MATH 210A, FALL 2017

HW 6 SOLUTIONS

WRITTEN BY DAN DORE

(If you find any errors, please email ddore@stanford.edu)

Question 1. Let A and B be two $n \times n$ matrices with entries in a field K.
Let L be a field extension of K, and suppose there exists $C \in \mathrm{GL}_n(L)$ such that $B = CAC^{-1}$.
Prove there exists $D \in \mathrm{GL}_n(K)$ such that $B = DAD^{-1}$.
(that is, turn the argument sketched in class into an actual proof)

Solution. We'll start off by proving the existence and uniqueness of rational canonical forms in a precise way. To do this, recall that the *companion matrix* for a monic polynomial $p(t) = t^d + a_{d-1}t^{d-1} + \cdots + a_0 \in K[t]$ is defined to be the $(d-1) \times (d-1)$ matrix

$$M(p) := \begin{pmatrix} 0 & 0 & \cdots & 0 & -a_0 \\ 1 & 0 & \cdots & 0 & -a_1 \\ 0 & 1 & \cdots & 0 & -a_2 \\ \vdots & & \ddots & & \vdots \\ 0 & 0 & \cdots & 1 & -a_{d-1} \end{pmatrix}$$

This is the matrix representing the action of t in the cyclic $K[t]$-module $K[t]/(p(t))$ with respect to the ordered basis $1, t, \dots, t^{d-1}$.

Proposition 1 (Rational Canonical Form)**.** For any matrix M over a field K, there is some matrix $B \in \mathrm{GL}_n(K)$ such that

$$BMB^{-1} \simeq M_{\text{can}} := M(p_1) \oplus M(p_2) \oplus \cdots \oplus M(p_m)$$

Here, $p_1, \dots, p_n$ are monic polynomials in $K[t]$ with $p_1 \mid p_2 \mid \cdots \mid p_m$. The monic polynomials p_i are uniquely determined by M,[1] so if for some $C \in \mathrm{GL}_n(K)$ we have $CMC^{-1} = M(q_1) \oplus \cdots \oplus M(q_k)$ for $q_1 \mid q_1 \mid \cdots \mid q_m \in K[t]$, then $m = k$ and $q_i = p_i$ for each i.

Moreover, if $\iota\colon K \hookrightarrow L$ is a field extension, then if we let $\iota(M)$ denote the matrix M considered as a matrix with coefficients in L, we have $\iota(M)_{\text{can}} = \iota(M_{\text{can}})$. In other words, the rational canonical form does not depend on which field we consider the coefficients of M to lie inside.

Before we give the proof, let's see why this implies the statement of the question. By the proposition, there exist matrices $\alpha, \beta \in \mathrm{GL}_n(K)$ such that $A_{\text{can}} = \alpha A \alpha^{-1}$ and $B_{\text{can}} = \beta B \beta^{-1}$. But then, letting $\iota\colon K \hookrightarrow L$ be the field extension, we have:

$$\iota(B_{\text{can}}) = \iota(\beta)\iota(B)\iota(\beta)^{-1} = (\iota(\beta)C)\iota(A)(C^{-1}\iota(\beta)^{-1})$$

Let $\gamma = (\iota(\beta)C) \in \mathrm{GL}_n(L)$, so we have $\iota(B_{\text{can}}) = \gamma\iota(A)\gamma^{-1}$. By the uniqueness part of the proposition, this implies that $\iota(B_{\text{can}}) = \iota(A)_{\text{can}}$. By the part of the proposition regarding field extensions, we have $\iota(A)_{\text{can}} = \iota(A_{\text{can}})$, so $\iota(A_{\text{can}}) = \iota(B_{\text{can}})$. Since ι is injective, this implies that $A_{\text{can}} = B_{\text{can}}$, i.e. that $\alpha A \alpha^{-1} = \beta B \beta^{-1}$. Letting $D = \beta\alpha^{-1} \in \mathrm{GL}_n(K)$, this shows us that $B = DAD^{-1}$.

Now, let's prove the proposition:

[1] In particular, $\prod_i p_i$ is the characteristic polynomial of M and p_m is the minimal polynomial of M.

Proof. This follows from the invariant factor form of the structure theorem for finitely generated modules over a principal ideal domain. Namely, if M is a $n \times n$ matrix, we may regard it as a linear transformation on $V := K^n$ with respect to the standard ordered basis of K^n. Defining $t \cdot v = Mv$, we may give V the structure of a $K[t]$-module. V is finitely generated as a $K[t]$-module (since it is finite-dimensional as a K-vector space), so by the invariant factor form of the structure theorem for finitely generated modules over a principal ideal domain, there is an isomorphism

$$\varphi\colon K[t]^r \oplus K[t]/(p_1(t)) \oplus K[t]/(p_2(t)) \oplus \cdots \oplus K[t]/(p_m(t)) \xrightarrow{\sim} V \tag{1}$$

for $p_1, \ldots, p_m \in K[t]$ with $p_1 \mid p_2 \mid \cdots \mid p_m$, and this isomorphism is as a t-module. Note that $r = 0$, so we may omit the $K[t]^r$ factor: $K[t]$ has infinite dimension as a K-vector space, but V has finite dimension n. Furthermore, this theorem tells us that the principal ideals (p_i) are uniquely determined by the $K[t]$-module V (i.e. they are uniquely determined by M). Thus, the p_i are unique up to multiplication by a unit of $K[t]$. We know that these units are exactly the non-zero elements of K, so if we add the requirement that the p_i are *monic*, they are uniquely determined as elements of $K[t]$.

If $d_i = \deg p_i$ and $e_1, \ldots, e_m$ are the images of $1 \in K[t]/(p_i(t))$ under the canonical inclusions into the direct sum, then $\mathbf{b} := \{e_1, t \cdot e_1, \ldots, t^{d_1 - 1} \cdot e_1, e_2, \ldots, t^{d_2 - 1} \cdot e_2, \ldots, e_m, \ldots, t^{d_m - 1} \cdot e_m\}$ is an ordered basis for the left-hand side. Since $M(p_i)$ is the matrix of the linear transformation t acting on $K[t]/(p_i(t))$ with respect to the ordered basis $1, \ldots, t^{d_i - 1}$, we see that the matrix of the linear transformation t acting on the left-hand-side of (1) with respect to $\mathbf{b}$ is M_{can}.

Now, since φ is an isomorphism of $K[t]$-modules, we see that $\mathbf{B} := \varphi(\mathbf{b})$ is an ordered basis of V and that the matrix of t with respect to $\mathbf{B}$ is M_{can}. Let $\mathbf{f}$ be the standard ordered basis of $K^n \simeq V$, and define a matrix (B_{ij}) by the equations $\mathbf{f}_i = \sum_{j=1}^n B_{ji} \mathbf{B}_j$, which are uniquely determined because $\mathbf{B}$ is a basis. Since $\mathbf{f}$ is also a basis, we may also uniquely write $\mathbf{B}_j = \sum_{k=1}^n B'_{kj} \mathbf{f}_k$, so this tells us that

$$\mathbf{f}_j = \sum_{i=1}^n \sum_{k=1}^n B'_{ki} B_{ij} \mathbf{f}_k$$

Since $\mathbf{f}$ is a basis, we may compare coefficients to see that $B' \cdot B = I$, so $B \in \mathrm{GL}_n(K)$. Note that $B \cdot B' = I$ as well, i.e. $B' = B^{-1}$. This says that:

$$\mathbf{f}_j = BB'\mathbf{f}_j = B\left(\sum_{k=1}^n B'_{kj} \mathbf{f}_k\right) = B\mathbf{B}_j$$

We want to show that $M_{\text{can}} = BMB^{-1}$. This says that in the standard ordered basis $\mathbf{f}$, $BMB^{-1}(\mathbf{f}_j) = \sum_{i=1}^n (M_{\text{can}})_{ij} \mathbf{f}_i$.

We have, using the fact that M_{can} is the matrix for M with respect to the basis $\mathbf{B}$:

$$\begin{aligned} BMB^{-1}(\mathbf{f}_j) &= BMB'(\mathbf{f}_j) \\ &= BM\left(\sum_{k=1}^{n} B'_{kj}\mathbf{f}_k\right) \\ &= BM(\mathbf{B}_j) \\ &= B\left(\sum_i (M_{\text{can}})_{ij}\mathbf{B}_i\right) \\ &= \sum_i (M_{\text{can}})_{ij} B(\mathbf{B}_i) \\ &= \sum_i (M_{\text{can}})_{ij}\mathbf{f}_i \end{aligned}$$

Now, assume that we have some $C \in \mathrm{GL}_n(K)$ with $CMC^{-1} = N := M(q_1) \oplus \cdots \oplus M(q_\ell)$ with $q_1 \mid q_2 \mid \cdots \mid q_m$ and q_i monic for all i. Let W be the $K[t]$ module K^n with t acting by the matrix N with respect to the standard basis $\mathbf{f}$ of K^n. Thus, $W \simeq K[t]/(q_1(t)) \oplus \cdots \oplus K[t]/(q_\ell(t))$. We consider C as an homomorphism $V \xrightarrow{\sim} W$, sending $\mathbf{f}_j$ to $\sum_i C_{ij}\mathbf{f}_i$. Then we have $C(t \cdot v) = CMv = NCv = t \cdot (C(v))$ for all $v \in V$, so C is an homomorphism of $K[t]$-modules. Since $C \in \mathrm{GL}_n(K)$, it is bijective and thus an isomorphism of $K[t]$-modules. Thus, we have an isomorphism:

$$K[t]/(p_1(t)) \oplus \cdots \oplus K[t]/(p_m(t)) \simeq K[t]/(q_1(t)) \oplus \cdots \oplus K[t]/(q_\ell(t))$$

Now, we may apply the uniqueness of the invariant factor form of the structure theorem for finitely generated modules over a principal ideal domain to conclude that $\ell = m$ and $(p_i) = (q_i)$ as principal ideals. Thus, q_i and p_i differ by multiplication by a non-zero element of K, and since they are both assumed to be monic, we have $p_i = q_i$.

Finally, we need to show that the rational canonical form is preserved by field extensions. This is a consequence of the above existence and uniqueness statements. Let $\iota\colon K \hookrightarrow L$ be a field extension and consider the matrix $\iota(M)$ with coefficients in L. By the existence of rational canonical form, we know that there is a matrix $B \in \mathrm{GL}_n(K)$ with $BMB^{-1} = M_{\text{can}} = \bigoplus_i M(p_i)$ with $p_1 \mid \cdots \mid p_m$ and $p_i \in K[t]$ monic. Applying ι, we have $\iota(B)\iota(M)\iota(B)^{-1} = \iota(M_{\text{can}}) = \bigoplus_i \iota(M(p_i)) = \bigoplus_i M(\iota(p_i))$, where $\iota(p_i)$ is the polynomial $p_i \in K[t]$ viewed as a polynomial with coefficients in L. Now, $\iota(p_i) = q_i$ is a monic polynomial in $L[t]$, and the condition that $p_i \mid p_{i+1}$, i.e. $p_{i+1} = p_i \cdot f_i$ with $f_i \in K[t]$, is also preserved by ι (since $\iota(p_{i+1}) = \iota(p_i) \cdot \iota(f_i)$ and $\iota(f_i) \in L[t]$). Thus, $\iota(M_{\text{can}})$ is conjugate to $\iota(M)$ by $\iota(B) \in \mathrm{GL}_n(L)$, and it is of the form $\bigoplus_{i=1}^{\ell} M(q_i)$ with q_i monic and $q_1 \mid \cdots \mid q_\ell$. Thus, by the above uniqueness statement, $\iota(M_{\text{can}})$ is in rational canonical form, so $\iota(M_{\text{can}}) = \iota(M)_{\text{can}}$.

Note that this invariance under field extensions is not true for other normal forms of matrices, such as the normal form obtained by using the elementary divisor version of the structure theorem for principal ideal domains: if $\pi(t)$ is irreducible as an element of $K[t]$ but not of $L[t]$, then the matrix corresponding to the $K[t]$-module $K[t]/(\pi^e)$ is not in its "elementary divisor form" as a matrix with coefficients in L. $\square$

Question 2. Let k be an algebraically closed field of characteristic $\neq 2$. Fix two nonzero elements $\lambda, \mu \in k$. Let V and W be 2-dimensional k-vector spaces. Let $\alpha\colon V \to V$ have matrix $\left(\begin{smallmatrix} \lambda & 0 \\ 1 & \lambda \end{smallmatrix}\right)$, and let $\beta\colon W \to W$ have matrix $\left(\begin{smallmatrix} \mu & 0 \\ 1 & \mu \end{smallmatrix}\right)$. Let $\gamma\colon V \otimes W \to V \otimes W$ be $\alpha \otimes \beta$ (defined on elementary tensors by $\gamma(v \otimes w) = \alpha(v) \otimes \beta(w)$).

Find the Jordan decomposition of γ (that is, give a list of blocks and their sizes, and prove your answer is correct). Give a basis for all eigenspaces of γ. What happens if char $k = 2$?

Solution.

Let e_1, e_2 be the standard basis for V and f_1, f_2 the standard basis for W. Then (since tensor product commutes with direct sum), $V \otimes W$ is spanned by the vectors $e_1 \otimes f_1, e_1 \otimes f_2, e_2 \otimes f_1, e_2 \otimes f_2$. In the e_1, e_2 basis, α acts by:

$$\alpha(e_2) = \lambda e_2, \qquad \alpha(e_1) = \lambda e_1 + e_2$$

and likewise, β acts by:

$$\beta(f_2) = \mu f_2, \qquad \beta(f_1) = \mu f_1 + f_2$$

Now, we can compute:

$$\begin{aligned}
\gamma(e_1 \otimes f_1) &= \big(\alpha(e_1)\big) \otimes \big(\beta(f_1)\big) \\
&= (e_2 + \lambda e_1) \otimes (f_2 + \mu f_1) \\
&= (\lambda\mu)(e_1 \otimes f_1) + \lambda(e_1 \otimes f_2) + \mu(e_2 \otimes f_1) + e_2 \otimes f_2
\end{aligned}$$

$$\begin{aligned}
\gamma(e_1 \otimes f_2) &= \big(\alpha(e_1)\big) \otimes \big(\beta(f_2)\big) \\
&= (e_2 + \lambda e_1) \otimes (\mu f_2) \\
&= (\lambda\mu)(e_1 \otimes f_2) + \mu(e_2 \otimes f_2)
\end{aligned}$$

$$\begin{aligned}
\gamma(e_2 \otimes f_1) &= \big(\alpha(e_2)\big) \otimes \big(\beta(f_1)\big) \\
&= (\lambda e_2) \otimes (f_2 + \mu f_1) \\
&= (\lambda\mu)(e_2 \otimes f_1) + \lambda(e_2 \otimes f_2)
\end{aligned}$$

$$\begin{aligned}
\gamma(e_2 \otimes f_2) &= \big(\alpha(e_2)\big) \otimes \big(\beta(f_2)\big) \\
&= (\lambda e_2) \otimes (\mu f_2) \\
&= (\lambda\mu)(e_2 \otimes f_2)
\end{aligned}$$

Thus, the matrix in the ordered basis $\{(e_1 \otimes f_1), (e_1 \otimes f_2), (e_2 \otimes f_1), (e_2 \otimes f_2)\}$ is:

$$\gamma = \begin{pmatrix} \lambda\mu & 0 & 0 & 0 \\ \lambda & \lambda\mu & 0 & 0 \\ \mu & 0 & \lambda\mu & 0 \\ 1 & \mu & \lambda & \lambda\mu \end{pmatrix}$$

We can see directly that the characteristic polynomial of this matrix is $(T - \lambda\mu)^4$, so $\lambda\mu$ is the only eigenvalue.

To find the Jordan decomposition, we need to look at the generalized eigenspaces for $\lambda\mu$. Consider the matrix for $\gamma - \lambda\mu$:

$$\gamma - \lambda\mu = \begin{pmatrix} 0 & 0 & 0 & 0 \\ \lambda & 0 & 0 & 0 \\ \mu & 0 & 0 & 0 \\ 1 & \mu & \lambda & 0 \end{pmatrix}$$

This matrix clearly has rank 2, since the bottom row is linearly dependent from the other rows and the middle two rows are nonzero scalar multiples of each other. The kernel, which is two-dimensional, is the eigenspace for $\lambda\mu$. We can take $\{\lambda(e_1 \otimes f_2) - \mu(e_2 \otimes f_1) + e_2 \otimes f_2, e_2 \otimes f_2\}$ as a basis. Since $\lambda\mu$ is the only eigenvalue of γ and its eigenspace is two-dimensional, the Jordan decomposition of γ must consist of exactly two Jordan blocks, both of which have eigenvalue $\lambda\mu$. Either both blocks have dimension 2 or one block has dimension 3 and the other has dimension 1. To decide which, consider $(\gamma - \lambda\mu)^2$:

$$(\gamma - \lambda\mu)^2 = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 2\mu\lambda & 0 & 0 & 0 \end{pmatrix}$$

If $\text{char}(k) \neq 2$, then this is non-zero. But $(\gamma - \lambda\mu)^2$ annihilates a Jordan block of size 2 and eigenvalue $\lambda\mu$, so in this case there must be a Jordan block of size 3 and another of size 1. On the other hand, if $\text{char}(k) = 2$, then $(\gamma - \lambda\mu)^2 = 0$, but $(\gamma - \lambda\mu)^2$ does not vanish on a Jordan block of size 3 and eigenvalue $\lambda\mu$; thus, in this case the Jordan decomposition for γ consists of two Jordan blocks of size 2.

We can see the bases for the generalized eigenspaces explicitly:

(a) The case when $\text{char}(k) \neq 2$:

Let $w = e_1 \otimes f_1$, which satisfies $(\gamma - \lambda\mu)^3(w) = 0$ but $(\gamma - \lambda\mu)^2(w) = 2\mu\lambda(e_2 \otimes f_2) \neq 0$. Thus, w generates a cyclic subspace of maximal dimension 3, spanned by $w, (\gamma - \lambda\mu)w, (\gamma - \lambda\mu)^2 w$. (we can see that these are linearly independent easily from the matrices for $(\gamma - \lambda\mu)$ and $(\gamma - \lambda\mu)^2$).

In terms of the ordered basis $w, (\gamma - \lambda\mu)w, (\gamma - \lambda\mu)^2 w$ for this subspace, we see that $\gamma = (\gamma - \lambda\mu) + \lambda\mu$ takes the form:

$$\gamma = \begin{pmatrix} \lambda\mu & 0 & 0 \\ 1 & \lambda\mu & 0 \\ 0 & 1 & \lambda\mu \end{pmatrix}$$

Now, consider the eigenvector $v = \lambda(e_1 \otimes f_2) - \mu(e_2 \otimes f_1) + e_2 \otimes f_2$. We have seen that this is an eigenvector for γ. In the ordered basis $\{(e_1 \otimes f_1), (e_1 \otimes f_2), (e_2 \otimes f_1), (e_2 \otimes f_2)\}$, we can write:

$$w = \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \quad (\gamma - \lambda\mu)w = \begin{pmatrix} 0 \\ \lambda \\ \mu \\ 1 \end{pmatrix}, \quad (\gamma - \lambda)^2 w = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 2\mu\lambda \end{pmatrix}, \quad v = \begin{pmatrix} 0 \\ \lambda \\ -\mu \\ 1 \end{pmatrix}$$

We can see that these are linearly independent, because $(\lambda, -\mu)$ is not a scalar multiple of (λ, μ) as $\lambda, \mu \neq 0$ and $\text{char}(k) \neq 2$. Thus, these four elements give a basis of $V \otimes W$ with respect to which γ takes its Jordan normal form.

(b) The case when $\mathrm{char}(k) = 2$:

We know that there are two Jordan blocks with size 2 and eigenvalue $\lambda\mu$, and we already have a basis for the $\lambda\mu$-eigenspace, so we need to find two linearly independent vectors w_1, w_2 with $(\gamma - \lambda\mu)v \neq 0, (\gamma - \lambda\mu)w \neq 0$.

We can let $w_1 = e_1 \otimes f_1$, since $(\gamma - \lambda\mu)w_1 = \lambda(e_1 \otimes f_2) + \mu(e_2 \otimes f_1) + e_2 \otimes f_2 \neq 0$. Note that $(\gamma - \lambda\mu)w_1 = \lambda(e_1 \otimes f_2) - \mu(e_2 \otimes f_1) + e_2 \otimes f_2 = v$, the eigenvector considered above. We can take $w_2 = \mu^{-1}(e_1 \otimes f_2$. This is clearly linearly independent from w_1, and we have $(\gamma - \lambda\mu)w_2 = e_2 \otimes f_2$, the other eigenvector considered above.

Thus, γ takes its Jordan normal form with respect to the ordered basis

$$\{e_1 \otimes f_1, \lambda(e_1 \otimes f_2) + \mu(e_2 \otimes f_1) + e_2 \otimes f_2, \mu^{-1}(e_1 \otimes f_2), e_2 \otimes f_2\}$$

Question 3. Let k be an algebraically closed field of characteristic $\neq 3$. Let V be an n-dimensional k-vector space, and suppose that $T\colon V \to V$ has minimal polynomial $(t-\lambda)^n$ for some nonzero $\lambda \in k$. Find the Jordan decomposition of T^3.

Solution. Consider the minimal polynomial $m(t)$ of T^3. Because k is algebraically closed, $m(t) = \prod_i (t - \mu_i)^{e_i}$ for some $\mu_i \in k$, and furthermore $\mu_i = \lambda_i^3$ for some $\lambda_i \in k$. We can write:

$$0 = m(T^3) = \prod_i (T^3 - \lambda_i^3)^{e_i} = \prod_i \Big((T - \lambda_i)^{e_i} \cdot (T - \omega\lambda_i)^{e_i} \cdot (T - \omega^2\lambda_i)^{e_i} \Big)$$

where ω is a primitive cube root of unity, i.e. $\omega \neq 1$ but $\omega^3 = 1$ or equivalently ω is a root of the polynomial $t^2 + t + 1 = \frac{(t-1)^3}{(t-1)}$. Now, since the minimal polynomial of T is $(t-\lambda)^n$, the only eigenvalue of T is λ, so if $\mu \neq \lambda$ is any element of k, $(T - \mu)$ has trivial kernel and thus is invertible. Thus, if $\lambda_i^3 \neq \lambda$, none of $\lambda_i, \omega\lambda_i, \omega^2\lambda_i$ are equal to λ, so $(T^3 - \lambda_i^3)^{e_i} = (T - \lambda_i)^{e_i} \cdot (T - \omega\lambda_i)^{e_i} \cdot (T - \omega^2\lambda_i)^{e_i}$ is invertible. Thus, λ_i is not an eigenvalue for T^3, so this term cannot appear in $m(t)$.

Therefore, $m(T^3) = (T^3 - \lambda^3)^m$ for some $m \leq n$, i.e. the only eigenvalue for T^3 is λ^3. Thus, the dimension of the λ^3-eigenspace of T^3 is equal to the number of Jordan blocks for T^3. Now, let v be an eigenvector for T^3. Then we have:

$$0 = (T^3 - \lambda^3)v = (T - \omega\lambda)(T - \omega^2\lambda)(T - \lambda)v$$

But since the only eigenvalue of T is λ, $(T - \omega\lambda)$ and $(T - \omega^2\lambda)$ are both invertible, so this implies that $(T - \lambda)v = 0$, i.e. that v is an eigenvector for T. But since the minimal polynomial of T is $(t-\lambda)^n$, the Jordan decomposition for T consists of a single Jordan block of size n, so the λ-eigenspace for T is one-dimensional. Thus, the λ^3-eigenspace for T^3 is one-dimensional, so T^3 has Jordan decomposition consisting of a single Jordan block of eigenvalue λ^3 and size n.

Note that the hypothesis that the characteristic of k is not equal to 3 is essential. If k is any field of characteristic 3, then the matrix $T = \left(\begin{smallmatrix} 1 & 1 \\ 0 & 1 \end{smallmatrix}\right)$ is a Jordan block and it has minimal polynomial $(t-1)^2$. However, $T^3 = \left(\begin{smallmatrix} 1 & 3 \\ 0 & 1 \end{smallmatrix}\right) = \left(\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix}\right)$ is the identity matrix, so this has two Jordan blocks of size 1.

Question 4. Let V be a finite-dimensional nonzero vector space over a field k.

(a) For each monic irreducible $\pi \in k[t]$, define

$$V(\pi) = \{\, v \in V \mid \exists k \in \mathbb{N} \text{ s.t. } (\pi(T))^k(v) = 0 \,\}.$$

(When k is algebraically closed, these are the *generalized eigenspaces* $V_\lambda = V(t-\lambda)$ of T.)

Prove that $V(\pi) \neq 0$ if and only if $\pi | m_T$, and that $V = \bigoplus_{\pi | m_T} V(\pi)$.

An endomorphism $T\colon V \to V$ is *semisimple* if every T-stable subspace of V admits a T-stable complementary subspace: i.e. for every $T(U) \subseteq U$ there exists a decomposition $V = U \oplus W$ with $T(W) \subseteq W$. (Keep in mind that such a complement is not unique in general; e.g. consider scalar multiplication by 2.)

(b) Use rational canonical form to prove that T is semisimple if and only if m_T has no repeated irreducible factor over k. (Hint: apply (a) to T-stable subspaces of V to reduce to the case when m_T has one monic irreducible factor.) Deduce that

 (i) a Jordan block of rank > 1 is never semisimple,

 (ii) if T is semisimple then m_T is the "squarefree part" of χ_T, and

 (iii) if T is semisimple and $U \subseteq V$ is a T-stable nonzero proper subspace then the induced endomorphisms $T_U\colon U \to U$ and $\overline{T}\colon V/U \to V/U$ are semisimple.

(c) Let V' be another nonzero finite-dimensional k-vector space, and let $T'\colon V' \to V'$ be another endomorphism. Prove that T and T' are semisimple if and only if the endomorphism $T \oplus T'$ of $V \oplus V'$ is semisimple.

Solution. (a) Consider the unique factorization $m_T = \prod_i \pi_i^{e_i}$ into products of powers of pairwise distinct monic irreducibles $\pi_i \in k[t]$, and let $f_1 | \dots | f_n$ be the invariant factors for (V, T) from the rational canonical form, so $f_n = m_T$ and hence for each j we have $f_j = \prod_i \pi_i^{e_{j,i}}$ where $0 \le e_{1,i} \le e_{2,i} \le \dots \le e_{n,i}$. Using the Chinese Remainder Theorem and viewing V as an $F[t]$-module via letting t act as T, shuffling terms in direct sums (and dropping the vanishing quotients $k[t]/(\pi_i^{e_{j,i}})$ with $e_{j,i} = 0$) gives $k[t]$-linear isomorphisms

$$V \simeq \bigoplus_j (k[t]/(f_j)) = \bigoplus_j \left(\bigoplus_{i | e_{j,i} > 0} (k[t]/(\pi_i^{e_{j,i}})) \right)$$
$$= \bigoplus_i \left(\bigoplus_{j | e_{j,i} > 0} (k[t]/(\pi_i^{e_{j,i}})) \right).$$

Although the initial isomorphism is *not* unique (just as choosing an eigenbasis is not unique when one exists), we shall now show that the subspaces of V corresponding to the summands

$$\bigoplus_{j | e_{j,i} > 0} (k[t]/(\pi_i^{e_{j,i}}))$$

under this composite isomorphism are exactly the $V(\pi_i)$ (which shows that they are *intrinsic* to (V, T)), and that $V(\Pi) = 0$ for any monic irreducible polynomial $\Pi \ne \pi_i$ for any i (i.e. $\Pi \nmid m_T$).

Going across the chain of $k[t]$-linear isomorphisms (which converts the action of $h(T)$ on V for $h \in k[t]$ over to the multiplication operator by h on the various cyclic-module quotients of $k[t]$), we just have to show that if π and Π are distinct monic irreducibles in $k[t]$ then (i) multiplication by π^N kills $k[t]/(\pi^e)$ for any $N \ge e > 0$, and (ii) multiplication by Π on $k[t]/(\pi^e)$ is an k-linear automorphism. This shows that $V(\Pi) = 0$ for $\Pi \ne \pi_i$ for any i, and that $V(\pi_i)$ is exactly the direct summand of V corresponding to $\oplus_{j | e_{j,i} > 0} (k[t]/(\pi_i^{e_{j,i}}))$.

The first of these two claims is obvious, and for the second we just have to prove injectivity. But injectivity is a very concrete statement: if $g \in k[t]$ and $\Pi \cdot g$ is divisible by π^e then g is divisible by π^e. Since Π and π are distinct *monic* irreducibles, this follows from consideration of the "unique" (up to scalars and rearrangement) factorization of $\Pi \cdot g$ into irreducibles; the trivial case $g = 0$ is treated separately.

Remark 2. Beware that if we work in terms of matrices and then consider the situation over an extension field F' of F, the irreducibles π_i in $F[t]$ may be reducible in $F'[t]$ (unless π_i is linear and so leaves no room to factor). Hence, the formation of the π_i-primary subspaces $V(\pi_i)$ is generally destroyed by extension of the coefficient field except in the case that all π_i's are linear, which is to say that m_T factors completely into linears in $F[t]$ (or the same for χ_T, since m_T and χ_T have the same monic irreducible factors in $F[t]$ in general).

Since m_T and χ_T have the same monic irreducible factors, the monic linear factors of χ_T are precisely the polynomials $t - \lambda$ for which $\chi_T(\lambda) = 0$, or in other words for which $\lambda \in F$ is an eigenvalue for T acting on V. Thus, the i for which the monic irreducible π_i is linear are precisely those of the form $t - \lambda$ with λ an eigenvalue in F for T acting on V.

(b) Before proving the first part, we consider the general structure of T-stable subspaces $W \subseteq V$. Using the primary decomposition as in part (a), if $m_T = \prod \pi^{e_\pi}$ is the monic irreducible factorization of m_T then we have a T-stable decomposition $V = \oplus_\pi V(\pi)$ as a direct sum of its primary components (with $V(\pi)$ killed by a high power of $\pi(T)$ and acted upon with vanishing kernel by $\Pi(T)$ for all monic irreducible $\Pi \neq \pi$ in the factorization of m_T). If $W \subseteq V$ is a nonzero T-stable subspace then since $m_T(T|_W)$ on W is the restriction of $m_T(T) = 0$, we have that $m_{T|_W}$ divides m_T in $k[t]$, so $m_{T|_W}$ has its monic irreducible factors given among the various π's in m_T (with multiplicity at most e_π, if it occurs at all in $m_{T|_W}$). Hence, we have compatible primary decompositions

$$W = \bigoplus_\pi W(\pi) \subseteq \bigoplus_\pi V(\pi) = V$$

where it is understood that $W(\pi)$ means 0 if π does not actually divide $m_{T|_W}$ (or equivalently, if $\pi(T)$ acts as an automorphism on W).

The upshot is this: if there is to be a T-stable complement W' to W in V then *necessarily* $W' = \oplus_\pi W'(\pi)$ as well, so the equality $W \oplus W' = V$ says exactly that $W(\pi) \oplus W'(\pi) = V(\pi)$ for all π. Since $V(\pi)$ is killed by a high power of $\pi(T)$, we conclude that W has a T-stable complement in V if and only if $W(\pi)$ has a T-stable complement in $V(\pi)$ for all π. Put another way: T acts semisimply on V if and only if it acts semisimply on every $V(\pi)$! This will reduce our problems to the primary components where things are easy to compute.

We will use the following fact:

Lemma 3. If π is any monic irreducible factor of m_T (or equivalently, of χ_T), then $\pi(T)$ kills $V(\pi)$ if and only if π appears without repetition in the irreducible factorization of m_T.

Proof. By construction, $V(\pi)$ considered as an $k[t]$-module is a direct sum of cyclic modules $k[t]/(\pi^e)$ for various positive exponents e that are exactly the multiplicities of π as an irreducible factor of the various invariant factors $f_1, \dots, f_n$ for T acting on V. Since $f_n = m_T$, it follows that π divides m_T only to first order if and only if π divides each of the f_j's to at most first order, which is to say that $V(\pi)$ is a direct sum of copies of $k[t]/(\pi)$ as an $k[t]$-module. Since $\pi(T)$ acting on $V(\pi)$ goes over to multiplication by π on the direct sum of various $k[t]/(\pi^e)$'s that arise from the monic irreducible factorizations of the f_j's, it follows that $\pi(T)$ kills $V(\pi)$ if and only if multiplication by π kills $k[t]/(\pi^e)$ as e ranges through the multiplicities of π that appear in the monic irreducible factorizations of the f_j's. Hence, our problem is reduced to the obvious fact that multiplication by π kills $k[t]/(\pi^e)$ with $e > 0$ if and only if $e = 1$. □

By Lemma 3, m_T has no repeated irreducible factors if and only if each $V(\pi)$ is killed by $\pi(T)$. But we have just seen that T is semisimple on V if and only if T acts semisimply on each $V(\pi)$. Hence, for the purposes of proving the theorem it is harmless to replace V and T with $V(\pi)$ and $T|_{V(\pi)}$ so as to reduce to the case when $m_T = \pi^e$ for some monic irreducible π and some $e > 0$. Clearly V as an $k[t]$-module may be viewed as a module over the quotient ring $k' = k[t]/(\pi^e)$. The T-stable subspaces are precisely the $k[t]$-submodules, or equivalently the k'-submodules. Hence, semisimplicity of T is equivalent to the statement that every k'-submodule of V has an k'-submodule complement. If m_T has no repeated factors then $e = 1$, so k' is a *field* and thus V is an k'-vector space with visibly finite dimension; the construction of complements to k'-subspaces is then clear by linear algebra over k'! Conversely, assuming T to be semisimple, if we let $W = \ker \pi(T)$ then there is an $k[t]$-submodule $W' \subseteq V$ complementary to W. Since $W \oplus W' = V$ but $W = \ker \pi(T)$, it follows that $\pi(T)$ acts injectively on W'. But $\pi(T)^e = 0$ on V, whence $W' = 0$. This gives $V = W = \ker \pi(T)$, so $\pi(T) = 0$. In other words, if T is semisimple then m_T has no repeated factors.

Now, let's show that this characterization of semisimplicity implies properties (i)-(iii):

(i) The minimal polynomial of a Jordan block of rank e and eigenvalue λ is $(t - \lambda)^e$, so if $e > 1$, this has repeated irreducible factors and is thus not semisimple. Note that in particular, using Jordan decomposition, a semisimple matrix over an algebraically closed field is exactly the same as a diagonalizable matrix.

(ii) The "squarefree part" of a monic polynomial is just the product of its monic irreducible factors taken without multiplicity. Since m_T and χ_T have the same irreducible factors, this follows immediately.

(iii) By the result we just proved, the minimal polynomial $m_T \in k[t]$ has no repeated irreducible factors. But $m_T(T) = 0$, and restricting this identity to W gives $m_T(T|_W) = 0$, so $m_{T|_W}$ divides m_T in $k[t]$. Hence, $m_{T|_W}$ has no repeated irreducible factors, so we conclude that $T|_W$ is semisimple. Likewise, $m_T(\overline{T})$ on V/W is induced by $m_T(T) = 0$, so $m_{\overline{T}}$ divides m_T in $k[t]$. Hence, as with $m_{T|_W}$, we conclude that $m_{\overline{T}}$ has no repeated irreducible factors, and so $\overline{T}$ on V/W is semisimple.

(c) The key is to show that $m_{T \oplus T'} = \operatorname{lcm}(m_T, m_{T'})$, since for any nonzero $g, h \in k[t]$ it is clear that $\operatorname{lcm}(g, h)$ has no repeated irreducible factors if and only if the same holds for g and h separately (and so the criterion for semisimplicity from part (b) then gives the desired result). Since $m_{T \oplus T'}(T \oplus T') = 0$ acts on $V \subseteq V \oplus V'$ as $m_{T \oplus T'}(T)$, it follows that $m_{T \oplus T'}(T) = 0$ on V; likewise, $m_{T \oplus T'}(T') = 0$ on V'. Hence, the monic $m_{T \oplus T'}$ is a multiple of both m_T and $m_{T'}$ in $k[t]$, so it is a multiple of the least common multiple of these two polynomials. For the reverse divisibility, if we let $h = \operatorname{lcm}(m_T, m_{T'})$ then we have to show that $h(T \oplus T')$ vanishes on $V \oplus V'$. But this restricts to $h(T)$ on V and $h(T')$ on V', so we have to prove the vanishing of these latter two operators. Since h is a multiple of both m_T and $m_{T'}$, we get the desired vanishing on V and V'.

Question 5. Let V be a n-dimensional k-vector space with $0 < n < \infty$, and let $T\colon V \to V$ be an endomorphism.

(a) Using rational canonical form and Cayley–Hamilton, prove the following are equivalent:

(1) $\exists k \geq 1$ such that $T^k = 0$.

(2) $T^n = 0$.

(3) There is an ordered basis of V w.r.t. which the matrix for T is upper triangular with 0's on the diagonal.

(4) $\chi_T = t^n$.

We call such T *nilpotent*.

(b) We say that T is *unipotent* if $T - 1$ is nilpotent. Formulate characterizations of unipotence analogous to the conditions in (a), and prove that a unipotent T is invertible.

(c) Assume k is algebraically closed. Using Jordan canonical form and generalized eigenspaces, prove that there is a unique expression

$$T = T_{\mathrm{ss}} + T_{\mathrm{n}}$$

where T_{ss} and T_{n} are a pair of *commuting* endomorphisms of V with T_{ss} semisimple and T_{n} nilpotent. (This is the *additive Jordan decomposition* of T.)

Show in general that $\chi_T = \chi_{T_{\mathrm{ss}}}$ (so T is invertible if and only if T_{ss} is invertible)

Show by example with $\dim V = 2$ that uniqueness fails if we drop the "commuting" requirement. (You just need to give the matrix T and the two decompositions $T = T_{\mathrm{ss}} + T_{\mathrm{n}}$ and $T = T'_{\mathrm{ss}} + T'_{\mathrm{n}}$; you do not need to prove these matrices are semisimple/nilpotent, as long as they are.)

(d) Assume k is algebraically closed and $S\colon V \to V$ is invertible. Using the existence and uniqueness of additive Jordan decomposition, prove that there is a unique expression

$$S = S_{\mathrm{ss}} S_{\mathrm{u}}$$

where S_{ss} and S_{u} are *commuting* endomorphisms of V with S_{ss} semisimple and S_{u} unipotent (so S_{ss} is necessarily invertible too). This is the *multiplicative Jordan decomposition* of T.

Solution. (a) If $T^r = 0$ for some r then the minimal polynomial m_T divides t^r, and so by unique factorization the monic polynomial m_T has only t as a monic irreducible factor. But χ_T and m_T have the same irreducible factors, so the monic χ_T with degree n must be t^n. By the Cayley-Hamilton theorem, we then get $T^n = 0$ in such cases. Hence, (1) implies (2). It is obvious that (2) implies (1). Of course, when these equivalent conditions hold, so $\chi_T = t^n$, it is clear that (4) holds. Conversely, if (4) holds then the only monic irreducible factor of χ_T is t (by uniqueness of prime factorization and the fact that t is irreducible). Thus, by the Cayley-Hamilton theorem, $T^n = 0$.

If (3) holds then it is clear that $\chi_T = t^n$ via the given matrix, so (4) holds. If (4) holds, so $\chi_T = t^n$, it follows that 0 is an eigenvalue of T. Let $L \subseteq V$ be a line in the 0-eigenspace, and consider the induced map $\overline{T}$ on V/L. Working with a basis of V adapted to L and V/L shows that $\chi_T = t\chi_{\overline{T}}$, so $\chi_{\overline{T}} = t^{n-1}$. Thus, (4) holds for $\overline{T}$ on V/L, so if we work by induction on the dimension then there is an ordered basis of V/L with respect to which $\overline{T}$ has an upper-triangular matrix having 0's on the diagonal. Lifting this to a linearly independent set in V and sticking a basis of L onto the beginning of the list gives an ordered basis of V with respect to which T has the type of matrix as required in (3).

Note that reversing the order of a basis does not affect properties (1), (2), and (4), but it swaps upper triangular and lower triangular matrices. Thus, we see that these properties are also equivalent to the property of being lower-triangular with 0's on the diagonal with respect to some basis.

(b) An endomorphism T is unipotent if and only if the following equivalent conditions apply:

(1) There is some $k \geq 1$ with $(T-1)^k = 0$.

(2) $(T-1)^n = 0$.

(3) There is an ordered basis of V with respect to which the matrix for T is upper-triangular with 1's on the diagonal.

(4) $\chi_T = (t-1)^n$.

The equivalence of (1), (2), and (3) follow immediately from the equivalence of the corresponding properties (1), (2), and (3) appearing in part (a), applied to the endomorphism $T-1$. We can see that (3) implies (4) immediately by looking at the matrix given by (3). Finally, the fact that (4) implies (2) follows immediately from the Cayley-Hamilton theorem.

(c) From Question 4(a), we know that there is a intrinsic and unique decomposition $V = \oplus_j V((\pi_j))$, where π_j runs through the distinct monic irreducible factors of m_T. Since k is algebraically closed, $\pi_j = (t-\lambda_j)$. We write $V(\lambda_j) := V((t-\lambda_j))$. Thus, the characteristic polynomial of T on V is the product of the characteristic polynomials of the restrictions of T to the $V(\lambda_j)$'s (think about "block matrices"). The idea is to reduce the problem to the individual generalized eigenspaces, on which everything is an easy calculation.

Recall that $V_j := V(\lambda_j)$ is the kernel of $(T-\lambda_j)^N$ for a sufficiently large integer N (that we may take to be $\dim V$, for example). Hence, any linear operator on V that commutes with T must preserve the V_j's (as it commutes with the $(T-\lambda_j)^{\dim V}$'s). Now if $T = S+N$ is an expression of a sum of commuting semisimple and nilpotent operators on V, then since S commutes with both S and N it follows that S must commute with T. Similarly, N must commute with T. Hence, such S and N necessarily preserve the V_j's. The restrictions S_j and N_j of such an S and N to each V_j are commuting operators with N_j obviously nilpotent (as $N_j^{\dim V}$ is the restriction of $N^{\dim V} = 0$) and S_j semisimple by Question 4(b)(iii). In other words, if these is to be an additive Jordan decomposition for T then it *must* be built up as a direct sum of such decompositions for the $T|_{V_j}$'s.

This shows that the problem of uniqueness for T is reduced to uniqueness for the $T|_{V_j}$'s. Conversely, if each $T|_{V_j}$ has an additive Jordan decomposition as $S_j + N_j$ then we may *define* $S = \oplus S_j$ and $N = \oplus N_j$ as visibly commuting operators on $V = \oplus V_j$ such that S is semisimple by Question 4(c), N is nilpotent (obvious!), and $S + N = \oplus T|_{V_j} = T$ as desired. In this way, we see that both the existence and uniqueness problems for T are reduced to the existence and uniqueness problems for each $T|_{V_j}$. We have noted above that $\chi_T = \prod_j \chi_{T|_{V_j}}$, and likewise $\chi_{\oplus S_j} = \prod_j \chi_{S_j}$ for such hypothetical S_j's, so if we can solve the problem on the V_j's with $T|_{V_j}$ and S_j having the same characteristic polynomial on V_j for all j then likewise T and $S = \oplus S_j$ on V have the same characteristic polynomial.

We may now choose j_0 and rename V_{j_0}, $T|_{V_{j_0}}$, and λ_{j_0} as V, T, and λ to reduce to the case when $V = V(\lambda)$ is an entire generalized eigenspace for some $\lambda \in k$. Hence, $(T-\lambda)^{\dim V} = 0$ and $\chi_T = (t-\lambda)^{\dim V}$. The expression $T = \lambda + (T-\lambda)$ is a sum of semisimple and *nilpotent* endomorphisms that clearly commute with each other (and the "semisimple part" $\lambda \cdot \mathrm{id}_V$ has the same characteristic polynomial as T). This settles the existence problem in this case, and for uniqueness suppose $T = S + N$ with commuting semisimple S and nilpotent N. Since S is semisimple and k is algebraically closed, its minimal polynomial is a product of pairwise distinct linear factors, and we may apply Lemma 3 to see that for any $\mu_j \in k$, $(T-\mu_j)$ is 0 on $V_S(\mu_j)$ (the generalized eigenspace for S with eigenvalue μ_j). Thus, $V_S(\mu_j)$ is the μ_j-eigenspace for S and $V \simeq \oplus V_S(\mu_j)$, so V is spanned by eigenvectors for S, i.e. S can be diagonalized. Now, we just have to show that S has only one eigenvalue, namely λ, for then the diagonalizable S must equal λ (consider an eigenbasis for S!) and thus $N = T - \lambda$, giving the required uniqueness. To show that any eigenvalue μ for S must equal λ, we just have to show that μ is an eigenvalue for T (since $V = V(\lambda)$ forces λ to be the only eigenvalue of T on V).

Consider the nonzero μ-eigenspace $W = \ker(S - \mu)$ in V. Since N commutes with S, it preserves W. The restriction $N|_W$ is clearly nilpotent, so any 0-eigenvector for $N|_W$ (these do exist!) is a μ-eigenvector of S and so is an eigenvector for $S + N = T$ with eigenvalue $\mu + 0 = \mu$ as well.

For a counterexample if we drop the commuting requirement, consider the additive decomposition

$$\begin{pmatrix} \lambda_1 & x \\ 0 & \lambda_2 \end{pmatrix} = \begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix} + \begin{pmatrix} 0 & x \\ 0 & 0 \end{pmatrix}.$$

Writing this as $T = S + N$, clearly S is semisimple (it is even diagonalized) and N is nilpotent (all eigenvalues are 0, and explicitly $N^2 = 0$). However,

$$SN = \begin{pmatrix} 0 & \lambda_1 x \\ 0 & 0 \end{pmatrix}, \quad NS = \begin{pmatrix} 0 & \lambda_2 x \\ 0 & 0 \end{pmatrix}.$$

Thus, if $x \neq 0$ and $\lambda_1 \neq \lambda_2$ then $SN \neq NS$, so this is *not* the additive Jordan decomposition of T. In fact, if $\lambda_1 \neq \lambda_2$ then T is diagonalizable, with eigenbasis $\{(1,0), (x/(\lambda_2 - \lambda_1), 1)\}$, so T is itself semisimple, i.e. $T_{\rm ss} = T$ and $T_{\rm n} = 0$. It follows that the Jordan decomposition of a square matrix is usually *not* easy to see in terms of an explicit matrix for the given map, even if this matrix is given in upper-triangular form.

(d) Assume T is invertible, and let $T = T_{\rm ss} + T_{\rm n}$ be its unique additive Jordan decomposition (so $T_{\rm ss}$ is invertible since T is invertible and $\chi_T = \chi_{T_{\rm ss}}$). The identity $T_{\rm n} T_{\rm ss} = T_{\rm ss} T_{\rm n}$ implies $T_{\rm ss}^{-1} T_{\rm n} = T_{\rm n} T_{\rm ss}^{-1}$, so $T_{\rm n}$ and $T_{\rm ss}^{-1}$ also commute. Thus, we can define $T'_{\rm ss} = T_{\rm ss}$ and $T'_{\rm u} = 1 + T_{\rm ss}^{-1} T_{\rm n}$, with $T_{\rm ss}^{-1} T_{\rm n}$ nilpotent because $(T_{\rm ss}^{-1} T_{\rm n})^{\dim V} = (T_{\rm ss}^{-1})^{\dim V} \cdot (T_{\rm n})^{\dim V} = 0$, by the fact that $T_{\rm ss}^{-1}$ and $T_{\rm n}$ commute. Hence, $T'_{\rm u}$ is a unipotent operator. By construction, $T'_{\rm ss}$ and $T'_{\rm u}$ clearly commute, and $T'_{\rm ss} T'_{\rm u} = T_{\rm ss} + T_{\rm n} = T$. This establishes existence of the multiplicative Jordan decomposition.

Uniqueness for multiplicative Jordan decomposition is a formal consequence of uniqueness in the additive case, as follows. Suppose SU is an expression for T as a product of commuting operators with semisimple S and unipotent U. Let $N = U - 1$, so N is nilpotent and commutes with S. Hence, $T = S(1 + N) = S + SN$ with S and SN commuting operators and SN a nilpotent operator (as its $(\dim V)$th-power is $S^{\dim V} N^{\dim V} = 0$ because $N^{\dim V} = 0$). Uniqueness in the additive case then gives that S must equal $T_{\rm ss}$ and SN must equal $T_{\rm n}$. This shows uniqueness for S (it has to equal the unique semisimple constituent from the additive Jordan decomposition), and also forces $N = S^{-1} T_{\rm n} = T_{\rm ss}^{-1} T_{\rm n}$. Hence, we are forced to have $U = 1 + N = 1 + T_{\rm ss}^{-1} T_{\rm n}$.