent because $T^3 = 0$.

Proposition 1.13. *Suppose $T \in \mathcal{L}(V)$ is nilpotent. Then $T^{dim(V)} = 0$.*

Proof. Since T is nilpotent, $G(0, T) = V$. Thus Proposition 1.12 implies that

$$N(T^{\dim(V)}) = V.$$

□

1.4.1. *Distinct Generalised Eigenspaces are Disjoint.* Let $T : V \to V$ be a linear operator and write G_λ for the generalised eigenspace of T with eigenvalue λ. In this section, we show that the generalised eigenspaces for distinct eigenvalues are disjoint. This is the heart of the Cayley-Hamilton theorem.

First we need a result about systems of linear equations of a special form.

Proposition 1.14. *The $n \times n$ Vandermonde matrix is*

$$\begin{bmatrix} 1 & x_1 & x_1^2 & \dots & x_1^{n-1} \\ 1 & x_2 & x_2^2 & \dots & x_2^{n-1} \\ \vdots & \vdots & \vdots & & \vdots \\ 1 & x_n & x_n^2 & \dots & x_n^{n-1} \end{bmatrix},$$

where $x_1, \dots, x_n \in \mathbb{C}$. The Vandermonde matrix is invertible if and only if all of the x_i are distinct.

Proof. (Sketch.) Recall that the determinant of a matrix is a linear combination of terms formed from taking a product of matrix entries, one from each row and column. The determinant is 0 if and only if the matrix is not invertible.

Clearly, if $x_i = x_j$ the Vandermonde matrix has a repeated row, and cannot be invertible. Hence, $(x_i - x_j)$ divides the determinant for any $i < j$. By inspection, the terms in the determinant all have degree $1 + 2 + \ldots + n - 1 = \binom{n}{2}$. The polynomial $\prod_{i<j} x_i - x_j$ is also of degree $\binom{n}{2}$ and divides the determinant. So these polynomials are equal. Hence the determinant is non-zero if and only if the x_i are all distinct. □

With the Vandermonde determinant in hand, we can prove the base case of our argument, which is interesting in its own right. We show that (proper) eigenvectors with distinct eigenvalues are linearly independent.

Proposition 1.15. *Suppose that $v_1, \dots, v_k$ are eigenvectors of T with distinct eigenvalues $\lambda_1, \dots, \lambda_t$. Then $v_1, \dots, v_k$ are linearly independent.*

Proof. Suppose that there were a linear dependence between the v_i:

$$\alpha_1 v_1 + \alpha_2 v_2 + \ldots + \alpha_k v_k = \mathbf{0}\,.$$

Multiplying both sides by T^j we obtain

$$\lambda_1^j \alpha_1 v_1 + \lambda_2^j \alpha_2 v_2 + \ldots + \lambda_k^j \alpha_k v_k = \mathbf{0}\,,$$

for any positive integer j. Taking the the first k of these equations we form a linear system

$$\begin{bmatrix} 1 & 1 & \dots & 1 \\ \lambda_1 & \lambda_2 & \dots & \lambda_k \\ \vdots & \vdots & \vdots & \vdots \\ \lambda_1^{k-1} & \lambda_2^{k-1} & \dots & \lambda_k^{k-1} \end{bmatrix} \begin{bmatrix} \alpha_1 v_1 \\ \alpha_2 v_2 \\ \vdots \\ \alpha_k v_k \end{bmatrix} = \mathbf{0}\,.$$

The first term is the (transpose of the) Vandermonde matrix, which is invertible since the λ_i are distinct. Hence the only solution of the linear system has $\alpha_i = 0$ for all $1 \leq i \leq k$. □

Finally, we prove our main result.

Theorem 1.4. *Suppose that $v_1, \dots, v_k$ are generalised eigenvectors of a linear operator T with distinct eigenvalues $\lambda_1, \dots, \lambda_k$. Then $v_1, \dots, v_k$ are linearly independent.*

Proof. By definition, a generalised eigenvector of height 1 is an (ordinary) eigenvector. Write e_j for the height of v_j, and observe that $T^{e_j - t} v_j$ is a generalised eigenvector of height t. In particular, $w_j = T^{e_j - 1} v_j$ is an ordinary eigenvector, with eigenvalue λ.

Next, observe that $(T - \lambda_i I) v_j = (\lambda_j - \lambda_i) v_j$. Set $M = \prod_{j=1}^{k} (T - \lambda_j)^{e_j - 1}$, then

$$M v_j = \prod_{i \neq j} (\lambda_i - \lambda_j)^{e_i - 1} w_j\,,$$

which is an eigenvector of T with eigenvalue λ_j. We write γ_j for the scalar such that $M v_j = \gamma_j w_j$.

Now, suppose that

$$\alpha_1 v_1 + \alpha_2 v_2 + \ldots + \alpha_k v_k = \mathbf{0}$$

is a linear dependence between the generalised eigenvectors v_i. Applying M to this equation, we obtain

$$\alpha_1 \gamma_1 w_1 + \alpha_2 \gamma_2 w_2 + \ldots + \alpha_k \gamma_k w_k = \mathbf{0}\,,$$

which is a linear equation between eigenvectors of T. By Proposition 1.15, the scalars α_i are all zero, so the generalised eigenvectors are linearly independent. □

An important corollary is that generalised eigenspaces are linearly independent. We make this precise below.

Corollary 1.2. *Let $G_1, \dots, G_k$ be generalised eigenspaces of $T : V \to V$, and let $v_i \in G_i$ for $i = 1, \dots, k$. If*

$$\alpha_1 v_1 + \ldots + \alpha_k v_k = \boldsymbol{0}$$

then $\alpha_1 = \alpha_2 = \ldots = \alpha_k = 0$.

Corollary 1.2 can be rearranged to show that no linear combination $\alpha_1 v_1 + \ldots + \alpha_{k-1} v_{k-1}$ belongs to G_k, which is a substantially stronger claim than that the generalised eigenspaces intersect pairwise in $\{0\}$.

1.5. **The Cayley-Hamilton Theorem.**

We complete the proof of the Cayley-Hamilton theorem, by showing that V is a direct sum of the generalised eigenspaces.

Theorem 1.5 (Cayley Hamilton, I)**.** *Let $T : V \to V$ be a linear operator over an algebraically closed field, and $G_1, \ldots, G_k$ the generalised eigenspaces of T. Then $T = \oplus_{i=1}^{k} G_i$.*

Proof. We showed in Theorem 1.4 that the generalised eigenspaces are disjoint. It will suffice to show that they span V. We will prove this by induction. The base case is $\dim V = 1$, in which case the result holds trivially.

Suppose that the result holds for all vector spaces with $\dim V \leq t$ and let V be a vector space of dimension $t+1$. Since T is defined over an algebraically closed field, T has an eigenvector, say $Tv = \lambda v$. So the generalised eigenspace G_λ is nonempty, and T-invariant by Proposition 1.10. Recall that $G_\lambda = \mathrm{N}(T - \lambda I)^{t+1}$ by Proposition 1.11, and that $U = \mathrm{R}(T - \lambda I)^{t+1}$ is disjoint from G_λ by Proposition 1.12. In fact, $V = G_\lambda \oplus U$ where each of these spaces is T-invariant. Hence any vector in V has a **unique** expression as $v = w + u$ where $w \in G_\lambda$ and $u \in U$.

Now, T has dimension $\leq t$ so the inductive hypothesis applies. Say $U = G_2 \oplus G_3 \oplus \cdots \oplus G_k$, where G_i is a generalised eigenspace with eigenvalue λ_i. We need to show that each generalised eigenspace of U is also a generalised eigenspace of V. It will suffice to show that a generalised eigenvector not contained in G_λ is contained in U. Suppose that $w \in G_\lambda$ and $u \in U$, such that $w + u$ is a generalised eigenvector of V with eigenvalue μ, distinct from λ (since otherwise $w + u$ would be contained in G_λ, by definition). By Proposition 1.11, $(T - \mu I)^{t+1}(w + u) = \mathbf{0}$. Hence $(T - \mu I)^{t+1} w = \mathbf{0}$, and $w \in G_\lambda \cap G_\mu$. But Theorem 1.4 forces $w = \mathbf{0}$ so the generalised eigenvector belongs to U, and by the induction hypothesis is contained in one of the G_i. Hence, G is a direct sum of generalised eigenspaces. □

Recalling Cauchy's theorem, any linear transformation over an algebraically closed field may be written in upper triangular form with respect to a suitable basis. The eigenvalues of such a linear transformation are the diagonal entries of the matrix, and the multiplicities of these entries are the dimensions of the generalised eigenspaces. The standard definition of the characteristic polynomial is as $\det(M - tI)$, which is a polynomial in t. (It is not easily seen from this definition that the characteristic polynomial is independent of the choice of basis, but it is.) The polynomial vanishes if and only if the matrix fails to be invertible. Hence its roots are the eigenvalues of M (as is easily seen by computing the characteristic polynomial with respect to a Cauchy-basis).

Definition 1.9. *Suppose $T \in \mathcal{L}(V)$. Let $\lambda_1, ..., \lambda_m$ denote the distinct eigenvalues of T, with corresponding algebraic multiplicities $d_1, ..., d_m$. The **characteristic polynomial** of T is defined to be*

$$\prod_{i=1}^{m} (x - \lambda_i)^{d_i}.$$

Now we can state the standard version of the Cayley-Hamilton theorem.

Theorem 1.6 (Cayley-Hamilton, II)**.** *Let M be a matrix over an algebraically closed field, and $\chi_M(t)$ the characteristic polynomial of M. Then $\chi_M(M) = \mathbf{0}$. That is, M satisfies its own characteristic polynomial.*

Proof. By the first version of the Cayley-Hamilton theorem, Theorem 1.5, we know that V is a direct sum of generalised eigenspaces. Suppose that G_λ is a generalised eigenspace of M with dimension d. By the definition of the characteristic polynomial, $(x-\lambda)^d$ is a divisor of the characteristic polynomial of M. We will show that $G_\lambda)$ is in the kernel of $(M-\lambda)^d$.

By Definition 1.5, every vector $v \in G_\lambda$ satisfies $(M-\lambda I)^m$ for some m depending on v. But vectors $(M-\lambda I)^i v$ for $1 \leq i \leq m$ are linearly independent by Proposition 1.11, so $m \leq d$. In particular, $(M-\lambda I)^d$ sends every vector in G_λ to 0. This argument can be applied to each generalised eigenspace in turn.

Since the first version of the Cayley-Hamilton theorem states that the generalised eigenvectors of M span V, this implies that $V \in \ker\chi(M)$, equivalently $\chi(M) = \mathbf{0}$. □

While logically equivalent to our statement, this one fundamentally obscures the main application of the Cayley-Hamilton theorem: for any $T \in \mathrm{Hom}(V,V)$ there exists a unique decomposition of V into generalised eigenspaces, G_i. Each G_i is T-invariant, and the decomposition $V = \oplus_{i=1}^t G_i$ is a direct sum. Hence, with respect to a basis of generalised eigenvectors, T can be expressed as a block-diagonal matrix,

$$\begin{bmatrix} M_1 & 0 & \dots & 0 \\ 0 & M_2 & \dots & 0 \\ \vdots & \vdots & & \vdots \\ 0 & 0 & \dots & M_t \end{bmatrix}$$

where M_i encodes the action of T on the generalised eigenspace G_i. Applying Cauchy's theorem to each M_i, these matrices are upper triangular with fixed diagonal λ_i.

Definition 1.10. *Suppose $T \in \mathcal{L}(V)$. The **minimal polynomial** of T is defined to be the unique monic polynomial $p(x)$ of smallest degree such that $p(T) = 0$.*

Proposition 1.16. *Suppose $T \in \mathcal{L}(V)$. The characteristic polynomial and minimal polynomial of T have the same zeros, which are precisely the eigenvalues of T.*

Proof. Let $q(x)$ and $p(x)$ be the characteristic polynomial and minimal polynomial of T, respectively. Since $p(x)$ divides $q(x)$ there is a polynomial $f(x)$ such that $q(x) = p(x)f(x)$. If a scalar λ is a zero of $p(x)$, then

$$q(x) = p(x)f(x) = 0 \cdot f(x) = 0$$

Thus λ is a zero of $q(x)$, and hence λ is an eigenvalue.

Conversely, suppose λ is an eigenvalue of T with corresponding eigenvector v. Then

$$0 = p(T)(v) = p(Tv) = p(\lambda v) = p(\lambda)v$$

Since $x \neq 0$, it follows that λ is a zero of $p(t)$. □

Yet another statement equivalent to the Cayley-Hamilton is that the minimal polynomial of a matrix divides its characteristic polynomial. We provide a proof of this in the appendix.

1.6. **Jordan Canonical Form.**

We conclude our analysis of a single operator $T : V \to V$ with an investigation of the action of T on a single generalised eigenspace. In fact, G_λ decomposes into a direct sum of chains: this is essentially the Jordan Canonical form.
The next theorem shows how we can select ordered bases whose union is an ordered basis B from generalized eigenspaces of a linear operator T such that $[T]_B$ is a block-diagonal matrix

$$\begin{bmatrix} A_1 & O & ... & O \\ O & A_2 & ... & O \\ \vdots & \vdots & & \vdots \\ O & O & ... & A_k \end{bmatrix}$$

called the **Jordan Canonical Form**, where each O is a zero matrix and the matrices A_i, called **Jordan blocks**, are of the form $[\lambda]$ or

$$\begin{bmatrix} \lambda & 1 & 0 & ... & 0 & 0 \\ 0 & \lambda & 1 & ... & 0 & 0 \\ \vdots & \vdots & \vdots & & \vdots & \vdots \\ 0 & 0 & 0 & ... & \lambda & 1 \\ 0 & 0 & 0 & ... & 0 & \lambda \end{bmatrix}.$$

Theorem 1.7. *Let G_λ be the generalised eigenspace of T with eigenvalue λ. There exist generalised eigenvectors $v_1, \ldots, v_d \in G_\lambda$ and integers $e_1, \ldots, e_d$ such that each chain*

$$C_i = \langle v_i, (T - \lambda I)v_i, \ldots, (T - \lambda I)^{e_i} v_i \rangle$$

is a T-invariant subspace of G_λ, and $G_\lambda = \oplus_{j=1}^{d} C_j$.

Proof. This proof is by induction on the dimension of G_λ. The base case holds trivially when G_λ has dimension 1: any non-zero vector is a basis, and there is no non-trivial condition to be satisfied. Suppose that all generalised eigenspaces of dimension $\leq r$ can be expressed as a direct sum of cyclic subspaces. For any $v \in G_\lambda$ the cyclic subspace $C_v = \langle (T-\lambda I)^t v \mid t \in \mathbb{N} \rangle$ is T-invariant (the argument is identical to the one given in Proposition 1.10).

We write M for the restriction of $T - \lambda I$ to G_λ. For a vector $v_i = v_{i,0}$ we write $M^j v_{i,0} = v_{i,j}$. Recall that the *height* of $v_{i,0}$ is the least j such that $v_{i,j} = \mathbf{0}$.

Now, suppose that G_λ has dimension $r + 1$. Since G_λ contains an eigenvector, M is neither injective nor surjective on G_λ. Hence the range of M is a proper subspace U of G_λ, of dimension $\leq k$. Applying the inductive hypothesis to U, we obtain a basis

$$u_{1,0}, \ldots, u_{1,e_1-1},\ u_{2,0}, \ldots, u_{2,e_2-1},\ \ldots,\ u_{d,0}, \ldots, u_{d,e_d-1}\,,$$

for U, where e_i is the height of u_i. Every vector $u \in U$ is of the form Mv for some vector $v \in G_\lambda$ (not necessarily unique). For each $1 \leq i \leq d$ choose a vector $v_i \in G_\lambda$ such that $Mv_i = u_{i,0}$. In particular, $v_{i,j+1} = u_{i,j}$ for any non-negative integer j.

We will show that the vectors

$$v_{1,0}, \ldots, v_{1,e_1},\ v_{2,0}, \ldots, v_{2,e_2},\ \ldots,\ v_{d,0}, \ldots, v_{d,e_d}$$

are linearly independent. Suppose that

$$\alpha_{1,0}v_{1,0} + \ldots + \alpha_{1,e_1}v_{1,e_1} + \ldots + \alpha_{d,e_d}v_{e,e_d} = \mathbf{0}\,.$$

Applying M to both sides of this equation (noting carefully that $Mv_{i,e_i} = 0$),

$$\alpha_{1,0}u_{1,0} + \ldots + \alpha_{1,e_1-1}u_{1,e_1-1} + \ldots + \alpha_{d,e_d-1}u_{d,e_d-1} = \mathbf{0}\,.$$

But the $u_{i,j}$ are linearly independent by the induction hypothesis, so $\alpha_{i,j} = 0$ for all $1 \leq i \leq d$ and $1 \leq j \leq e_i - 1$. What remains is an equation

$$\alpha_{1,e_1}v_{1,e_1} + \ldots + \alpha_{1,e_d}v_{d,e_d} = \mathbf{0}\,,$$

or equivalently, since $v_{i,e_i} = u_{i,e_i-1}$,

$$\alpha_{1,e_1}u_{1,e_1-1} + \ldots + \alpha_{1,e_d}u_{d,e_d-1} = \mathbf{0}\,,$$

Again, by the inductive hypothesis these vectors are linearly independent and so all $\alpha_{i,j}$ are 0.

A careful inspect of the proof thus far shows that we have constructed a basis for a subspace consisting of vectors which belong to $R(M)$, or have a non-zero image in $R(M)$. By the Replacement theorem, we can extend the linearly independent set $v_{i,j}$ given above to a basis of V. Any additional vectors satisfy $x_i \notin R(N)$ and $x_i \in \text{Null}(N)$. These are precisely the eigenvectors which do not belong to any cycle of dimension greater than 1. Hence a basis for G_λ is given by

$$v_{1,0}, \ldots, v_{1,e_1},\ v_{2,0}, \ldots, v_{2,e_2},\ \ldots,\ v_{d,0}, \ldots, v_{d,e_d}\ x_1, \ldots, x_\ell\,.$$

This completes the proof. □

1.7. **Rational Canonical Form.**

Definition 1.11 ([DF04]). *Let $p(x) = x^m + a_{m-1}x^{m-1} + ... + a_1x + a_0$ be any monic polynomial in $\mathbb{F}_q[X]$. The **companion matrix** of $p(x)$, denoted $\mathcal{C}_p(x)$, is the $m \times m$ matrix with 1's down the first subdiagonal, $-a_0, -a_1, ..., -a_{m-1}$ down the last column, and zeros everywhere else.*

Example 1.2. *The companion matrix of the polynomial $p(x) = x^3 + 4x^2 - 3x + 5$ is*

$$\mathcal{C}_p(x) = \begin{bmatrix} 0 & 0 & -5 \\ 1 & 0 & 3 \\ 0 & 1 & -4 \end{bmatrix}$$

Definition 1.12 ([DF04]). *A matrix is said to be in **rational canonical form** if it is the direct sum of companion matrices for monic polynomials $a_1(x), ..., a_m(x)$ of degree at least one which satisfy $a_i(x) \mid a_{i+1}(x)$ for $i = 1, ..., m-1$. The polynomials $a_i(x)$ are called the **invariant factors** of the matrix.*

Note that the RCF has the form

$$\begin{bmatrix} \mathcal{C}_{a_1}(x) & & & \\ & \mathcal{C}_{a_2}(x) & & \\ & & \ddots & \\ & & & \mathcal{C}_{a_m}(x) \end{bmatrix}$$

Example 1.3. *Consider the matrix*

$$M = \begin{bmatrix} 3 & 0 & 1 \\ 0 & 2 & 7 \\ 0 & 0 & 3 \end{bmatrix}$$

This has characteristic and minimal polynomial $(2-x)(3-x)^2 = -x^3 + 8x^2 - 21x + 18$*, so the invariant factors consist only of the characteristic polynomial. This has companion matrix* $\begin{bmatrix} 0 & 0 & -18 \\ 1 & 0 & 21 \\ 0 & 1 & -8 \end{bmatrix}$ *which is the RCF for* M*.*

1.8. **Necessary and sufficient conditions for diagonalizability.**

In this Chapter, we have reviewed the structure of a linear operator in detail. In the Cayley-Hamilton theorem, we showed that a linear operator T on vector space V induces a decomposition of V into a direct sum of generalised eigenspaces, while in the Jordan Canonical Form theorem, we gave a description of the action of T on any generalised eigenspace. Together, these results provide a complete understanding of a single linear operator. In this thesis, we are interested in necessary and sufficient conditions for diagonalizability. We outline some of these here.

Theorem 1.8. *Suppose* $T \in \mathcal{L}(V)$*. Let* $\lambda_1, ..., \lambda_k$ *be the distinct eigenvalues of* T*. Then* T *is diagonalizable if and only if the minimal polynomial of* T *is of the form*

$$p(x) = \prod_{i=1}^{k} (x - \lambda_i)$$

Proof. Suppose T is diagonalizable, and hence let $B = \{v_1, ..., v_n\}$ be an ordered basis for V consisting of eigenvectors of T. Define

$$p(x) = \prod_{i=1}^{k} (x - \lambda_i)$$

Consider eigenvector $v_i \in B$ with corresponding eigenvalue λ_j. Then we have $(T - \lambda_j I)(v_i) = 0$. Since $(x - \lambda_j)$ divides $p(x)$, there is a polynomial $f(x)$ for which $p(x) = f(x)(x - \lambda_j)$. Thus

$$p(T)(v_i) = f(T)(T - \lambda_j I)(v_i) = 0$$

Then $p(T) = 0$, since its null space contains a basis for V. Furthermore, the factorization of the minimal polynomial has at least as many linear factors as distinct eigenvalues (Proposition 1.16), and hence has degree at least k. Thus $p(x)$ is the minimal polynomial.

Conversely, suppose $\lambda_1, ..., \lambda_k$ are distinct scalars such that the minimal polynomial $p(x)$ factors as

$$p(x) = \prod_{i=1}^{k}(x - \lambda_i)$$

By Proposition 1.16, the λ_is are eigenvalues of T. We apply induction on $n = \dim(V)$. Clearly, T is diagonalizable when $n = 1$. Now assume T is diagonalizable when $\dim(V) < n$ for some $n > 1$. Let $W = R(T - \lambda_k I)$ and $U = N(T - \lambda_k I)$. Obviously $W \neq V$, since v_k is an eigenvalue of T. If $W = \{0\}$, then $T = \lambda_k I$ is clearly diagonalizable. So suppose $0 < \dim(W) < n$. Then W is T-invariant, and for any $x \in W$,

$$\left(\prod_{i=1}^{k-1}(T - \lambda_i I)\right)(x) = 0.$$

It follows that the minimal polynomial of $T|_W$ divides $\prod_{i=1}^{k-1}(x - \lambda_i)$. Hence by the induction hypothesis, $T|_W$ is diagonalizable. Furthermore, λ_k is not an eigenvalue of $T|_W$ (by Proposition 1.16). Therefore $W \cap U = \{0\}$. Now consider the disjoint bases $B_1 = \{v_1, .., v_m\}$ for W and $B_2 = \{w_1, ..., w_p\}$ for U. By the rank-nullity theorem, $m + p = n$. We show that $B = B_1 \cup B_2$ is linearly independent. Consider scalars $a_1, .., a_m$ and $b_1, ..., b_p$ such that $x + y = 0$ for

$$x = \sum_{i=1}^{m} a_i v_i \in W \quad \text{and} \quad y = \sum_{i=1}^{p} b_i w_i \in U.$$

It follows that $x = -y \in W \cap U$, thus $x = 0$. Given that each basis is linearly independent, each scalar a_i and b_i is 0, and we conclude that B is a linearly independent subset of V consisting of n eigenvectors. Thus B is a basis for V consisting of eigenvectors of T, and hence T is diagonalizable. □

Corollary 1.3. *If the eigenvalues of T are distinct then T is diagonalizable.*

Proof. Clearly, if the eigenvalues of T are distinct then each has algebraic multiplicity 1, and the characteristic polynomial of T is of the form

$$p(x) = \prod_{i=1}^{k}(x - \lambda_i).$$

That is, the minimal polynomial is square-free. □

A matrix with non-trivial generalised eigenvectors is never diagonalizable, though it can be difficult to decide whether a matrix has such a generalised eigenvector.

Theorem 1.9. *[Art10] Let T be a linear operator on a finite-dimensional complex vector space. The following conditions are equivalent:*

(1) *T is diagonalizable,*
(2) *every generalized eigenvector is an eigenvector,*
(3) *all of the blocks in the JFC of T are 1×1 blocks.*

Proof. $(1 \implies 2)$ Suppose that T is diagonalizable, say T with respect to basis $B = (v_1, .., v_n)$ is the diagonal matrix D with diagonal entries $\lambda_1, ..., \lambda_n$. Let v be a generalized eigenvector in V such that $(T - \lambda_i)^k v = 0$ for some λ_i and some $k > 0$. We write $S = T - \lambda_i I_n$ to reduce to the case $S^k v = 0$. Let $[v]_B = (x_1, ..., x_n)^t$. The coordinates of $S^k v$ will be $\lambda_i^k x_i$. Since $S^k v = 0$, either $\lambda_i = 0$ or $x_i = 0$, and in either case $\lambda_i^k x_i = 0$. Thus $Sv = 0$.
$(2 \implies 3)$ We prove the contrapositive. Recall that the $k \times k$ Jordan block J_0 operates on the standard basis of $\mathbb{C}^k$ as

$$e_1 \to e_2 \to ... \to e_k \to 0.$$

If the JFC of T has a $k \times k$ Jordan block with $k > 1$, then looking at the action of $J_\lambda - \lambda I$, we see that there is a generalized eigenvector that is not an eigenvector.
Finally, it is clear that $(3 \implies 1)$. □

In the next chapter, we will focus on the theory of operators on an inner product space, with the goal of showing that a symmetric matrix over a field of characteristic 0 is diagonalizable. We will see in chapter 4 that this proof does not generalise to positive characteristic.

2. Inner Product Spaces and the Spectral Theorem in characteristic 0

In this chapter we introduce the spectral theorem, by which we know that symmetric matrices are diagonalizable in characteristic 0. We will explore positive characteristic in future chapters. Material in this chapter is drawn from [Rom08].

Definition 2.1. *Let V be a vector space over $\mathbb{R}$ or $\mathbb{C}$. A mapping $\langle , \rangle : V \times V \to \mathbb{F}$ is called an **inner product** if it has the following properties:*

1. *Linearity in the first argument: $\langle u + v, w\rangle = \langle u, w\rangle + \langle v, w\rangle$ and $\langle \alpha u, v\rangle = \alpha\langle u, v\rangle$.*
2. *Conjugate symmetry: $\langle u, v\rangle = \overline{\langle v, u\rangle}$.*
3. *Positive definiteness: $\langle v, v\rangle \geq 0$ for all $v \in V$ and $\langle v, v\rangle = 0$ only if $v = 0$.*

*A vector space equipped with an inner product is called an **inner product space** (or a **metric vector space**).*

Example 2.1. *Let $V = \mathbb{F}^n$. For any $A \in M_n(\mathbb{F})$, the map $\langle , \rangle : V \times V \to \mathbb{F}$ defined as the dot product*

$$\langle x, y\rangle = x^T y \quad \text{for } x, y \in V$$

is an inner product.

We review some terminology relating to adjoints next.

2.1. Adjoints, Self-Adjoint Matrices.

Definition 2.2. *Let T be a linear operator on an inner product space V over $\mathbb{R}$ or $\mathbb{C}$. The **adjoint** of T is the unique linear operator T^* on V defined by the condition*

$$\langle Tu, v\rangle = \langle u, T^* v\rangle$$

for all $u, v \in V$.
*Furthermore, T is called **normal** if $TT^* = T^*T$, and **self-adjoint** if $T = T^*$.*

Proposition 2.1. *The eigenvalues of a self-adjoint operator over $\mathbb{R}$ or $\mathbb{C}$ are real.*

Proof. Let $v \in V$ be an eigenvalue of T with eigenvalue λ. Then

$$\lambda\langle v, v\rangle = \langle Tv, v\rangle = \langle v, Tv\rangle = \langle v, \lambda v\rangle = \lambda^*\langle v, v\rangle.$$

Since $\langle v, v\rangle \neq 0$, it follows that $\lambda = \lambda^*$. Hence λ is real. □

Proposition 2.2. *A self-adjoint operator over $\mathbb{R}$ or $\mathbb{C}$ is diagonalizable.*

Proof. Let T be a self-adjoint operator and v a generalized eigenvector of T. By the previous proposition the corresponding eigenvalue is real. Then $(T - \lambda I)$ is also self-adjoint, since

$$(T - \lambda I)^* = T^* - \lambda^* I = T - \lambda I.$$

Let k be the smallest integer for which $(T - \lambda I)^k v = 0$. Then $w = (T - \lambda I)^{k-1} v$ is an eigenvector of T with corresponding eigenvalue λ. Now consider the inner product

$$\langle w, w\rangle = \langle (T - \lambda I)^{k-1} v, (T - \lambda I)^{k-1} v\rangle = \langle v, (T - \lambda I)^{2k-2} v\rangle.$$

If $2k-2 \geq k$ then $(T-\lambda I)^{2k-2}v = 0$, which would contradict the positive-definite property of inner products. Hence $2k-2 < k$, so $k = 1$. Then $(T-\lambda I)v = 0$ and $w = v$. We conclude that every generalized eigenvector is already an eigenvector. $\square$

Definition 2.3. *Let V be an inner product space.*

- *A vector $x \in V$ is **orthogonal** to a vector y, denoted $x \perp y$, if $\langle x, y\rangle = 0$.*
- *A vector $x \in V$ is **orthogonal** to a subset $S \subseteq V$, denoted $x \perp S$, if $\langle x, s\rangle = 0$ for all $s \in S$.*

Definition 2.4. *Let T be a linear operator on an inner product space V. If $\|Tv\| = \|v\|$ for all $v \in V$, T is called **orthogonal** if the underlying field is $\mathbb{R}$ and **unitary** if it is $\mathbb{C}$. Furthermore, a matrix M is called **orthogonal** if $M^TM = I_n$ and **unitary** if $M^*M = I_n$.*

Definition 2.5. *Let V and W be inner product spaces. A bijective linear operator $T : V \to W$ is called an **isometry** if*

$$\langle Tu, Tv\rangle = \langle u, v\rangle$$

for all $u, v \in V$.

The set of isometries of V onto V form a group, with function composition as its group operation. Thus, if T and S are isometries of V onto V, then so is TS.

Proposition 2.3. *Let T be a linear operator on an inner produce space V. The following statements are equivalent.*

(1) *T is unitary or orthogonal.*
(2) *T is an isometry.*
(3) *$TT^* = T^*T = I$.*

Proof. ($1 \implies 2$) For any $x \in V$, we have

$$\langle Tx, Tx\rangle = \|Tx\|^2 = \|x\|^2 = \langle x, x\rangle$$

($2 \implies 3$) For any $x \in V$, we have

$$\langle x, x\rangle = \langle Tx, Tx\rangle = \langle x, T^*Tx\rangle = 0$$

Thus $\langle x, (I - T^*T)x\rangle = 0$ for all $x \in V$. This implies that $I - T^*T = T_0$, and therefore $T^*T = I$. Since $\langle x, T^*Tx\rangle = \langle TT^*x, x\rangle$, applying similar steps gives us $TT^* = I$.
($3 \implies 1$) For any $x \in V$, we have

$$\|x\|^2 = \langle x, x\rangle = \langle x, T^*Tx\rangle = \langle TT^*x, x\rangle = \langle Tx, Tx\rangle = \|Tx\|^2.$$

$\square$

Proposition 2.4. *The matrix $M \in M_n(\mathbb{R})$ is an isometry if and only if the columns of M form an orthonormal basis of $\mathbb{R}^n$.*

Proof. By definition M is orthogonal if and only if $MM^\top = I_n$. But the (i, j) entry of $MM^\top$ is the inner product of the i^{th} and j^{th} columns of M. So the definition requires that the columns form an orthonormal basis. $\square$

Proposition 2.5. *Let V be a vector space over a field of characteristic 0. Then any symmetric bilinear form on V is diagonalizable by an orthogonal matrix.*

Proof. Since a symmetric matrix is diagonalizable, it is self-adjoint and so diagonalizable by Proposition 2.2. In particular, there are no generalised eigenvectors, and each eigenspace admits an orthonormal basis of eigenvectors. So it suffices to show that eigenvectors from distinct eigenspaces are orthogonal. To see this, suppose that $Mv = \lambda v$ and that $Mu = \mu u$ where $\lambda \neq \mu$ and M is self-adjoint.

Then

$$\lambda\langle v, u\rangle = \langle Mv, u\rangle = \langle v, Mu\rangle = \mu\langle v, u\rangle .$$

But $\lambda \neq \mu$ then forces $\langle v, u\rangle = 0$ and eigenvectors from distinct eigenspaces are orthogonal. As a result, M admits a basis of eigenvectors, and is diagonalized by an orthogonal matrix. □

Theorem (The Spectral Theorem). *Let T be a linear operator on an inner product space V over $\mathbb{R}$ or $\mathbb{C}$. The following are equivalent:*

- *T is normal.*
- *$\|Tv\| = \|T^*v\|$ for all $v \in V$.*
- *Every eigenvector of T is an eigenvector of T^*. Every generalised eigenvector of T is an eigenvector, and eigenvectors with distinct eigenvalues are orthogonal.*
- *T is diagonalizable by an isometry.*

Proof. ($1 \implies 2$) Suppose T is normal. Then

$$\|Tv\|^2 = \langle Tv, Tv\rangle = \langle v, T^*Tv\rangle = \langle v, TT^*v\rangle = \langle T^*v, T^*v\rangle = \|T^*v\|^2$$

T being normal gives us the third equality, and that $T^{**} = T$ the fourth.
($2 \implies 3$) Let v be an eigenvector of T with corresponding eigenvalue λ. Observing that $T - \lambda I$ is normal if and only if T is, we have

$$0 = |(T - \lambda I)v| = |(T - \lambda I)^*v| = |(T^* - \lambda^* I)v| = 0$$

So v is an eigenvector of T with eigenvalue λ^*.

Suppose that v is a generalised eigenvector of T with eigenvalue λ, so that $(T-\lambda I)^k v = 0$, with k being the smallest integer for which this holds. Let $S = (T - \lambda I)(T^* - \lambda^* I)$. Then $S^* = (T^* - \lambda^* I)(T - \lambda I)^* = S$, so S is self-adjoint. Since T is normal,

$$S^k v = (T^* - \lambda^* I)^k (T - \lambda I)^k v = 0,$$

for any integer k. Recall that by Proposition 2.2, self-adjoint operators over $\mathbb{R}$ or $\mathbb{C}$ are diagonalizable. Thus S is diagonalizable and so $N(S) = N(S^k)$ for all $k \geq 1$. Observe that

$$\langle Sv, v\rangle = \langle (T - \lambda I)v, (T - \lambda I)v\rangle,$$

which is nonzero unless $k = 1$ and v is an eigenvector of T. Thus $Sv \neq 0$ for $k > 1$.

Similarly, $S^d v = 0$ if and only if $(T - \lambda I)^d v = 0$. We consider the expression

$$\langle S^{k-1}v, S^{k-1}v\rangle = \langle v, S^{2k-2}v\rangle.$$

By hypothesis, the left-hand side of the equation is nonzero, so $2k - 2 < k$ which forces $k = 1$. Hence every generalised eigenvector of T is an eigenvector of T. Since T admits

a basis of generalized eigenvectors by the Cayley-Hamilton theorem, and every generalised eigenvector is an eigenvector, T is diagonalizable.

Finally, let v_1 and v_2 be eigenvectors of T with distinct eigenvalues λ_1 and $lambda_2$. Then

$$\lambda_1\langle v_1, v_2\rangle = \langle Tv_1, v_2\rangle = \langle v_1, T^*v_2\rangle = \lambda_2\langle v_1, v_2\rangle.$$

Hence $\langle v_1, v_2\rangle = 0$.
$(3 \implies 4)$ By the previous section, T admits an orthonormal basis of eigenvectors. We assemble an orthonormal basis of eigenvectors of T by constructing an orthonormal basis for each eigenspace using the Gram-Schmidt process. The union of these bases gives an orthonormal basis for the space, with respect to which T is diagonal.
$(4 \implies 1)$ Suppose T is diagonalizable by a matrix M having orthonormal columns. Hence $M^*M = I_n$. So M^*TM is a diagonal matrix, and $(M^*TM)^* = M^*T^*M$ is the adjoint of T with respect to this basis. Then

$$(M^*TM)(M^*T^*M) = M^*TT^*M.$$

Since diagonal matrices commute, we also have

$$(M^*TM)(M^*T^*M) = (M^*T^*M)(M^*TM) = M^*T^*TM.$$

Cancelling invertible matrices, we conclude that $TT^* = T^*T$ and so T is normal. □

The spectral theorem is a powerful result for giving conditions about diagonalizability. It tells us that symmetric, orthogonal, and self-adjoint matrices are all diagonalizable when the base field has characteristic 0. Next we will investigate what happens over finite fields, in the simplest nontrivial case of 2×2 matrices.

3. Diagonalizability of 2x2 Matrices Over a Finite Field

3.1. General 2x2 Matrices.

In this section we explore the case of 2x2 matrices over finite fields to see how it differs from the characteristic zero case. This will require more abstract algebra than previous chapters. We classify four types of 2x2 matrices over a field $\mathbb{F}_q$, and provide counts for each up to change of basis. Our cases derive from the specification of a matrix as being diagonalizable or non-diagonalizable, and having repeated or distinct eigenvalues. We use the following notation:

	Diagonalizable	Non-diagonalizable
Repeated eigenvalues	R-D	R-ND
Non-repeated eigenvalues	NR-D	NR-ND

It will be convenient to impose an ordering on the elements of $\mathbb{F}_p$ and to insist that the eigenvalues of A are always listed in increasing order.

First, we introduce some concepts from group theory.

Definition 3.1. *Let G be a group and X a finite set. Then a function $\phi : G \times X \to X$ is an **action** (of G on X) if and only if the following conditions are satisfied:*

(1) $\phi(x, 1_G) = x$ *for all* $x \in X$.
(2) $\phi(\phi(x, g), h) = \phi(x, gh)$ *for all* $x \in X$ *and all* $g, h \in G$.

An action of G on a set X is essentially the same thing as a homomorphism from G into the symmetric group on X. In these notes we will mostly be interested in the conjugation action of the general linear group $GL_2(\mathbb{F}_q)$ on the set of 2×2 matrices over $\mathbb{F}_q$.

Definition 3.2. *Let $\phi : G \times X \to X$ be a group action. For any $\alpha \in X$, we call*

$$stab_G(\alpha) = \{g \in G \mid g(\alpha) = \alpha\}$$

*the **stabilizer** of α in G. Similarly,*

$$orb_G(\alpha) = \{g(\alpha) \mid g \in G\}$$

*the **orbit** of α under G.*

The next result is one of the more useful theorems in finite group theory.

The Orbit-Stabilizer Theorem ([Gal86]). *Let G be a finite permutation group acting on a set S. Then, for any $\alpha \in S$, $|G| = |orb_G(\alpha)||stab_G(\alpha)|$.*

Before we begin our computations with matrices of the four cases, we recall some properties of the general linear group.

Definition 3.3. *A matrix is **invertible** if all of its eigenvalues are non-zero. The product of invertible matrices is invertible, so that the invertible $n \times n$ matrices over a field $\mathbb{F}$ form a group. This is the **general linear group** of dimension n over $\mathbb{F}$, normally denoted $\mathrm{GL}_n(\mathbb{F})$.*

Proposition 3.1. *The group* $\mathrm{GL}_2(\mathbb{F}_p)$ *has size* $(p^2-1)(p^2-p)$.

Proof. Recall from elementary linear algebra that matrix is invertible if and only if the rows of the matrix are linearly independent. Furthermore, two matrices are distinct if and only if their rows are distinct (as an ordered set). Let us choose two linearly independent vectors from a two-dimensional vector space over $\mathbb{Z}_p$ in all possible ways.

The first row of the matrix can be any non-zero vector, so there are p^2-1 choices. The second row must not be a scalar multiple of the first row, so there are p^2-p choices for the second row. Hence the size of $\mathrm{GL}_2(\mathbb{F}_p)$ is $(p^2-1)(p^2-p)$. □

We will apply the Orbit-Stabilizer theorem to matrices from each of the four types in turn. Before we begin, let us specify the group action.

Definition 3.4. *The* ***conjugation action*** *of the general linear group* $\mathrm{GL}_n(\mathbb{Z}_p)$ *on the set of* $n \times n$ *matrices is given by*

$$\phi(M, A) = M^{-1}AM\,,$$

where $M \in \mathrm{GL}_n(\mathbb{Z}_p)$ *and* $A \in M_n(\mathbb{Z}_p)$.

It can be easily verified that the conjugation action **is** a group action, as follows:

$$\begin{aligned}\phi(M, \phi(N, A)) &= \phi(M, N^{-1}AN) = M^{-1}N^{-1}ANM \\ &= (MN)^{-1}A(MN) = \phi(MN, A).\end{aligned}$$

The verification that the identity matrix acts trivially on all matrices M is routine, and so is omitted. It is well known that conjugation by an invertible matrix corresponds to a change of basis operation which preserves the eigenvalues and Jordan Canonical Form of a matrix. Hence it is the natural equivalence operation to consider when representing a linear transformation by a matrix.

3.1.1. *R-D Matrices.*

Every R-D matrix is of the form

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} \alpha & 0 \\ 0 & \alpha \end{bmatrix} \begin{bmatrix} a & b \\ c & d \end{bmatrix}^{-1} \text{ where } a, b, c, d, \alpha, \beta \in \mathbb{F}_q.$$

We will count the R-D matrices using the Orbit-Stabilizer theorem.

Proposition 3.2. *The number of R-D matrices in* $M_2(\mathbb{F}_q)$ *is* q. *(This includes the zero matrix.)*

Proof. Observe that since A is a scalar matrix, it commutes with all other matrices in the general linear group:

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} \alpha & 0 \\ 0 & \alpha \end{bmatrix} = \begin{bmatrix} \alpha & 0 \\ 0 & \alpha \end{bmatrix} \begin{bmatrix} a & b \\ c & d \end{bmatrix}$$

Hence, the stabilizer of A is all of $GL_2(q)$ and the orbit of A is just $\{A\}$. There are q R-D matrices, one for each element of the field $\mathbb{F}_q$. □

3.1.2. *NR-D Matrices.*

Next, we consider the diagonalizable matrices with distinct eigenvalues. We will assume an ordering on the field elements, so that all Jordan Canonical Forms are unique.

Every NR-D matrix is of the form

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} \alpha & 0 \\ 0 & \beta \end{bmatrix} \begin{bmatrix} a & b \\ c & d \end{bmatrix}^{-1} \text{ where } a, b, c, d, \alpha, \beta \in \mathbb{F}_q.$$

Proposition 3.3. *The number of NR-D matrices is* $\frac{q^2(q^2-1)}{2}$.

Proof. Again, we compute the the size of the stabilizer of a matrix with distinct eigenvalues. To avoid working with matrix inverses, we observe that $M^{-1}AM = A$ if and only if $AM = MA$, since the matrix M is assumed to be invertible. We will solve the system of linear equations which come from this matrix equation.

$$\begin{bmatrix} \alpha & 0 \\ 0 & \beta \end{bmatrix} \begin{bmatrix} a & b \\ c & d \end{bmatrix} = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} \alpha & 0 \\ 0 & \beta \end{bmatrix}$$

We can rewrite this matrix equation as a set of linear equations, obtained by setting the corresponding matrix entries equal:

$$\begin{aligned} \alpha a - \alpha a &= 0 \\ \beta b - \alpha b &= 0 \\ \alpha c - \beta c &= 0 \\ \beta d - \beta d &= 0 \end{aligned}$$

Clearly the first and last of these equations are trivial. We are assuming that $\alpha \neq \beta$, so since we work over a field, we must have $b = 0$ and $c = 0$. Hence the stabilizer of the matrix A under the conjugation action is the group

$$G_A = \{M \mid MAM^{-1} = A\} = \left\{ \begin{bmatrix} a & 0 \\ 0 & d \end{bmatrix} \;\middle|\; a, d \neq 0 \right\}$$

which has size $(p-1)^2$. Now we apply the Orbit-Stabilizer Theorem: the orbit of a NR-D matrix must be of size

$$\frac{q(q-1)^2(q+1)}{(q-1)^2} = q^2 + q\,.$$

Now, the number of diagonal matrices with non-repeated eigenvalues is $\binom{q}{2}$. So the total number of NR-D matrices is

$$\frac{(q+1)q^2(q-1)}{2} = \frac{q^2(q^2-1)}{2}\,.$$

□

3.1.3. *R-ND Matrices.*

Proposition 3.4. *The number of R-ND matrices in* $M_2(\mathbb{F}_q)$ *is* $q^3 - q$.

Proof. Let

$$M = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \quad \text{and} \quad A = \begin{bmatrix} \alpha & 1 \\ 0 & \alpha \end{bmatrix}.$$

Every R-ND matrix is of the form

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} \alpha & 1 \\ 0 & \alpha \end{bmatrix} \begin{bmatrix} a & b \\ c & d \end{bmatrix}^{-1} \text{ where } a, b, c, d, \alpha, \in \mathbb{F}_q.$$

We want

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} \alpha & 1 \\ 0 & \alpha \end{bmatrix} = \begin{bmatrix} \alpha & 1 \\ 0 & \alpha \end{bmatrix} \begin{bmatrix} a & b \\ c & d \end{bmatrix}$$

We can rewrite this matrix equation as a set of linear equations, obtained bysetting the corresponding matrix entries equal:

$$\begin{aligned} a\alpha &= a\alpha + c \\ a + b\alpha &= b\alpha + d \\ c\alpha &= c\alpha \\ c + d\alpha &= d\alpha \end{aligned}$$

Clearly we must have $c = 0$ and $a = d$. Thus if M commutes with A then M is of the form

$$\begin{bmatrix} a & b \\ 0 & a \end{bmatrix}, \quad a \neq 0.$$

Hence the stabilizer under the conjugation group action is

$$G_A = \{M \mid MAM^{-1} = A\} = \left\{ \begin{bmatrix} a & b \\ 0 & a \end{bmatrix} \,\middle|\, a \neq 0 \right\}.$$

which has size $q(q-1)$. Applying the Orbit-Stabilizer theorem, the orbit of a R-ND matrix must be of size

$$\frac{|GL_2(q)|}{|G_A|} = \frac{q(q-1)^2(q+1)}{q(q-1)} = (q-1)(q+1) = q^2 - 1.$$

Hence there are precisely $q^2 - 1$ matrices in $M_2(\mathbb{F}_q)$ with the same JFC as A. With q choices for α, there are $q^3 - q$ R-ND matrices in $M_2(\mathbb{F}_q)$. $\square$

3.1.4. *NR-ND Matrices.*

For NR-ND matrices, we use the rational canonical form.

Proposition 3.5. *The number of NR-ND matrices in* $M_2(\mathbb{F}_q)$ *is* $q^2(q-1)^2/2$.

Proof. Recall that a matrix is diagonalizable if and only if its minimal polynomial can be expressed as a product of linear factors. Thus an NR-ND matrix has an irreducible minimal polynomial. Since the minimal and characteristic polynomials have the same zeroes, the characteristic polynomial of an NR-ND matrix is also irreducible.
Say that $x^2 - \beta x + \gamma$ an irreducible quadratic. Observe that the matrix with characteristic polynomial $\chi(A) = x^2 - \beta x + \gamma$ is

$$A = \begin{bmatrix} 0 & -\gamma \\ 1 & \beta \end{bmatrix}.$$

We want

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} 0 & -\gamma \\ 1 & \beta \end{bmatrix} = \begin{bmatrix} 0 & -\gamma \\ 1 & \beta \end{bmatrix} \begin{bmatrix} a & b \\ c & d \end{bmatrix}$$

Again, we rewrite this as a system of linear equations.

$$\begin{aligned} b &= -c\gamma \\ d &= a + \beta c \\ -a\gamma + b\beta &= -d\gamma \\ -c\gamma + d\beta &= b + \beta d \end{aligned}$$

We see that an element of the stabilizer is of the form

$$\begin{bmatrix} a & -c\gamma \\ c & a + \beta c \end{bmatrix}$$

This matrix is invertible with characteristic polynomial $a^2 + a\beta c + c^2\gamma$. This polynomial has a root if and only if the discriminant $(\beta c)^2 - 4c^2\gamma = c^2(\beta^2 - 4\gamma)$ is a square. Since we assumed $x^2 - \beta x + \gamma$ is irreducible, there is no square root of $\beta^2 - 4\gamma$ in the field. So we just need that a and c are non-zero. Thus stabilizer is

$$G_A = \left\{ \begin{bmatrix} a & -c\gamma \\ c & a + \beta c \end{bmatrix} \,\middle|\, a, c \neq 0 \right\}$$

and has size $|G_A| = q^2 - 1$.
Applying the Orbit-Stabilizer theorem, the orbit of a NR-ND matrix must be of size

$$\frac{|GL_2(q)|}{|G_A|} = \frac{(q^2-1)(q^2-q)}{q^2-1} = q^2 - q$$

So the number of matrices with an irreducible minimal polynomial is

$$\binom{q}{2}(q^2 - q) = \frac{q^2(q-1)^2}{2}.$$

□

We can now state the final counts for the four types of matrices.

Proposition 3.6. *The counts for the matrices in $M_2(\mathbb{F}_q)$ are*

	R	*NR*
D	q	$q^2(q^2-1)/2$
ND	q^3-q	$q^2(q-1)^2/2$

We make some observations here: the number of matrices with a repeated eigenvalue is q^3, almost all of these matrices fail to be diagonalizable. It is also interesting to note that the total number of diagonalizable matrices among all 2×2 matrices is $\frac{1}{2}q^4 - \frac{1}{2}q^2 + q$ for any prime power q. In particular, as $q \to \infty$ the probability that a randomly chosen 2×2 matrix is diagonalizable tends to $\frac{1}{2}$. In the next section we introduce some specialised results on finite fields, before considering 2×2 symmetric matrices in positive characteristic.

3.2. **Cyclotomy.**

Here we review some results on cyclotomic numbers. This algebra is needed to work through our analysis of symmetric matrices.

Definition 3.5. *Let R be a commutative ring. The set of formal symbols*

$$R[x] = \{a_n x^n + a_{n-1}x^{n-1} + ... + a_1 x + a_0 \mid a_i \in R, n \text{ is a nonnegative integer}\}$$

*Is called the **ring of polynomials over** R **in the indeterminate x**.*
Two elements

$$a_n x^n + a_{n-1}x^{n-1} + ... + a_1 x + a_0$$

and

$$b_n x^n + b_{n-1}x^{n-1} + ... + b_1 x + b_0$$

of $R[x]$ are considered equal if and only if $a_i = b_i$ for all nonnegative integers i. (Define $a_i = 0$ when $i > n$ and $b_i = 0$ when $i > m$.)

Definition 3.6. *Recall the quadratic formula gives the roots of the polynomial $p(x) = x^2 + rx + s$ as*

$$\frac{-r \pm \sqrt{r^2 - 4s}}{2}.$$

*The **discriminant** is the term beneath the radical.*

The discriminant has the following properties:

(1) The discriminant is zero if and only if the roots of $p(x)$ are equal.
(2) The polynomial splits into linear factors if and only if the discriminant is a square.
(3) The polynomial is irreducible if and only if the discriminant is not a square.

Definition 3.7. *Let q be an odd prime power. The **quadratic residues** of $\mathbb{F}_q$ are the elements of the unique subgroup of index 2 in $\mathbb{F}_q^*$. The **quadratic non-residues** are the unique coset of the quadratic residues in $\mathbb{F}_q^*$.*

Proposition 3.7. *The quadratic equation $x^2 - bx + c$ is solvable in $\mathbb{F}_q$ if and only if the discriminant $b^2 - 4c$ is a quadratic residue in $\mathbb{F}_q$ (i.e. has a square root).*

Proposition 3.8. *There are $(q-1)/2$ elements in $\mathbb{F}_q$ that are quadratic residues and $(q-1)/2$ elements in $\mathbb{F}_q$ that are quadratic non-residues.*

Proposition 3.9. *If $q \equiv 1 \mod 4$ then x is a QR if and only if $-x$ is a QR. If $q \equiv 3 \mod 4$ then x is a QR if and only if x is **not** a QR.*

Proof. Let g be a generator of the multiplicative group of the finite field. Then g has order $q-1$, which is $2 \mod 4$ if $q \equiv 3 \mod 4$ and $0 \mod 4$ otherwise. Clearly g^k is a Quadratic Residue if and only if k is even, and the QRs form a subgroup of $\mathbb{F}_q^*$.

Observe that $\left(g^k\right)^{p-1/2} = 1$ if k is even (by Lagrange's theorem) and -1 if k is odd (because there are at most $p - 1/2$ solutions to $x^{p-1/2} - 1$ in a field, but squaring must produce 1 by Lagrange). So -1 is a quadratic residue precisely when $q - 1/2$ is even; that is when $q \equiv 1 \mod 4$. Now the result follows from closure of the quadratic residues. □

Proposition 3.10. *Suppose $gcd(q, 2) = 1$. Then there are $(q-1)/2$ values of $b^2 - 4c$ which are non-residues for each value of b. Thus there are $q(q-1)/2$ irreducible quadratics over $\mathbb{F}_q$.*

Let g be a generator for $\mathbb{F}_q^*$. The e^{th} powers in $\mathbb{F}^q$ are the subgroup

$$C_{e,0} = \langle g^e \rangle .$$

More generally, the **cyclotomic classes** of order e are the cosets of $C_{e,0}$, given as $C_{e,t} = g^t C_{e,0}$.

Definition 3.8. *The **cyclotomic numbers** of order e are defined as follows.*

$$(i, j)_e = |\{C_{e,i} + 1\} \cap C_{e,j}|$$

That is, $(i, j)_e$ is the number of solutions to the equation $x^2 + 1 = y^2$ where $x \in C_{e,i}$ and $y \in C_{e,j}$.

Proposition 3.11. *The following identities hold.*

(1) $(i, j)_e = (-i, j - i)_e$.
(2) $\sum_j (i, j)_e = \frac{q-1}{e} - n_i$ *where* $n_0 = 1$ *when* $\frac{q-1}{e}$ *is even and* $n_{e/2} = 1$ *when this quantity is odd, with* $n_i = 0$ *otherwise.*
(3) $(j, i)_e = (i, j)_e$ *if* $\frac{q-1}{e}$ *is even, and* $(i + e/2, j + e/2)$ *if* $\frac{q-1}{e}$ *is odd.*

Proof. Found in [Hal86]. □

From Proposition 3.11 we establish the cyclotomic numbers of order 2.

Proposition 3.12. *Suppose that $q \equiv 3 \mod 4$, say $q = 4t+3$. Then the cyclotomic numbers of order 2 are as follows.*

$$(0, 0)_2 = (1, 0)_2 = (1, 1)_2 = t, \;\; (0, 1)_2 = t + 1 .$$

If $q \equiv 1 \mod 4$ and $q = 4t + 1$ then

$$(0, 0)_2 = t - 1, \;\; (1, 1)_2 = (1, 0)_2 = (0, 1)_2 = t .$$

Proof. Suppose that $q = 4t + 3$. Then $\frac{q-1}{e}$ is odd and by we have

$$(0,0)_2 + (0,1)_2 = \frac{q-1}{e} - n_0 = \frac{4t+3-1}{2} - 0 = 2t+1$$
$$(1,0)_2 + (1,1)_2 = \frac{q-1}{e} - n_1 = \frac{4t+3-1}{2} - 1 = 2t$$
$$(0,0)_2 = (1,0)_2 = (1,1)_2 = t,\ (0,1)_2 = t+1.$$

Now suppose $q = 4t + 1$. Then $\frac{q-1}{e}$ is even and by we have

$$(0,0)_2 + (0,1)_2 = \frac{q-1}{e} - n_0 = \frac{4t+1-1}{2} - 1 = 2t-1$$
$$(1,0)_2 + (1,1)_2 = \frac{q-1}{e} - n_1 = \frac{4t+1-1}{2} - 0 = 2t$$
$$(0,0)_2 = t-1,\ (1,1)_2 = (1,0)_2 = (0,1)_2 = t.$$

□

3.3. **2x2 Symmetric Matrices.**
Here we consider matrices of the form

$$M = \begin{bmatrix} a & b \\ b & d \end{bmatrix}$$

over $\mathbb{F}_q$, giving counts for diagonalizable and non-diagonalizable symmetric matrices with non-repeated eigenvalues. There are q^3 such matrices.
The characteristic polynomial of M is

$$\chi(M) = \lambda^2 - (a+d)\lambda + ad - b^2.$$

By a little algebra, the discriminant of $\chi(M)$ is

$$\Delta(M) = (a+d)^2 - 4(ad - b^2) = (a-d)^2 + (2b)^2.$$

The case where $b = 0$ is easy: there are q^2 matrices in total, q of which have a repeated eigenvalue.

Proposition 3.13. *Suppose that $b \neq 0$, that $a - d \neq 0$ and that $\Delta(M) \neq 0$. Under these conditions, the number of diagonalizable matrices is as follows.*

(1) *If $q \equiv 1 \mod 4$, there are $q(q-1)(q-3)$ matrices which satisfy the hypotheses. Of these $\frac{1}{2}q(q-1)(q-5)$ diagonalizable matrices, and $\frac{1}{2}q(q-1)(q-1)$ non-diagonalizable matrices.*
(2) *If $q \equiv 3 \mod 4$ there are $q(q-1)^2$ matrices which satisfy the hypotheses. Of these $\frac{1}{2}q(q-1)(q-3)$ diagonalizable matrices and $\frac{1}{2}q(q-1)(q+1)$ non-diagonalizable matrices.*

Proof. By hypothesis $a-d \neq 0$ and $b \neq 0$ so there are at most $q(q-1)^2$ matrices to consider. Since $b \neq 0$ we can rewrite the discriminant as

$$\left(\frac{a-d}{2b}\right)^2 + 1 = \Delta(M)\left(\frac{1}{2b}\right)^2.$$

If $q \equiv 3 \mod 4$, -1 is not a quadratic residue, the left hand side never vanishes and the total number of matrices satisfying the conditions is $q(q-1)^2$. If $q \equiv 1 \mod 4$ then -1 is a quadratic residue, and has two square roots. Since $a - d/2b$ assumes every non-zero value in the field precisely $q(q-1)$ times, there are $2q(q-1)(q-3)$ matrices satisfying the conditions.

In any case, $(a - d/2b)^2$ assumes the value of every non-zero quadratic residue $2q(q-1)$ times. We care only about whether $\Delta(M)$ is a residue or a non-residue to determine whether M is diagonalizable. Recalling the definition of the cyclotomic numbers, we see that the number of matrices which are diagonalizable is

$$2q(q-1)(0,0)_2.$$

For $q \equiv 1 \mod 4$ this evaluates to

$$2q(q-1)(t-1) = 2q(q-1)\frac{(q-5)}{4} = \frac{1}{2}q(q-1)(q-5)$$

and for $q \equiv 3 \mod 4$ this evaluates to

$$2q(q-1)t = 2q(q-1)\frac{(q-3)}{4} = \frac{1}{2}q(q-1)(q-3).$$

The number of matrices which are not diagonalizable is

$$2q(q-1)(0,1)_2.$$

For $q \equiv 1 \mod 4$ this evaluates to

$$2q(q-1)t = 2q(q-1)\frac{(q-1)}{4} = \frac{1}{2}q(q-1)(q-1)$$

and for $q \equiv 3 \mod 4$ this evaluates to

$$2q(q-1)(t+1) = 2q(q-1)\frac{(q+1)}{4} = \frac{1}{2}(q-1)(q+1).$$

□

Next we look at matrices with the conditions that $a - d = 0$ and $b \neq 0$.

Proposition 3.14. *The set of matrices of the form*

$$M = \begin{bmatrix} a & b \\ b & -a \end{bmatrix}$$

with $b \neq 0$ decomposes as follows.

(1) *If $q \equiv 1 \mod 4$ there are $\frac{1}{2}(q-1)(q-3)$ diagonalizable matrices and $\frac{1}{2}(q-1)^2+2(q-1)$ non-diagonalizable matrices, of which $2(q-1)$ have a repeated eigenvalue, all are nilpotent.*

(2) *If $q \equiv 3 \mod 4$ there are $\frac{1}{2}(q-1)^2$ diagonalizable matrices and $\frac{1}{2}(q-1)(q+1)$ non-diagonalizable matrices. None have repeated eigenvalues and none are nilpotent.*

Proof. Since the trace of M is zero, the eigenvalues sum to 0 and the characteristic polynomial of M is $\chi(M) = \lambda^2 - (a^2 + b^2)$. M is nilpotent if $\chi(M) = \lambda^2$. This is true when $b^2 = -a^2$. For a fixed $b \neq 0$, there are two solution to this equation and hence $2(q-1)$ nilpotent matrices when $q \equiv 1 \mod 4$, and none when $q \equiv 3 \mod 4$.
If $(a^2 + b^2)$ is a nonzero square then $\chi(M)$ splits as $(\lambda + \sqrt{a^2 + b^2})(\lambda - \sqrt{a^2 + b^2})$, and the matrix is diagonalizable with distinct eigenvalues. If $a = 0$ this always occurs, giving $q - 1$ matrices. Otherwise, we wish to solve the equation $a^2 + b^2 = c^2$ in the field, which we can rewrite as

$$\left(\frac{a}{b}\right)^2 + 1 = \left(\frac{c}{b}\right)^2.$$

As a and b take the values of the non-zero field elements, each quadratic residue occurs on the left side $2(q-1)$ times. This equation has $2(q-1)(0,0)_2$ solutions. Thus the number of diagonalizable matrices of the given form is

$$2(q-1)(0,0)_2 + (q-1).$$

This evaluates as $\frac{1}{2}(q-1)(q-3)$ if $q \equiv 1 \mod 4$ and as $\frac{1}{2}(q-1)^2$ if $q \equiv 3 \mod 4$.
Finally, if $a^2 + b^2$ is a non-square, the matrix will not be diagonalizable over the base field. The number of matrices in this case will be $2(q-1)(0,1)_2$. This is $(q-1)^2/2$ if $q \equiv 1 \mod 4$ and $(q-1)(q+1)/2$ if $q \equiv 3 \mod 4$. To summarise: the total count for $q \equiv 1 \mod 4$ is:

- $\frac{1}{2}(q-1)(q-3)$ diagonalizable matrices, all with distinct eigenvalues.
- $\frac{1}{2}(q-1)^2 + 2(q-1)$ non-diagonalizable matrices, of which $2(q-1)$ have a repeated root, all of which are nilpotent.

It can easily be verified that these figures sum to $q(q-1)$, as required.

The total count for $q \equiv 3 \mod 4$ is:

- $\frac{1}{2}(q-1)^2$ diagonalizable matrices, all with distinct eigenvalues.
- $\frac{1}{2}(q-1)(q+1)$ non-diagonalizable matrices, all with distinct eigenvalues.

Again, the total count of matrices here is $q(q-1)$. This completes the proof. □

We now characterise the cases where the discriminant is zero.

Proposition 3.15. *Suppose that M is symmetric with $b \neq 0$ and $a - d \neq 0$ (that is: M is not diagonal, and the trace is non-zero), and that $\Delta(M) = 0$. If $q \equiv 1 \mod 4$ then there are $2q(q-1)$ such matrices with a repeated eigenvalue. If $q \equiv 3 \mod 4$ there are no such matrices with $\Delta(M) = 0$.*

Proof. The discriminant vanishes precisely when $(a-d)^2 = -(2b)^2$. By Proposition 3.9 there are **no** non-trivial solutions when $q \equiv 3 \mod 4$.

When $q \equiv 1 \mod 4$ there are solutions. There are q choices for a, leaving $q-1$ choices for d to ensure that the trace does not vanish. For each choice of a and d, there are two solutions to $(a-d)^2 = -(2b)^2$, given by $\pm 2\epsilon b$ where $\epsilon^2 = -1$. Hence there are $2q(q-1)$ matrices satisfying the conditions of the proposition. □

We provide an example to demonstrate the computation of symmetric nilpotent matrices over a finite field, since this phenomenon does not arise in characteristic 0.

Example 3.1. *It seems more convenient to fix the off diagonal elements in this computation: suppose that M is a matrix with diagonal entries a and d and off diagonal entry $b \equiv 7 \mod 17$. A matrix with repeated eigenvalues necessarily has $(a-d)^2 = -(14)^2$, that is: $(a-d)^2 \equiv 8 \mod 14$. So $(a-d) \equiv \pm 5 \mod 17$. In particular, every matrix with a repeated eigenvalue and off diagonal entry 7 is of the form*

$$\begin{pmatrix} a & 7 \\ 7 & a+5 \end{pmatrix}, \quad \begin{pmatrix} a & 7 \\ 7 & a-5 \end{pmatrix}.$$

There are $2q$ matrices with a repeated eigenvalue with off diagonal entries equal to b. The matrices come in pairs (as displayed) with the same eigenvalues. Now, since the eigenvalues are equal and their sum is the trace of the matrix, we have that $2a + 5 = 2\lambda$. Hence $\lambda = a + 9 \cdot 5 \equiv a + 11 \mod 17$. So for $a = 6$, we obtain a symmetric nilpotent matrix.

Now we summarise the results of our computations in two theorems.

Theorem 3.1. *If $q \equiv 1 \mod 4$, the number of symmetric diagonalizable 2×2 matrices is*

$$\frac{1}{2}q(q-1)(q-5) + q^2 + \frac{1}{2}(q-1)(q-3) = \frac{1}{2}\left(q^3 - 3q^2 + q + 3\right)$$

The number of symmetric non-diagonalizable 2×2 matrices is

$$\frac{1}{2}q(q-1)(q-1) + \frac{1}{2}(q-1)(q-1) + 2(q-1) + 2q(q-1) = \frac{1}{2}\left(q^3 + 3q^2 - q - 3\right).$$

Of these matrices, $2q^2 - q$ have a repeated eigenvalue. Of these matrices, the q scalar matrices are diagonalizable, and the remainder have a non-trivial Jordan block. Of those having a non-trivial Jordan block, $2q$ are nilpotent.

Theorem 3.2. *If $q \equiv 3 \mod 4$, the number of symmetric diagonalizable 2×2 matrices is*

$$\frac{1}{2}q(q-1)(q-3) + q^2 + \frac{1}{2}(q-1)^2 = \frac{1}{2}\left(q^3 - q^2 + q + 1\right)$$

The number of symmetric non-diagonalizable 2×2 matrices is

$$\frac{1}{2}q(q-1)(q+1) + \frac{1}{2}(q-1)(q+1) = \frac{1}{2}\left(q^3 + q^2 - q - 1\right).$$

The q scalar matrices are the only ones having a repeated eigenvalue.

To summarise, our counts for $q \equiv 1 \mod 4$ are

	R	NR
D	q	$\frac{1}{2}(q^3 + 3q^2 - q + 3)$
ND	$2q^2 - 2q$	$\frac{1}{2}(q^3 - q^2 - 5q - 3)$

and our counts for $q \equiv 3 \mod 4$ are

	R	NR
D	q	$\frac{1}{2}(q^3 - q^2 - q + 1)$
ND	0	$\frac{1}{2}(q^3 + q^2 - q - 1)$

We conclude this chapter with some further observations on the structure of a 2×2 symmetric nilpotent matrix.

Proposition 3.16. *Suppose that M is a nilpotent matrix of the form*

$$\begin{pmatrix} x & 1 \\ 1 & y \end{pmatrix}$$

Then $x + y = 0$ and $xy = 1$.

Proof. A non-zero nilpotent matrix has characteristic polynomial λ^2. So the trace and determinant of M must be zero. Hence $x + y = 0$ and $xy - 1 = 0$. The result follows. □

Eliminating y in the above Proposition, we find that the matrix M can be written as follows:

$$\begin{pmatrix} x & 1 \\ 1 & -x \end{pmatrix}$$

It is now clear that the determinant vanishes if and only if $x^2 = -1$ and hence we have proved the following result.

Proposition 3.17. *There exists a nilpotent symmetric matrix over a field $\mathbb{F}$ if and only if -1 is a square in $\mathbb{F}$.*

Furthermore, every nilpotent 2×2 matrix is of the form

$$\lambda \begin{pmatrix} x & 1 \\ 1 & -x \end{pmatrix}$$

where x is a square root of -1. In fact there is close connection between the failure of a matrix to be diagonalizable and the existence of nilpotent matrices of a particular form. We will investigate these issues further in the next chapter.

4. Nilpotent Self-adjoint matrices in Positive Characteristic

Recall that the Jordan Canonical Form gives a basis with respect to which a given linear transformation can be written in near diagonal form. We can be a little more precise about the structure of this form.

Proposition 4.1. *Any matrix can be written as a sum*

$$M = S + N$$

where S is diagonalizable, N is nilpotent and the matrices S and N have common generalised eigenspaces.

Proof. By the Jordan Canonical form theorem, there exists a basis such that M has non-zero entries only on and directly above the diagonal. Furthermore, the entries above the diagonal are 0 or 1, and an entry is 1 only if the entries below and to the right are equal.

Let S be the matrix containing the diagonal entries of M (with respect to the Jordan basis) and N the matrix containing the enties above the diagonal. Then S is diagonalizable (even diagonal!) and N is strictly upper triangular, and hence nilpotent. To show that S and N commute, it is enough to verify this for a single generalised eigenspace: but here S restricts to a scalar matrix, so the result is immediate. □

We have shown that every matrix is the sum of a diagonalizable and of a nilpotent matrix. Suppose that there existed a nilpotent symmetric matrix of dimension d over the field k. Clearly such a matrix is not diagonalizable. Conversely if every symmetric matrix of dimension d over k is diagonalizable, then there are no symmetric nilpotent matrices. We conclude the following.

Proposition 4.2. *The following are equivalent:*

- *There exists a nilpotent symmetric matrix in $M_d(k)$.*
- *There exists a symmetric matrix in $M_d(k)$ which is not diagonalizable.*

This motivates an exploration of nilpotent symmetric matrices.

4.1. Bilinear and Quadratic Forms.

Definition 4.1. *Let V be a vector space over $\mathbb{F}$. A mapping $\langle , \rangle : V \times V \to \mathbb{F}$ is called a **bilinear form** if for any $x, y, z \in V$ and $\alpha, \beta \in \mathbb{F}$,*

$$\langle \alpha x + \beta y, z \rangle = \alpha \langle x, z \rangle + \beta \langle y, z \rangle$$

and

$$\langle z, \alpha x + \beta y \rangle = \alpha \langle z, x \rangle + \beta \langle z, y \rangle$$

We denote the set of bilinear forms on V by $\mathcal{B}(V)$.
Furthermore, a bilinear form is said to be

(1) ***symmetric** if $\langle x, y \rangle = \langle y, x \rangle$*
(2) ***skew-symmetric** if $\langle x, y \rangle = -\langle y, x \rangle$*
(3) ***alternating** if $\langle x, x \rangle = 0$*

for all $x, y \in V$.

Definition 4.2. *A bilinear form* $\langle , \rangle$ *on a finite-dimensional vector space* V *is called* ***diagonalizable*** *if there is an ordered basis* B *for* V *such that* M_B *is a diagonal matrix.*

Proposition 4.3. *Every diagonalizable bilinear form on a finite-dimensional vector space* V *is symmetric.*

Proof. Suppose $\langle , \rangle$ is a diagonalizable bilinear form. Then there exists an ordered basis B such that M_B is a diagonal matrix. Trivially, M_B is a symmetric matrix, and hence, by the previous proposition, $\langle , \rangle$ is symmetric. □

To prove the converse, which holds only in fields not of characteristic 2, we use the following lemma.

Lemma 4.1. *Let* V *be a vector space over a field not of characteristic 2. Then no nonzero symmetric bilinear form on* V *is alternating.*

Proof. Suppose $\langle , \rangle$ is a nonzero symmetric bilinear for on V. It suffices to show there exists a vector $x \in V$ such that $\langle x, x \rangle \neq 0$. Since $\langle , \rangle$ is nonzero, there are vectors $u, v \in V$ such that $\langle u, v \rangle \neq 0$. If $\langle u, u \rangle \neq 0$ or $\langle v, v \rangle \neq 0$, the result is immediate. Otherwise, set $x = u + v$. Then

$$\langle x, x \rangle = \langle u, u \rangle + \langle u, v \rangle + \langle v, u \rangle + \langle v, v \rangle = 2\langle u, v \rangle \neq 0$$

since we assumed $\langle u, v \rangle \neq 0$. □

Example 4.1. *Let* $V = \mathbb{R}^n$. *For* $x = (a_1, a_2, .., a_n)$ *and* $y = (b_1, b_2, .., b_n)$ *in* V, *the dot product, the map* $\langle , \rangle : V \times V \to \mathbb{R}$ *defined by*

$$\langle x, y \rangle = \sum_{i=1}^{n} a_i b_i$$

is a symmetric bilinear form.

The next proposition tells us we really need only consider symmetric and alternating forms. Here we note that we do not consider fields of characteristic 2: this will be a recurring theme.

Proposition 4.4. *Let* V *be a vector space over a field* $\mathbb{F}$. *If* $char(\mathbb{F}) \neq 2$, *then a bilinear form on* V *is alternating if and only if it is skew-symmetric.*

Proof. First, suppose $\langle , \rangle$ is an alternating inner product. Then

$$0 = \langle x + y, x + y \rangle = \langle x + y \rangle + \langle x + y \rangle$$

and so

$$\langle x, y \rangle = -\langle y, x \rangle$$

which shows that $\langle , \rangle$ is skew-symmetric.
Now suppose $\langle , \rangle$ is skew-symmetric. Then

$$\langle x, x \rangle = -\langle x, x \rangle \quad \Longrightarrow \quad 2\langle x, x \rangle = 0 \quad \Longrightarrow \quad \langle x, x \rangle = 0$$

and hence $\langle , \rangle$ is alternating. □

An inner product space with a symmetric or an alternating form is called an **orthogonal geometry** or a **symplectic geometry**, respectively.
Like linear transformations, bilinear forms can be represented as matrices.

Definition 4.3. *Let V be an inner product with an ordered basis $B = \{b_1, ..., b_n\}$. A bilinear form is completely determined by the matrix $M_B \in M_{n\times n}(\mathbb{F})$ defined by*

$$(M_B)_{ij} = \langle b_i, b_j \rangle.$$

*This is called the **matrix of the form** with respect to B. Moreover, any $A \in M_n(\mathbb{F})$ is the matrix of some bilinear form on V.*

Example 4.2. *The Minkowski space M_4 is the four-dimensional real orthogonal geometry $\mathbb{R}^4$ with inner product defined by*

$$\begin{aligned}\langle e_1, e_1\rangle &= \langle e_2, e_2\rangle = \langle e_3, e_3\rangle = 1\\ \langle e_4, e_4\rangle &= -1\\ \langle e_i, e_j\rangle &= 0 \text{ for } i \neq j\end{aligned}$$

where $e_1, ..., e_4$ is the standard basis for $\mathbb{R}^4$.
The matrix of the form with respect to the standard basis is

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

Intuitively, symmetric forms have symmetric matrices and vice-versa.

Proposition 4.5. *Let V be a finite-dimensional vector space, and let B be an ordered basis for V. A bilinear form $\langle , \rangle$ on V is symmetric if and only if its matrix with respect to B is symmetric.*

Proof. Let $B = \{v_1, .., v_n\}$. First assume the form is symmetric. Then for $1 \leq i, j \leq n$,

$$(M_B)_{ij} = \langle v_i, v_j\rangle = \langle v_j, v_i\rangle = (M_B)_{ij}$$

and hence M_B is symmetric.

Conversely, suppose M_B is symmetric. Let $\langle , \rangle' : V \times V \to \mathbb{F}$ be the mapping defined by $\langle x, y\rangle' = \langle y, x\rangle$ for all $x, y \in V$. Observe that $\langle , \rangle'$ is a bilinear form, and let M'_B be its matrix with respect to B. Then for $1 \leq i, j \leq n$,

$$(M'_B)_{ij} = \langle v_i, v_j\rangle' = \langle v_j, v_i\rangle = (M_B)_{ji} = (M_B)_{ij}$$

and hence $M'_B = M_B$. Since there is a one-to-one correspondence between bilinear forms and their matrices, we have $\langle , \rangle' = \langle , \rangle$. Hence $\langle y, x\rangle = \langle x, y\rangle' = \langle x, y\rangle$ for all $x, y \in V$, and therefore $\langle , \rangle$ is symmetric. □

Definition 4.4. *Two matrices $A, B \in M_n(\mathbb{F})$ are said to be **congruent** if there exists an invertible matrix P for which*

$$A = P^T BP.$$

Observe that congruence is an equivalence relation.

Proposition 4.6. *Two matrices A and B represent the same bilinear form on a vector space V if and only if they are congruent.*

Proof. Congruence is an equivalence relation, so if two matrices represent the same bilinear form on V, they must be congruent. Conversely, if $Q = [M]_B$ represents a bilinear form on V and $R = P^T Q P$ where P is invertible, then there is an ordered basis C for V for which

$$P = [M]_C^B$$

and so

$$Q = ([M]_C^B)^T [M]_B [M]_C^B$$

Thus, $Q = [M]_C$ represents the same form with respect to C. □

Symmetric bilinear forms are associated with functions called quadratic forms.

Definition 4.5. *Let V be a vector space over $\mathbb{F}$. A map $Q : V \to \mathbb{F}$ is called a **quadratic form** if there exists a symmetric bilinear form $\langle , \rangle_Q$ on V such that*

$$Q(v) = \langle v, v \rangle_Q \quad \textit{for all } v \in V.$$

If $\mathbb{F}$ is not of characteristic 2, there is a one-to-one correspondence between symmetric bilinear forms and quadratic forms.

4.2. **Isotropic Vectors.**

Isotropic vectors are important to understand why the spectral theorem fails over fields of positive characteristic.

Definition 4.6. *Let V be an inner product space.*

- *A nonzero $x \in V$ is **isotropic** if $\langle x, x \rangle = 0$.*
- *V is **isotropic** if it contains at least one isotropic vector.*
- *A subspace U of V is **isotropic** if all vectors in U are pairwise orthogonal (including $\langle x, x \rangle = 0$).*

Definition 4.7. *Let V be an inner product space. A vector $v \in V$ is called **degenerate** if $v \perp V$. We call the set of all degenerate vectors the **radical** of V, denoted $rad(V)$. Thus, $rad(V) = V^\perp$.*
*We call V **nonsingular** if $rad(V) = \{0\}$ and **singular** otherwise.*

Proposition 4.7. *If V is nonsingular and T is an isometry of V, then $\det T = \pm 1$. If $\det T = 1$, we call T a **rotation**. If $\det T = -1$, we call T a **reflection**. The rotations form an invariant subgroup of the isometry group of V whose index is at most 2.*

Proof. The last part follows from the fact that the map $T \mapsto \det T$ is a homomorphism of the isometry group of V, whose kernel are the rotations and whose image is ± 1. □

Proposition 4.8. *[Rom08] Let V be a vector space with a bilinear form. The following are equivalent:*

(1) *Orthogonality is a symmetric relation on V, that is,*

$$x \perp y \iff y \perp x.$$

(2) *The form on V is symmetric or alternate.*

Proof. It is clear that orthogonality is symmetric if the form is symmetric or alternate (for the latter case, recall that we have shown that alternate forms are also skew-symmetric).
Now assume that orthogonality is symmetric. Let $x \bowtie y$ mean that $\langle x, y\rangle = \langle y, x\rangle$ and let $x \bowtie V$ mean that $\langle x, v\rangle = \langle v, x\rangle$ for all $v \in V$. If $x \bowtie V$ for all $x \in V$, then we are done. So assume $x \not\bowtie V$. We wish to show that

$$x \not\bowtie V \implies x \text{ is isotropic} \quad \text{and} \quad (x \bowtie y \implies x \perp y) \tag{2}$$

Note that if the second conclusion holds, then since $x \bowtie x$ it follows that x is isotropic. So suppose $x \bowtie y$. Since $x \not\bowtie V$, there exists $z \in V$ such that $\langle x, z\rangle \neq \langle z, x\rangle$. Thus $x \perp y$ if and only if

$$\langle x, y\rangle(\langle x, z\rangle - \langle z, x\rangle) = 0.$$

This can be rewritten as

$$\begin{aligned}\langle x, y\rangle(\langle x, z\rangle - \langle z, x\rangle) &= \langle x, y\rangle\langle x, z\rangle - \langle x, y\rangle\langle z, x\rangle \\ &= \langle y, x\rangle\langle x, z\rangle - \langle x, y\rangle\langle z, x\rangle \\ &= \langle x, \langle y, x\rangle z - y\langle z, x\rangle\rangle.\end{aligned}$$

Reversing the arguments in the last expression gives us

$$\langle\langle y, x\rangle z - y\langle z, x\rangle, x\rangle = \langle y, x\rangle\langle z, x\rangle - \langle y, x\rangle\langle z, x\rangle = 0.$$

By our assumption that orthogonality is symmetric, the last expression is 0 and we have proven (2).
Let us assume the form on V is not symmetric and show this implies all vectors in V are isotropic (and hence the form on V is alternate). By our assumption, there exist $u, v \in V$ such that $u \not\bowtie v$, and so $u \not\bowtie V$ and $v \not\bowtie V$. Hence u and v are isotropic and for all $y \in V$,

$$\begin{aligned} y \bowtie u &\implies y \perp u \\ y \bowtie v &\implies y \perp v\end{aligned}$$

Since every $w \in V$ for which $w \not\bowtie V$ is isotropic, let $w \bowtie V$. Then $w \bowtie u$ and $w \bowtie v$ and hence $w \perp u$ and $w \perp v$. Now write

$$w = (w - u) + u$$

where $(w - u) \perp u$, since u is isotropic. Since the sum of two orthogonal isotropic vectors is isotropic, it follows that w is isotropic if $w - u$ is isotropic. But

$$\langle w + u, v\rangle = \langle u, v\rangle \neq \langle v, u\rangle = \langle v, w + u\rangle$$

and so $(w + u) \not\bowtie V$, which implies that $w + u$ is isotropic. Thus w is isotropic and so all vectors in V are isotropic. □

We now note that the absence of isotropic vectors is a crucial part of the spectral theorem. The spectral theorem requires the result that a self-adjoint operator over $\mathbb{R}$ or $\mathbb{C}$ is diagonalizable (Proposition 2.2), which uses the positive-definite property of inner products, and hence the absence of isotropic vectors. It is known as a consequence of the following two theorems that a quadratic form in positive characteristic has an isotropic vector in $\dim \geq 3$, thus the spectral theorem fails over fields of positive characteristic.

Theorem 4.1. *(Chevalley-Warning) Let $f_i(x_1, x_2, ..., x_n), i = 1, 2, ..., r$ be a polynomial over $\mathbb{F}_q$. If $n > \sum_{i=1}^{r} \deg(f_i)$, then the number of solutions to $f_1 = f_2 = ... = f_r = 0$ is divisible by q.*

This immediately leads to Chevalley's theorem, since q is at least 2.

Theorem 4.2. *(Chevalley) If the system of equations $f_1 = f_2 = ... = f_r = 0$ has the trivial solution, i.e. the polynomials have no constant terms, then the system also has a non-trivial solution.*

4.3. **Counting Nilpotent Matrices.**

In this section we give a theorem of Hall on the total number of nilpotent matrices in dimension d over a finite field. Afterwards, we will discuss the enumeration of symmetric nilpotent matrices, which is more involved (similar to the enumeration we carried out in Chapter 3). The material in this section is due to Brouwer, Gow and Sheekey, [SGB14].

Definition 4.8. *Let V be a finite-dimensional vector space and let M be a linear transformation on V. The **Fitting decomposition** of M is the unique decomposition $V = U \oplus W$ as unique sum of M-invariant subspaces U and W, such that $M\,|_U$ is nilpotent and $M\,|_W$ is invertible.*

Again, we can use the Jordan Canonical Form to prove that the Fitting decomposition of a matrix exists and is unique: the nilpotent portion of M is the generalised eigenspace at the eigenvalue 0 (which is an invariant subspace of M) and the invertible portion is the direct sum of all other generalised eigenspaces.

For a subspace S will use $\mathcal{N}(S)$ to denote the number of nilpotent operators on S, and $\mathcal{I}(S)$ to denote the number of invertible operators on S. Since these only depend on the dimension k of S, we may write $\mathcal{N}(S) = \mathcal{N}(k)$ and $\mathcal{I}(S) = \mathcal{I}(k)$. Furthermore, note that $\mathcal{I}(n) = |GL_n(q)|$.

Definition 4.9. *Treating q as a variable, define $[n] = q^n - 1$ for any integer n. In analogy to the usual factorial, define the q-factorial to be*

$$[n]! = [n] \cdot [n-1] \cdots [1] = \prod_{i=1}^{n} (q^i - 1)\,.$$

*The **Gaussian binomial coefficient** is*

$$\begin{bmatrix} n \\ k \end{bmatrix} = \frac{[n]!}{[k]![n-k]!}\,.$$

Just as binomial coefficients count subsets of an n set and factorials count the elements of a symmetric group, their q analogues arise with respect to counting substructures and symmetries of vector spaces. In fact, the Gaussian binomial coefficients count k-dimensional subspaces of an n dimensional space and the formula for the order of $\mathrm{GL}_n(q)$ may be expressed as $q^{\binom{n}{2}}[n]!$.

Theorem 4.3. *[SGB14] The number of nilpotent matrices in an n-dimensional vector space over $\mathbb{F}_q$ is $q^{n(n-1)}$.*

Proof. For an n-dimensional vector space V over $\mathbb{F}_q$, there are q^{n^2} matrices which each have a unique fitting decomposition. This yields the equality

$$q^{n^2} = \sum_{V=U\oplus W} \mathcal{N}(U)\mathcal{I}(W).$$

The number of ways to write as the direct sum of an m-space and an $(n-m)$-space is

$$\frac{\mathcal{I}(n)}{\mathcal{I}(m)\mathcal{I}(n-m)} = \frac{|GL_n(q)|}{|GL_m(q)||GL_{n-m}(q)|} = q^{m(n-m)} \begin{bmatrix} n \\ m \end{bmatrix}_q$$

Thus we obtain

$$\begin{aligned} q^{n^2} &= \frac{\mathcal{I}(n)}{\mathcal{I}(m)\mathcal{I}(n-m)} \sum_{m=0}^{n} \mathcal{N}(m)\mathcal{I}(n-m) \\ &= \mathcal{I}(n) \sum_{m=0}^{n} \frac{\mathcal{N}(m)}{\mathcal{I}(m)} \end{aligned}$$

We proceed by induction on n.

- $n = 0$: By definition $\mathcal{I}(0) = \mathcal{N}(0) = 1$.
- $n = 1$: $\mathcal{I}(1) = q - 1$, $\mathcal{N}(1) = 1$ and we have
$$q = (q-1)\left(\frac{1}{1} + \frac{1}{q-1}\right).$$
- $n = 2$: $\mathcal{I}(2) = |GL_2(q)| = (q^2-1)(q^2-q)$. We solve for $\mathcal{N}(2)$:
$$\begin{aligned} q^4 &= (q^2-1)(q^2-q)\left(\frac{1}{1} + \frac{1}{q-1} + \frac{\mathcal{N}(2)}{(q^2-1)(q^2-q)}\right) \\ &= (q^2-1)(q^2-q) + (q+1)(q^2-q) + \mathcal{N}(2) \end{aligned}$$
Thus
$$\mathcal{N}(2) = q^4 - (q^2+q)(q^2-q) = q^4 - (q^4-q^2) = q^2.$$

Now assume that $\mathcal{N}(n) = q^{n(n-1)}$. Then

$$\begin{aligned} q^{n^2} &= |GL_n(q)| \sum_{m=0}^{n} \frac{q^{m(m-1)}}{|GL_m(q)|} \\ &= \prod_{i=0}^{n-1}(q^n - q^i) \sum_{m=0}^{n} \frac{q^{m(m-1)}}{\prod_{j=0}^{m-1}(q^m - q^j)} \end{aligned}$$

The induction step is

$$\begin{aligned}
q^{(n+1)^2} &= |GL_{n+1}(q)|\left[\left(\sum_{m=0}^{n}\frac{q^{m(m-1)}}{|GL_m(q)|}\right)+\frac{\mathcal{N}(n+1)}{|GL_{n+1}(q)|}\right]\\
&= \prod_{i=0}^{n}(q^{n+1}-q^i)\sum_{m=0}^{n}\frac{q^{m(m-1)}}{\prod_{j=0}^{m-1}(q^m-q^j)}+\mathcal{N}(n+1)\\
&= (q^{n+1}-1)\prod_{m=1}^{n}q(q^n-q^{m-1})\sum_{m=0}^{n}\frac{q^{m(m-1)}}{\prod_{j=0}^{m-1}(q^m-q^j)}+\mathcal{N}(n+1)\\
&= q^n(q^{n+1}-1)\prod_{m=0}^{n-1}(q^n-q^m)\sum_{m=0}^{n}\frac{q^{m(m-1)}}{\prod_{j=0}^{m-1}(q^m-q^j)}+\mathcal{N}(n+1)\\
&= q^n(q^{n+1}-1)|GL_n(q)|\sum_{m=0}^{n}\frac{q^{m(m-1)}}{|GL_m(q)|}+\mathcal{N}(n+1)\\
&= q^n(q^{n+1}-1)q^{n^2}+\mathcal{N}(n+1)
\end{aligned}$$

Hence

$$\begin{aligned}
\mathcal{N}(n+1) &= q^{n^2+2n+1}-(q^{n+1}-1)q^{n^2+n}\\
&= q^{n^2+2n+1}-q^{n^2+2n+1}+q^{n^2+n}\\
&= q^{(n+1)n}
\end{aligned}$$

So $\mathcal{N}(n+1)=q^{(n+1)((n+1)-1)}$ by induction. □

4.4. **Symmetric Nilpotent matrices.** In this section, we give an outline of a method used by Brouwer, Gow and Sheekey to count symmetric nilpotent matrices with a given Jordan Canonical Form. We begin by describing a bijection between JCFs and integer partitions.

Definition 4.10. *A partition of n is an unordered sequence of positive integers $(a_1,\ldots,a_t)$ such that $a_1+a_2+\ldots+a_t=n$.*

There is a bijection between JCFs of nilpotent matrices and partitions, given by the sizes of the Jordan blocks. To be completely explicit we give this bijection for the 4×4 matrices. Each 4×4 nilpotent is conjugate to one of the following matrices:

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

We associate to each of these matrices a *partition* of 4, in which the parts are the nilpotency classes of the blocks along the diagonal. Thus the displayed matrices correspond to the partitions

$$1{+}1{+}1{+}1,\qquad 2{+}1{+}1,\qquad 2{+}2,\qquad 3{+}1\qquad 4$$

respectively. In general there is a bijection between conjugacy classes of nilpotent matrices in $M_n(k)$ and partitions of n, where the parts are the heights of the generalised eigenvectors of the corresponding matrix. A standard result in enumerative combinatorics shows that the number of partitions of n grows proportionally to $e^{\sqrt{2n}}$, which is faster than any polynomial function but slower than exponential.

Next, we must determine the size of each conjugacy class of nilpotent matrices in $M_4(q)$. This is, as usual, an application of the Orbit-Stabilizer Theorem.

Example 4.3. *We count the number of matrices in $M_4(q)$ with JCF*

$$J = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix}.$$

We use the Orbit-Stabilizer Theorem to compute this.

$$|\textit{Orbit}||\textit{stabilizer}| = |GL_4(q)|$$

The size of the stabilizer is the number of invertible matrices $M = \begin{bmatrix} a & b & c & d \\ e & f & g & h \\ i & j & k & l \\ m & n & o & p \end{bmatrix}$ which commute with J. Observe that

$$\begin{bmatrix} a & b & c & d \\ e & f & g & h \\ i & j & k & l \\ m & n & o & p \end{bmatrix} \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix} = \begin{bmatrix} 0 & a & 0 & c \\ 0 & e & 0 & g \\ 0 & i & 0 & k \\ 0 & m & 0 & o \end{bmatrix}$$

$$\begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} a & b & c & d \\ e & f & g & h \\ i & j & k & l \\ m & n & o & p \end{bmatrix} = \begin{bmatrix} e & f & g & h \\ 0 & 0 & 0 & 0 \\ m & n & o & p \\ 0 & 0 & 0 & 0 \end{bmatrix}.$$

These $MJ = JM$ if and only if

$$\begin{bmatrix} 0 & a & 0 & c \\ 0 & e & 0 & g \\ 0 & i & 0 & k \\ 0 & m & 0 & o \end{bmatrix} = \begin{bmatrix} e & f & g & h \\ 0 & 0 & 0 & 0 \\ m & n & o & p \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

i.e.,

$$\begin{array}{cccc} e = 0 & a = f & g = 0 & c = h \\ m = 0 & n = i & o = 0 & k = p. \end{array}$$

Thus M has the form

$$\begin{bmatrix} a & b & c & d \\ 0 & a & 0 & c \\ i & j & k & l \\ 0 & i & 0 & k \end{bmatrix}.$$

M is invertible if and only if

$$\begin{aligned} 0 \neq \det(M) &= a \begin{vmatrix} a & 0 & c \\ j & k & l \\ i & 0 & k \end{vmatrix} + i \begin{vmatrix} b & c & d \\ a & 0 & c \\ i & 0 & k \end{vmatrix} \\ &= ak \begin{vmatrix} a & c \\ i & k \end{vmatrix} - ci \begin{vmatrix} a & c \\ i & k \end{vmatrix} \\ &= ak(ak - ci) - ci(ak - ci) \\ &= (ak - ci)^2. \end{aligned}$$

$0 \neq (ak - ci)^2$ if and only if $0 \neq ak - ci = \det \begin{bmatrix} a & i \\ c & k \end{bmatrix}$. $|GL_2(q)|$ counts the possibilities for a, i, c, k. With four other free variables, the total count is

$$q^4|GL_2(q)| = q^4(q^2 - 1)(q^2 - q).$$

Thus the orbit is

$$\frac{(q^4 - 1)(q^4 - q)(q^4 - q^2)(q^4 - q^3)}{q^4(q^2 - 1)(q^2 - q)} = \frac{q^6(q^4 - 1)(q^3 - 1)(q^2 - 1)(q - 1)}{q^5(q^2 - 1)(q - 1)} = q(q^4 - 1)(q^3 - 1).$$

Hence the number of matrices in $M_4(q)$ with JCF J is $q(q^4 - 1)(q^3 - 1)$.

In fact, this computation is generalised by Brouwer Gow and Sheekey to give a formula for the number of matrices commuting with an arbitrary nilpotent matrix. This makes use of the following object.

Definition 4.11. *The **Young diagram** Y_λ of a partition λ is the shape obtained by placing left-adjusted squares of length $\lambda_1, .., \lambda_m$ below each other.*

Example 4.4. *If $\lambda = (5, 4, 2, 2)$, then*

$Y_\lambda =$ 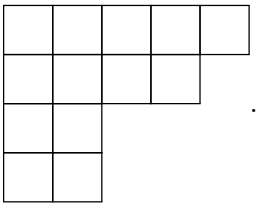 .

Proposition 4.9. *Suppose that N is a nilpotent matrix in Jordan Canonical Form with Jordan blocks ordered by size. Suppose further that blocks $C_i, \ldots, C_j$ all have size c. These blocks contribute to the number of matrices commuting with N as follows:*

(1) *A power of q. The exponent of q in this contribution is $\binom{j}{2} - \binom{i-1}{2}$.*
(2) *A part coprime to q: $[j - i + 1]!$*

(3) *After computing contributions from the cells of each size, multiply by q^A, where A is as follows: for each square in the Young diagram, add the number of cells in the column immediately to the left.*

To illustrate this proposition, we re-compute the number of matrices commuting with the example prior: there are two Jordan blocks of size 2, so the computation involves a 2×2 square. So $j = 2$, and $(i-1) = 0$ and $A = 2 + 2 = 4$. We obtain the quantity

$$q^{1-0}[2][1] \cdot q^4 = q^5(q^2-1)(q-1)$$

just as in the direct computation.
We provide another example demonstrating this type of computation on a more complex matrix.

Example 4.5. *Suppose that $N = \oplus_{i=1}^{6} A_i$ is a nilpotent matrix in Jordan Canonical Form with Jordan blocks*

$$A_1 = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix}, \quad A_2 = A_3 = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix}, \quad \textit{and} \quad A_4 = A_5 = A_6 = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}.$$

The Young diagram of N is

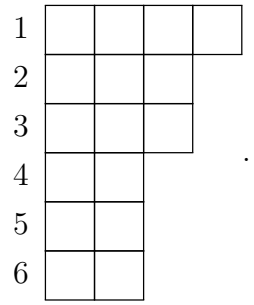

The number of commuting matrices contributed by each of the Jordan blocks is

$$\left.\begin{matrix} A_1 \end{matrix}\right\} \quad q^{\binom{1}{2}-\binom{0}{1}}[1-1+1]! = [1]!$$

$$\left.\begin{matrix} A_2 \\ A_3 \end{matrix}\right\} \quad q^{\binom{3}{2}-\binom{1}{2}}[3-2+1]! = q^3[2]!$$

$$\left.\begin{matrix} A_4 \\ A_5 \\ A_6 \end{matrix}\right\} \quad q^{\binom{6}{2}-\binom{3}{2}}[6-4+1]! = q^{15}[3]!$$

Finally, we find q^A:

$$A = 36 + 18 + 3 = 57$$

	6	6	3
	6	6	
	6	6	
	6		
	6		
	6		

Thus the number of matrices commuting with N is

$$q^{57}[1]!q^3[2]!q^{15}[3]! = q^{75}(q^3-1)(q^2-1)^2(q-1)^3.$$

Definition 4.12. *Nonzero even dimension n nondegenerate symmetric bilinear forms over fields of odd characteristic have two types. Let M be the matrix of such a form. We call the form **hyperbolic** precisely when $(-1)^{n/2}\det(M)$ is a square. The standard form is hyperbolic when $q \equiv 1 \mod 4$, and also when $4|n$, and **elliptic** otherwise.*

Next, Brouwer Gow and Sheekey compute the number of forms of elliptic and hyperbolic type that exist over a finite field.

Theorem 4.4. *The number of non-degenerate elliptic forms over $\mathbb{F}_q^4$ is*

$$F_e = \frac{1}{2}q^4(q^3-1)(q^2-1)(q-1)$$

and the number of non-degenerate hyperbolic quadratic forms over $\mathbb{F}_q^4$ is

$$F_h = \frac{1}{2}q^4(q^3-1)(q^2+1)(q-1)\,.$$

The next step is perhaps the most conceptually difficult in this entire process: we count the number of forms for which a given nilpotent transformation will be nilpotent.

The following formula is given by Brouwer, Gow and Sheekey: suppose that N is in Jordan Canonical Form, and that N contains c_i blocks of size i for each $1 \leq i \leq k$. Then the blocks of size i contribute to the total as follows:

(1) If c_i is even, then $c_i = 2t$ and the contribution to the sum is $q^{t(t+1)}[1][3]\cdots[2t-1]$, where the product is all odd terms up to $2t-1$. If c_i is odd write $c_i = 2t+1$, the contribution is $q^{t(t+1)}[1][3]\cdots[2t+1]$, where all the terms in the product are odd.
(2) Secondly, for each terminal block in a row of length i in the Young diagram, one multiplies by a factor of q for each block *strictly to the left, and at the same level or below* the terminal block.

Thus in our example, the Young diagram is a 2×2 square, so $c_2 = 2$ and all other c_i are zero. The contribution from the two blocks of size 2 is $q^2[1] = q^2(q-1)$. Since there are two blocks to the left of the upper right block, and one to the left of the lower right, we multiply by a factor of q^3: so there are $q^5(q-1)$ forms for which the Jordan form matrix is self adjoint.

Theorem 4.5 ([SGB14]). *Let V be n dimensional over $\mathbb{F}_q$, let F_e be the number of non-degenerate elliptic forms and F_h be the number of non-degenerate hyperbolic forms on V.*

Let N be a nilpotent linear transformation in Jordan Canonical Form, and let $ccl(N)$ be the size of the conjugacy class of N. Let $F_e(N)$ be the number of non-degenerate elliptic forms for which N is self-adjoint, and define $F_h(N)$ similarly. Then the number of symmetric conjugates of N is

$$ccl(N)\frac{F_e(N)}{F_e}$$

if I_n is elliptic and

$$ccl(N)\frac{F_h(N)}{F_h}$$

if I_n is hyperbolic.

In the result, $\mathrm{ccl}(N)$ is computed from Proposition 4.9 (recalling that we seek the size of an orbit, and this result gives the size of the stabilizer in $\mathrm{GL}_n(q)$). The quantity F_e is computed in Theorem 4.4 and the quantity $F_e(N)$ is described directly before the Theorem. All may be computed directly from the Young diagram associated with N.

The logic behind Theorem 4.5 is rather straightforward. First: a symmetric matrix is just a matrix self-adjoint with respect to the standard bilinear form. To count the number of matrices in a Jordan class which are self-adjoint to a particular bilinear form, it is enough to multiply the total number of matrices with that Jordan Canonical Form by the average number of forms for which the matrix is selfadjoint. These other quantities can be computed using sophisticated but standard methods for working with classical groups. Let us illustrate this theorem with the example we have been working with.

- For the given matrix N we computed that
$$|ccl(N)| = q(q^4-1)(q^3-1)\,.$$
- Since 4 is doubly even, I_4 is a hyperbolic form. The total number of hyperbolic forms on a four dimensional vector space is
$$F_h = \frac{1}{2}q^4(q^3-1)(q-1)(q^2+1)\,.$$
- The number of forms with respect to which N is self adjoint is
$$F_h(N) = q^5(q-1)\,.$$

Plugging these values into the formula, we find that N has $2q^2(q^2-1)$ symmetric conjugates. Since the dimension is $n \equiv 0 \mod 4$ this result is independent of q. Over the field of order 3 we did a brute force computation of the number of symmetric matrices conjugate to N. As predicted by the formula we obtained $144 = 2 \cdot 3^2 \cdot (3^2-1)$.

5. Conclusion

In this thesis, we set out to understand the relationship between symmetric and diagonalizable matrices. In Chapter 1 we explored Canonical Forms for matrices. Up to details involving computational representations of the underlying field, computation of these Canonical forms gives an effective method of deciding diagonalizability of a matrix. These methods typically have computational complexity at least cubic in the order of the matrix.

In Chapter 2 we saw that symmetric matrices are self-adjoint with respect to the standard form. Over a field of characteristic 0 a self-adjoint matrix is diagonalizable. In particular, a nilpotent symmetric matrix in characteristic zero must be the zero matrix.

In Chapter 3 we explored in detail the case of 2×2 matrices over a finite field of odd characteristic. We saw that the number of diagonalizable matrices depends on $q \mod 4$, and that the number of symmetric matrices which fail to be diagonalizable is related to solutions of the equation $x^2 + 1 = y^2$ in the field. In particular, the number of 2×2 symmetric nilpotents in characteristic p is 1 if $p \equiv 3 \mod 4$ and $2p - 1$ is $p \equiv 1 \mod 4$.

Finally in Chapter 4 we explore the enumeration of symmetric nilpotent matrices using methods developed by Brouwer Gow and Sheekey. They give an algorithm involving combinatorics on Young diagram to evaluate the number of symmetric nilpotent matrices with a given Jordan Canonical Form in any dimension and for any finite field.

As we conclude this project we are left with several ideas for future research.

(1) Can the methods of Brouwer-Gow-Sheekey be extended to compute the number of symmetric non-diagonalizable matrices over a finite field? Does this proportion tend to $1/2$ as $q \to \infty$ for every n?
(2) Our original motivation was to find sufficient criteria for diagonalizability over a finite field which would be visible to the naked eye. (To be more precise: computable in linear time in the number of matrix entries, like the symmetric condition.) Does such a criterion exist?
(3) From the theory of the Jordan Canonical Form, any matrix can be written uniquely as the sum of a diagonalizable matrix and a nilpotent matrix. As proved in Chapter 4, the number of $n \times n$ nilpotent matrices over a finite field of order q is $q^{n(n-1)}$ while the number of symmetric matrices is $q^{n(n+1)}$. We have seen that in characteristic 0, these sets intersect trivially, while in positive characteristic they seem to intersect as two random sets would. Are there other large subspaces of matrices which are all diagonalizable? Equivalently: are there large subspaces of matrices in positive characteristic which are disjoint from the non-zero nilpotent matrices?

6. Appendix: Alternate Proof of Cayley-Hamilton

The material in this section is drawn from [Axl15] and [FIS97].
Throughout, we assume:

- $\mathbb{F}$ is an algebraically closed field.
- V is a finite-dimensional nonzero vector space over $\mathbb{F}$.
- $\mathcal{L}(V)$ denotes the set of linear operators on V.

Remark: for $n = \dim(V)$, we can identify $\mathcal{L}(V)$ with $M_n(\mathbb{F})$ up to choice of basis. In particular, since the identity matrix I_n is unique, for any $T \in \mathcal{L}(V)$ we can define $T^0 = I^n$.

Proposition 6.1. *Suppose $T \in \mathcal{L}(V)$. Then for every nonnegative integer k,*

$$N(T^k) \subseteq N(T^{k+1})$$

Proof. If $k = 0$, then $\{0\} = N(T^0) \subseteq N(T^1)$ is immediate. Now suppose $v \in N(T^k)$. Then $T^{k+1}v = T(T^k v) = T(0) = 0$, hence $v \in N(T^{k+1})$. □

Proposition 6.2. *Suppose $T \in \mathcal{L}(V)$. Let $n = \dim(V)$. Then for every positive integer k,*

$$N(T^n) = N(T^{n+k})$$

Proof. By Proposition 2.1, clearly $N(T^n) \subseteq N(T^{n+k})$. Conversely, suppose that $v \in N(T^{n+k})$ but $v \notin N(T^n)$. Let $j+1$ be the smallest integer such that $T^{j+1}v = 0$. Then $T^n(T^{j-n}v) \neq 0$ but $T^{n+1}(T^{j-n}v) = 0$, contradicting Proposition 2.1. □

Lemma 6.1. *Suppose $U_1, ..., U_m$ are subspaces of V. Then $\sum_{i=1}^{k} U_i$ is a direct sum if and only if*

$$\dim\left(\sum_{i=1}^{k} U_i\right) = \sum_{i=1}^{k} \dim(U_i)$$

Example 6.1. *Consider $T \in \mathcal{L}(\mathbb{F}^4)$ defined by*

$$T(x_1, x_2, x_3, x_4) = (0, 2x_1, x_2, 4x_3).$$

Then $N(T^4) = \{(x_1, x_2, x_3, x_4) \in \mathbb{F}^4\}$ and $R(T^4) = \{(0, 0, 0, 0)\}$, satisfying $N(T^4) \oplus R(T^4) = \mathbb{F}^4$.

Corollary 6.1.

(a) *each $G(\lambda_j, T)$ is T-invariant*

(b) *each $(T - \lambda_j I)\,|_{G(\lambda_j, T)}$ is nilpotent*

Proof. Let $n = \dim(V)$. Recall that $G(\lambda_j, T) = N(T - \lambda_j I)^n$ for each j (Proposition 2.4). By Proposition 2.7, with $p(x) = (x - \lambda_j)^n$, we get (a). (b) follows from definitions. □

Proposition 6.3. *Suppose $T \in \mathcal{L}(V)$. Let $\lambda_1, ..., \lambda_m$ be distinct eigenvalues of T. Then*

$$V = \bigoplus_{j=1}^{m} G(\lambda_j, T).$$

Proof. Let $n = \dim(V)$. We use induction on n. For our base case, note that the result holds if $n = 1$. Now for our induction hypothesis assume that $n > 1$ and the result holds for all vector spaces of smaller dimension. Since our vector space is over an algebraically closed field, T has an eigenvalue; thus $m \geq 1$. Let

$$U = R(T - \lambda_1 I)^n.$$

Applying Proposition 1.12 to $(T - \lambda_1 I)$ shows that

$$V = G(\lambda_1, T) \oplus U \tag{3}$$

Using Proposition 1.7, with $p(x) = (x - \lambda_1)^n$, U is T-invariant. Because $G(\lambda_1, T) \neq \{0\}$, $\dim(U) < n$. Thus we can apply our induction hypothesis to $T\mid_U$. No generalized eigenvectors of $T\mid_U$ corresponding λ_1 are in U, because they are all in $G(\lambda_1, T)$. Thus each eigenvalue of $T\mid_U$ is in $\{\lambda_1, ..., \lambda_m\}$. By the induction hypothesis, $U = \bigoplus_{j=2}^{m} G(\lambda_j, T\mid_U)$. Thus it suffices to prove that $G(\lambda_k, T\mid_U) = G(\lambda_k, T)$ for $k = 2, ..., m$.

Fixing $k \in \{2, ..., m\}$, the inclusion $G(\lambda_k, T\mid_U) \subseteq G(\lambda_k, T)$ is clear. To prove the other direction, suppose $v \in G(\lambda_k, T)$. For $v_1 \in G(\lambda_1, T)$ and $u \in U$, we can write $v = v_1 + u$ using equation (3). By our induction hypothesis, $u = \sum_{j=2}^{m} v_j$ where each $v_j \in G(\lambda_j, T\mid_U) \subseteq G(\lambda_j, T)$. Thus $v = \sum_{j=1}^{m} v_j$. This equation, in combination with Proposition 1.4, implies that each v_j equals 0 unless $j = k$. Thus $v_1 = 0$ and $v = u \in U$, hence we can conclude that $v \in G(\lambda_k, T\mid_U)$. □

Definition 6.1. *Suppose $T \in \mathcal{L}(V)$ and let λ be an eigenvalue of T. We call $d = \dim(G(\lambda, T))$ the **algebraic multiplicity** of λ.*

Theorem (Cayley-Hamilton #1). *Suppose $T \in \mathcal{L}(V)$. Let $q(x)$ denote the characteristic polynomial of T. Then $q(T) = 0$.*

Proof. Let $\lambda_1, ..., \lambda_m$ be the distinct eigenvalues of T, and $d_1, ..., d_m$ the dimensions of the corresponding generalized eigenspaces $G(\lambda_1, T), ..., G(\lambda_m, T)$. For each $j \in \{1, ..., m\}$, corollary 6.1 states that $(T - \lambda_j I)\mid_{G(\lambda_j, T)}$ is nilpotent. Thus by Proposition 1.13,

$$(T - \lambda_j I)^{d_j}\mid_{G(\lambda_j, T)} = 0 \tag{4}$$

By Proposition 6.3, V can be decomposed as a direct sum of generalized eigenspaces, hence for every $v \in V$ there exist unique vectors $v_i \in G(\lambda_i, T)$ such that $v = \sum_{i=1}^{m} v_i$. Then to prove that $q(T) = 0$, it suffices to prove $q(T)\mid_{G(\lambda_j, T)} = 0$ for each j. We have

$$q(T) = \prod_{j=1}^{m} (T - \lambda_j I)^{d_j}.$$

Restricting both sides of the equation to $G(\lambda_j, T)$ and applying (4) proves our result.

$$q(T)\mid_{G(\lambda_j, T)} = \prod_{j=1}^{m} (T - \lambda_j I)^{d_j}\mid_{G(\lambda_j, T)} = 0$$

□

This definition lets us give an equivalent formulation of the Cayley-Hamilton theorem.

Theorem (Cayley-Hamilton #2). *Suppose $T \in \mathcal{L}(V)$. Let $q(x)$ and $p(x)$ denote the characteristic polynomial and minimal polynomial of T, respectively. Then $p(x)$ divides $q(x)$.*

Proof. We prove this equivalent to Cayley-Hamilton #1.

(#1 $\implies$ #2) Suppose that $q(T) = 0$. Using the division algorithm for polynomials, there exist $s(x), r(x) \in \mathbb{F}[x]$ such that

$$q(x) = p(x)s(x) + r(x)$$

and $\deg(r(x)) < \deg(p(x))$. Then

$$0 = q(T) = p(T)s(T) + r(T) = r(T)$$

This equation implies $r(x) = 0$; otherwise, $r(x)$ divided by its highest-degree coefficient would be a monic polynomial of smaller degree than $p(x)$ that when applied to T gives 0, a contradiction. Thus $p(x)$ divides $q(x)$.

(#2 $\implies$ #1) Suppose that $q(x)$ is a polynomial multiple of the minimal polynomial. Then there exists $s(x) \in \mathbb{F}[x]$ such that $q(x) = p(x)s(x)$. Thus

$$q(T) = p(T)s(T) = 0$$

□

REFERENCES

[Gal86] Joseph Gallian. “Contemporary Abstract Algebra”. In: 1986.

[Hal86] Marshall Hall. *Combinatorial Theory*. 1986.

[FIS97] Stephen H. Friedberg, Arnold J. Insel, and Lawrence E. Spence. *Linear algebra*. Third. Prentice Hall, Inc., Upper Saddle River, NJ, 1997, pp. xiv+557. ISBN: 0-13-233859-9.

[DF04] David S. Dummit and Richard M. Foote. *Abstract algebra*. Third. John Wiley & Sons, Inc., Hoboken, NJ, 2004, pp. xii+932. ISBN: 0-471-43334-9.

[Rom08] Steven Roman. *Advanced Linear Algebra*. 2008.

[Art10] Michael Artin. *Algebra*. 2010.

[SGB14] John Sheekey, Rod Gow, and Andries Brouwer. “Counting Symmetric Nilpotent Matrices”. In: *Electronic Journal of Combinatorics* 21 (Apr. 2014). DOI: 10.37236/3810.

[Axl15] Sheldon Axler. *Linear algebra done right*. Third. Undergraduate Texts in Mathematics. Springer, Cham, 2015, pp. xviii+340. ISBN: 978-3-319-11079-0; 978-3-319-11080-6. DOI: 10.1007/978-3-319-11080-6. URL: https://doi.org/10.1007/978-3-319-11080-6.