4/7/16

HW 3 extended to next Tues.
[New Hint : You will need the result of problem 3(b) to solve problem 4(c).]

We have completed the significant (but unfortunately named) Fundamental Theorem of Finitely Generated Modules Over a Principal Ideal Domain. Here is a summary.

☆ F.T.F.G.M.P.I.D.

Let R be a PID and let M be a finitely generated R-module. Then we have

$$M \approx R^{\oplus k} \oplus T$$

where k is the rank of M (size of a maximal R-linearly independent set) and $T = \text{Tor}_R(M)$ is the collection of R-torsion elements (which is a submodule because R is an integral domain).

Furthermore we can decompose T as a direct sum of cyclic modules

$$T \approx R/(d_1) \oplus R/(d_2) \oplus \cdots \oplus R/(d_n)$$

and if we assume that $(d_n) \subseteq \cdots \subseteq (d_1) \neq R$ then this collection of ideals is unique (called the "invariant factors" of T or M).

Alternatively if $d_n = p_1^{\alpha_{1n}} p_2^{\alpha_{2n}} \cdots p_m^{\alpha_{mn}}$ is the unique prime factorization of d_n [recall that PID $\Rightarrow$ UFD] then we can use the Chinese Remainder Theorem to decompose T further as

$$T \approx \bigoplus_{i,j} R/(p_i^{\alpha_{ij}})$$

where $d_j = p_1^{\alpha_{1j}} p_2^{\alpha_{2j}} \cdots p_m^{\alpha_{mj}}$. The ideals $(p_i^{\alpha_{ij}})$ are also uniquely determined (they are called the "elementary divisors" of T or M). ///

Remarks:

- The Chinese Remainder Theorem and the uniqueness of invariant factors / elementary divisors have not been proved yet. ↓

Hopefully some combination of you & I will prove them on HW 4. The proof of uniqueness I have in mind involves "tensoring" with the fields $R/(p_i)$.

- The original version of the FTFGMPID is due to Weierstrass (1868). If Cayley & Sylvester's "matrix" notation provided the form of Linear Algebra, then Weierstrass' Theorem provided its content. It is one of the most important theorems of 19th century mathematics.

Now we should discuss applications, I guess. Recall that there are basically two kinds of PIDs: $\mathbb{Z}$ & $K[x]$, for K a field. Applying the F.T. to $\mathbb{Z}$ gives

☆ F.T.F.G.A.G:

Recall that abelian groups are just $\mathbb{Z}$-modules. Thus if G is a finitely generated abelian group then we have a direct sum (and hence direct product) decomposition

$$G \approx \mathbb{Z}^{\oplus k} \bigoplus_{i=1}^{n} \mathbb{Z}/(d_i)$$

for some unique integers $k \geq 0$ and $1 < d_1 \mid d_2 \mid \cdots \mid d_n$.

[I stated this in MTH 761 on 12/3/15 and now it is done.]

Because of this, the structure of finite abelian groups is <u>much</u> easier than the structure of all finite groups. A classification such as this often provides shortcuts to theorems, such as the following.

☆ Primitive Root Theorem:

Let K be a field. Then any finite subgroup of $K^\times = K \setminus \{0\}$ is <u>cyclic</u>. In particular, if K is finite then $K^\times$ itself is cyclic.

↓

Proof: Let $G \leq K^{\times}$ be a finite subgroup so by the F.T. we have

$$G \approx \mathbb{Z}/(d_1) \oplus \cdots \oplus \mathbb{Z}/(d_n).$$

with $2 \leq d_1 \mid d_2 \mid \cdots \mid d_n$. Observe that for all $g \in G$ we have $g^{d_n} = 1$ (written multiplicatively). If G is <u>not cyclic</u> then we have $n \geq 2$ and hence $d_n < d_1 d_2 \cdots d_n = |G|$. It follows that

(*) $$\#\{ g \in G : g^{d_n} = 1 \} = |G| > d_n.$$

But now consider the polynomial $x^{d_n} - 1 \in K[x]$. Since this polynomial has degree d_n, you proved on HW1 [using Descartes' Theorem] that it has <u>at most d_n roots</u>, which contradicts (*).

We conclude that G is cyclic. ///

[I really like this proof. All the other proofs I know involve a nontrivial amount of number theory.]

Next we turn to the other PID, namely, $K[x]$. This was the case that Weierstrass originally had in mind.

So let K be a field and let M be a $K[x]$-module defined by a ring homomorphism

$$\lambda : K[x] \longrightarrow \text{End}_{Ab}(M).$$

Let's look at this carefully. By restricting λ to the subring $K \hookrightarrow K[x]$ we obtain a ring homomorphism

$$\lambda : K \longrightarrow \text{End}_{Ab}(M),$$

which shows us that M is, in particular, a K-vector space. But note that this is automatically a homomorphism of K-<u>algebras</u>:

$$\begin{tikzcd}
K \arrow[rr, "\lambda"] & & \text{End}_K(M). \\
& K \arrow[ul, "id"] \arrow[ur, "\lambda"'] &
\end{tikzcd}$$

Indeed, for all scalars $a, b \in K$ and for all vectors $m, n \in M$ we have

$$\begin{aligned} \lambda_a(m + \lambda_b(n)) &= \lambda_a(m) + \lambda_a(\lambda_b(n)) \\ &= \lambda_a(m) + \lambda_{ab}(n) \\ &= \lambda_a(m) + \lambda_{ba}(n) \\ &= \lambda_a(m) + \lambda_b(\lambda_a(n)), \end{aligned} \tag{$\circledast$}$$

which shows that $\lambda_a \in \mathrm{End}_K(M)$. For all $\varphi \in \mathrm{End}_K(M)$ we also have (from the <u>definition</u> of K-linearity) that

$$\lambda_a(\varphi(m)) = \varphi(\lambda_a(m))$$

and hence $\lambda_a \in Z(\mathrm{End}_K(M))$. ///.

Furthermore, note that the image of the full map $\lambda : K[x] \longrightarrow \mathrm{End}_{Ab}(M)$ is also in the subring $\mathrm{End}_K(M)$. Indeed, since the variable x is (by <u>definition</u>) in the center of $K[x]$ we have

$$\lambda_x(m + \lambda_a(n)) = \lambda_x(m) + \lambda_a(\lambda_x(n)),$$

just as in $\circledast$ above.

↓

However, I want to emphasize that the K-linear map $\lambda_x : M \to M$ need not be in the center $Z(\text{End}_K(M))$.

We have shown that if M is a $K[x]$-module then we obtain a homomorphism of K-algebras

$$\begin{tikzcd}
K[x] \arrow[rr, "\lambda"] & & \text{End}_K(M) \\
& K \arrow[ul] \arrow[ur, "\lambda"'] &
\end{tikzcd}$$

where λ_x is some (possibly non-central) K-linear map.

Conversely, if $\varphi \in \text{End}_K(M)$ is any K-linear map, then since $K[x]$ is the "free K-algebra generated by one element" there exists a unique K-algebra homomorphism $\lambda : K[x] \to \text{End}_K(M)$ such that $\lambda_x = \varphi$.

In summary, we have proved the following (tautological) result.

☆ Characterization of $K[x]$-modules :

The data of a $K[x]$-module is equivalent to a choice of (V, φ) where V is a K-vector space and $\varphi : V \to V$ is a K-linear endomorphism.

///

Remarks :

- Actually, the same argument works over any commutative ring R : the data of an $R[x]$-module is equivalent to a choice of (M, φ) where M is an R-module and $\varphi : M \to M$ is an R-linear endomorphism

- Moreover, the data of an $R[x_1, \dots, x_n]$-module is equivalent to a choice of R-module M and an n-tuple of commuting R-linear endomorphisms.

///

However, the case of $K[x]$-modules is the most interesting to us for two reasons.

1. Consider a $K[x]$-module corresponding to a pair (V, φ) where V is a finite dimensional vector space. If $V \cong K^{\oplus n}$ then after choosing a basis we can think of φ as an $n \times n$ matrix $[\varphi]$.

2. If V is finite dimensional then the corresponding $K[x]$-module will be finitely generated. Then since $K[x]$ is a PID we can apply the Fundamental Theorem to obtain some kind of "canonical decomposition" for the matrix $[\varphi]$.

Stay Tuned.

11/12/16

HW 3 due now.
HW 4 : TBA, due on Wed Apr 27.
Final Exam : Wed May 4, 11:00 – 1:30.

Right now we are applying the FTFGMPID to $K[x]$-modules when K is a field.

So let V be a $K[x]$-module with defining ring homomorphism

$$\lambda : K[x] \longrightarrow \text{End}_{Ab}(V).$$

Restricting λ to the subring $K \subseteq K[x]$ defines a K-vector space structure on the abelian group V (that's why I'm calling it V today). By examining the definitions carefully we see that λ is really a homomorphism of K-algebras

$$\lambda : K[x] \longrightarrow \text{End}_K(V)$$

an then since $K[x]$ is the "free K-algebra generated by one element",

↓

we see that λ is uniquely determined by the K-linear endomorphism $\lambda_x \in \mathrm{End}_K(V)$.

In summary,

☆ There is a bijection between $K[x]$-modules and pairs (V, α) where V is a K-vector space and $\alpha \in \mathrm{End}_K(V)$ is a K-linear endomorphism. ///

Now consider two $K[x]$-modules

$$\lambda : K[x] \to \mathrm{End}_K(U) \quad \& \quad \mu : K[x] \to \mathrm{End}_K(V)$$

such that $\lambda_x = \alpha$ & $\mu_x = \beta$. Recall that a $K[x]$-module homomorphism is defined to be a homomorphism of abelian groups $\varphi : U \to V$ such that for all $f(x) \in K[x]$ we have a commutative square

(*)

$$\begin{tikzcd}
U \arrow[r, "\varphi"] \arrow[d, "\lambda_{f(x)}"'] & V \arrow[d, "\mu_{f(x)}"] \\
U \arrow[r, "\varphi"'] & V
\end{tikzcd}$$

.

By taking $f(x) \in K \subseteq K[x]$, the diagram $\circledast$ tells us that $\varphi : U \to V$ is a K-linear map. Then substituting $f(x) = x$ into $\circledast$ gives

$$\begin{tikzcd} U \arrow[r, "\varphi"] \arrow[d, "\alpha"'] & V \arrow[d, "\beta"] \\ U \arrow[r, "\varphi"'] & V \end{tikzcd} \quad \text{, i.e., } \varphi \circ \alpha = \beta \circ \varphi .$$

If, moreover, φ is an isomorphism of $K[x]$-modules we conclude that $\varphi : U \to V$ is an isomorphism of K-vector spaces such that

$$\varphi \circ \alpha \circ \varphi^{-1} = \beta .$$

In summary,

☆ Two $K[x]$-modules defined by (U, α) and (V, β) are isomorphic if and only if there exists an isomorphism $\varphi : U \to V$ of K-vector spaces such that

$$\varphi \circ \alpha \circ \varphi^{-1} = \beta .$$

///

This leads to a surprising idea. Let V be a K-vector space and recall that the general linear group $GL_K(V) := Aut_K(V)$ acts on endomorphisms by conjugation

$$GL_K(V) \circlearrowright End_K(V)$$

$$\varphi \cdot \alpha := \varphi \circ \alpha \circ \varphi^{-1}$$

If V is finitely generated of rank n then by choosing a basis this gives us an action of the group of <u>invertible</u> $n \times n$ matrices on the algebra of <u>all</u> $n \times n$ matrices

$$GL_n(K) \circlearrowright Mat_n(K).$$

$$P \cdot A := PAP^{-1}.$$

[Remark: If $K \in \{\mathbb{R}, \mathbb{C}\}$ then we would call this the "adjoint action" of the Lie group $GL_n(K)$ on its Lie algebra $Mat_n(K)$.]

In general we say that two matrices in the same orbit of this action are "similar".

Our goal is now to classify matrices in $\mathrm{Mat}_n(K)$ up to "similarity". This problem is analogous to the problems solved by the RREF, Hermite & Smith Normal Forms, so we are looking for some new kind of "normal form".

I consider it surprising that the answer to this problem follows from the Smith Normal Form (via the FTFGMPID).

To be specific, consider any matrix $A \in \mathrm{Mat}_n(K)$. This defines a $K[x]$-module structure on $V := K^{\oplus n}$ via the map

$$\begin{aligned} \lambda : K[x] &\longrightarrow \mathrm{End}_K(V) \\ x &\longmapsto A. \end{aligned}$$

Since V is finite dimensional this module is finitely generated and since $K[x]$ is a PID we can apply the FTFGMPID to obtain an isomorphism of $K[x]$-modules

$$(*) \qquad \varphi : V \xrightarrow{\sim} \bigoplus_{i=1}^{d} \frac{K[x]}{(f_i(x))}$$

for some polynomials $f_i(x) \in K[x]$. By thowing away zero summands in $\circledast$ we can assume that none of these polynomials is a unit. Furthermore, I claim that none of the polynomials is zero. Indeed, the direct sum of $K[x]$-modules in $\circledast$ restricts to a direct sum of K-vector spaces, and we know that dimension adds over direct sums. Since we have assumed that V is finite dimensional this implies that each direct summand on the right is finite dimensional. Finally, we observe that

$$K[x] = K[x]/(0)$$

is <u>not</u> finite dimensional (e.g. because the set $1, x, x^2, \ldots$ is K-linearly independent).

In summary we can assume that the polynomials are non<u>constant</u>, <u>monic</u>, and satisfy

$$f_1(x) \mid f_2(x) \mid \cdots \mid f_d(x)$$

and in this case they are <u>unique</u>.

[Jargon : The polynomial $f_d(x)$ is called the minimal polynomial of A and the product

$$f_1(x) f_2(x) \cdots f_d(x)$$

is called the characteristic polynomial.]

///

Now we can think of the isomorphism φ as a "change of basis" on the vector space V, in which case it is represented by some invertible matrix $P \in GL_n(K)$. Furthermore since each of the summands in $\circledast$ is a $K[x]$-<u>sub</u>module we know in particular that it is preserved by the action of "x", i.e., by the action of the matrix A. Thus if we choose a K-basis for V by concatenating bases for the summands we obtain a block-diagonal matrix

$$PAP^{-1} = \begin{pmatrix} \boxed{B_1} & & & 0 \\ & \boxed{B_2} & & \\ & & \ddots & \\ 0 & & & \boxed{B_d} \end{pmatrix}$$

It remains only to choose some "canonical" bases for the summands. I think you'll agree that the following choice is reasonable.

☆ The Companion Matrix:

Let $f(x) \in K[x]$ be a monic polynomial

$$f(x) = r_0 + r_1 x + \cdots + r_{m-1} x^{m-1} + x^m.$$

And consider the basis $1, x, x^2, \ldots, x^{m-1}$ for the K-vector space $K[x]/(f(x))$. In terms of this basis the action of "x" satisfies

$$\begin{aligned} x \cdot 1 &= x \\ x \cdot x &= x^2 \\ &\vdots \\ x \cdot x^{m-2} &= x^{m-1} \\ x \cdot x^{m-1} &= x^m = -r_0 - r_1 x - \cdots - r_{m-1} x^{m-1} \end{aligned}$$

Thus if we think of "multiplying by x" as a K-linear endomorphism

$$K[x]/(f(x)) \longrightarrow K[x]/(f(x)),$$

its matrix in this basis is the following so-called "companion matrix"

$$C_{f(x)} := \begin{pmatrix} 0 & & & & -r_0 \\ 1 & 0 & & & -r_2 \\ & 1 & 0 & & \vdots \\ & & \ddots & 0 & -r_{m-2} \\ & & & 1 & -r_{m-1} \end{pmatrix}$$

In summary, we have the following.

☆ Theorem on Rational Canonical Form (RCF):

Let K be a field and consider any $n \times n$ matrix $A \in \mathrm{Mat}_n(K)$. Then there exist unique monic polynomials

$$f_1(x) \mid f_2(x) \mid \cdots \mid f_d(x)$$

and a unique invertible matrix $P \in GL_n(K)$ such that

$$PAP^{-1} = \left(\begin{array}{c|c|c} C_{f_1(x)} & & \bigcirc \\ \hline & C_{f_2(x)} & \\ \hline \bigcirc & & C_{f_d(x)} \end{array}\right)$$

We call this block diagonal matrix the rational canonical form of A : $RCF(A)$.

Furthermore, the RCF is a complete invariant for the action of $GL_n(K)$ on $Mat_n(K)$ by conjugation. ///

The RCF essentially corresponds to the "invariant factor" decomposition of the $K[x]$-module defined by A.

If K is algebraically closed (or if we're willing to pass to an algebraic extension) then we also have a canonical form corresponding to the "elementary divisor" decomposition.

Let $f_1(x) \mid f_2(x) \mid \cdots \mid f_d(x)$ be the invariant factors of A and suppose that the minimal polynomial splits as

$$f_d(x) = \prod_i (x - a_i)^{\alpha_{id}}$$

with $a_i \in K$.

Then we have an isomorphism $K[x]$-modules

$$V \approx \bigoplus_{i,j} \frac{K[x]}{((x-a_i)^{q_{ij}})}$$

where the action of A on the left corresponds to "multiplication by x" on the right. Our goal is to choose a canonical basis for each summand on the right. I think you'll agree that the following choice is reasonable.

☆ Jordon Blocks :

Consider the basis $1, (x-a), \ldots, (x-a)^{m-1}$ for the K-vector space $K[x]/((x-a)^m)$. In terms of this basis, "multiplication by x" is given by

$$x \cdot 1 = x = (x-a) + a \cdot 1,$$

$$\begin{aligned} x \cdot (x-a) &= x^2 - ax \\ &= (x-a)^2 + a \cdot (x-a), \end{aligned}$$

$$\vdots$$

$$x \cdot (x-a)^{m-2} = (x-a)^{m-1} + a \cdot (x-a)^{m-2},$$

and finally

$$x \cdot (x-a)^{m-1} = \overset{0}{\cancel{(x-a)^{m}}} + a \cdot (x-a)^{m-1}$$
$$= (x-a)^{m-1},$$

Thus it is represented by the $m \times m$ matrix

$$J_a(m) := \begin{pmatrix} a & & & & \\ 1 & a & & \bigcirc & \\ & 1 & a & & \\ & & \ddots & \ddots & \\ & \bigcirc & & & a \\ & & & 1 & a \end{pmatrix},$$

which we call a "Jordon block".

In summary, we have the following.

☆ Theorem on Jordan Canonical Form (JCF):

Let K be an algebraically closed field and consider any $n \times n$ matrix $A \in \mathrm{Mat}_n(K)$. Then there exist elements $a_i \in K$ (the eigenvalues of A), positive integers $a_{ij} \in \mathbb{N}$, and an invertible matrix $P \in GL_n(K)$ such that

$$PAP^{-1} = \begin{pmatrix} \ddots & & \\ & \boxed{J_{n_i}(\alpha_{ij})} & \\ & & \ddots \end{pmatrix}$$

If we order the diagonal Jordan blocks in some reasonable way then the matrix $JCF(A) := PAP^{-1}$ is unique. We call it <u>the</u> Jordan Canonical Form of A.

Furthermore, the JCF is a complete invariant for the action of $GL_n(K)$ on $Mat_n(K)$ by conjugation. ///

=

We'll discuss consequences of the RCF & JCF next time.

4/14/16

HW4 due Wed Apr 27.
Final Exam on Wed May 4, 11:00-1:30.

Last time we derived the Rational & Jordan Canonical forms of a matrix from the FTFGM over the PID $K[x]$.

Recall: The ring $\mathrm{Mat}_n(K)$ is actually a K-algebra via the ring homomorphism

$$\lambda : K \longrightarrow \mathrm{Mat}_n(K)$$

defined by $\lambda(a) := aI_n \in Z(\mathrm{Mat}_n(K))$. Then for any matrix $A \in \mathrm{Mat}_n(K)$ there exists a canonical "evaluation map" from the free algebra

$$\varphi_A : K[x] \longrightarrow \mathrm{Mat}_n(K)$$

defined by $\varphi_A(x) := A$ and $\varphi_A(a) = \lambda(a) = aI_n$ for all $a \in K \subseteq K[x]$. For convenience will often write

$$\varphi_A(f(x)) = \text{“} f(A) \text{”}.$$

To be explicit, if $f(x) = a_0 + a_1 x + \cdots + a_m x^m$ then we have

$$f(A) = a_0 I_n + a_1 A + \cdots + a_m A^m.$$

Now since φ_A is, in particular, a ring homomorphism we know that its kernel

$$\ker \varphi_A = \{ f(x) \in K[x] : f(A) = 0 \}$$

is an ideal of $K[x]$. Then since $K[x]$ is a PID (recall that $K[x]$ is a Euclidean domain via long division of polynomials and that Euclidean $\Rightarrow$ PID) we conclude that

$$\ker \varphi_A = (m_A(x))$$

for some unique monic polynomial $m_A(x)$ in $K[x]$ called the <u>minimal polynomial</u> of A.

Compare this to the minimal polynomial $m_a(x) \in K[x]$ of an element $a \in L \supseteq K$ in a field extension. If $m_a(x) \neq 0$ (i.e., if a is algebraic over K)

$\downarrow$

Then you proved on HW3 that $m_a(x)$ is irreducible. Proof: Suppose that we have $m_a(x) = f(x)g(x)$ for some $f(x), g(x) \in K[x]$ and evaluate at a to get

$$0 = m_a(a) = f(a)g(a) \tag{*}$$

Since L is a domain we can assume without loss that $f(a) = 0$. This means that $f(x) \in \ker \varphi_a = (m_a(x))$ and hence we have $f(x) = m_a(x)q(x)$ for some $q(x) \in K[x]$. But then since $m_a(x) \neq 0$ and since $K[x]$ is a domain we have

$$\begin{aligned}
m_a(x) &= f(x)g(x)\\
m_a(x) &= m_a(x)q(x)g(x)\\
m_a(x)\big(1 - q(x)g(x)\big) &= 0\\
1 - q(x)g(x) &= 0\\
1 &= q(x)g(x),
\end{aligned}$$

which implies that $g(x)$ is a unit. ///

However, the minimal polynomial of a matrix may be <u>reducible</u>. In this case the above proof fails at step (*)

because $\mathrm{Mat}_n(K)$ is not an integral domain. Indeed, even the commutative subring

$$K[A] := \mathrm{im}\,\varphi_A \subseteq \mathrm{Mat}_n(K)$$

might have zero divisors. In a sense, HW4 is all about exploring this possibility.

Now since $K[x]$ is a unique factorization domain [recall that PID $\Rightarrow$ UFD] we have a unique factorization

$$m_a(x) = p_1(x)^{m_1} \cdots p_d(x)^{m_d}$$

where the polynomials $p_i(x) \in K[x]$ are monic & irreducible.

There is another important polynomial associated to A called the characteristic polynomial $\chi_A(x) \in K[x]$, which is defined via the determinant. I was hoping to introduce determinants the correct way (via exterior powers), but we ran out of time, so here's the quick and dirty way.

☆ Let R be a commutative ring and consider an $n \times n$ matrix $A = (a_{ij}) \in \mathrm{Mat}_n(R)$. We define the *determinant* of A as an alternating sum of products of entries, indexed by permutations $\pi \in S_n$:

$$\det(A) := \sum_{\pi \in S_n} (-1)^{\pi} \prod_{i=1}^{n} a_{i,\pi(i)} \in R.$$

[Determinants were first mentioned in the year 1683, independently by Seki in Japan and Leibniz in Europe, and they are the reason that matrices were invented. Indeed, Sylvester's word "matrix" means a "womb" that gives birth to determinants!]

One can easily check that the determinants of elementary matrices are

$$\begin{aligned} \det(E_{ij}(r)) &= 1, \quad i \neq j \\ \det(E_{ii}(r)) &= r \\ \det(P_{ij}) &= -1, \quad i \neq j. \end{aligned}$$

Then for any matrix $A \in \mathrm{Mat}_n(R)$ and for any elementary matrix $E \in \mathrm{Mat}_n(R)$ one can check that

$$\det(EA) = \det(AE) = \det(E)\det(A).$$

When $R = K$ is a field we showed on 3/29/16 that every matrix $A \in \mathrm{Mat}_n(K)$ can be reduced to the form

$$\left(\begin{array}{c|c} I_k & 0 \\ \hline 0 & 0 \end{array}\right)$$

by multiplying on the right and left by elementary matrices. Thus we obtain the following result.

☆ Theorem on Determinants:

Let K be a field and consider matrices $A, B \in \mathrm{Mat}_n(K)$. Then we have

- $A \in GL_n(K) \iff \det(A) \neq 0$.
- $\det(AB) = \det(A)\det(B)$.

Proof : There exist $E, F \in GL_n(K)$ such that E & F are products of (invertible) elementary matrices and such that

$$EAF = \left(\begin{array}{c|c} I_k & 0 \\ \hline 0 & 0 \end{array}\right)$$

$$A = E^{-1}\left(\begin{array}{c|c} I_k & 0 \\ \hline 0 & 0 \end{array}\right)F^{-1}$$

It follows that A is invertible if and only if $k = n$. Then from the above observations we have

$$\det(A) = \det(E^{-1})\det\left(\begin{array}{c|c} I_k & 0 \\ \hline 0 & 0 \end{array}\right)\det(F^{-1})$$

and it follows that $\det(A) \neq 0 \iff k = n$.

Now if at least one of A & B is not invertible then AB is not invertible and we have

$$\det(AB) = 0 = \det(A)\det(B).$$

So assume that $A, B \in GL_n(K)$. Then we have $E_1, E_2, F_1, F_2 \in GL_n(K)$ such that E_1, E_2, F_1, F_2 are products of elementary matrices and such that

$$E_1 A F_1 = I_n \quad \& \quad E_2 B F_2 = I_n.$$

Then we have $A = E_1^{-1} F_1^{-1}$ & $B = E_2^{-1} F_2^{-1}$, and it follows that

$$\begin{aligned} \det(AB) &= \det(E_1^{-1} F_1^{-1} E_2^{-1} F_2^{-1}) \\ &= \det(E_1^{-1} F_1^{-1}) \det(E_2^{-1} F_2^{-1}) \\ &= \det(A) \det(B). \end{aligned}$$

///

Remarks :

- As long as $GL_n(R)$ is generated by elementary matrices [for example if R is a Euclidean domain] then the same proof works to show that

- $\det(AB) = \det(A)\det(B)$
- $A \in GL_n(R) \iff \det(A) \in R^\times$.

- In fact, the same result holds when R is an arbitrary commutative ring but it requires a new proof. Here's my favorite proof (details omitted):

Let $\Lambda^k R^{\oplus n}$ be the "k-th exterior power" of the free R-module $R^{\oplus n}$. Then $\Lambda^k R^{\oplus n}$ is a free R-module of rank $\binom{n}{k} = n!/(k!(n-k)!)$. In particular we have

$$\Lambda^n R^{\oplus n} \approx R.$$

Now any R-linear map $\varphi : R^{\oplus n} \to R^{\oplus n}$ induces a map $\Lambda^n \varphi : R \to R$ and if we choose a basis then this induced map is "multiplication by the determinant":

$\circledast$

$$\begin{tikzcd}
R^{\oplus n} \arrow[r, "A"] \arrow[d] & R^{\oplus n} \arrow[d] \\
\Lambda^n R^{\oplus n} \arrow[r, "\det(A)\cdot"] & \Lambda^n R^{\oplus n}
\end{tikzcd}$$

From this point of view the theorem is obvious, but then one must prove that the combinatorial definition of the determinant is equivalent to $\circledast$. ///

Now we can define the "other" polynomial.

☆ Let K be a field and consider a matrix $A \in \mathrm{Mat}_n(K)$. We define the characteristic polynomial as a determinant

$$\chi_A(x) := \det(x I_n - A) \in K[x].$$

Note that this is a monic polynomial of degree n. Its main pupose is for detecting the eigenvalues of A. Recall that $\lambda \in K$ is called an eigenvalue of A [from the German "eigen" = "belonging to"] if there exists some nonzero vector $u \in K^{\oplus n}$ such that

$$A u = \lambda u.$$

In other words, λ is on eigenvalue

$$\iff \exists\, u \neq 0 \text{ with } Au = \lambda I_n u$$

$$\iff \exists\, u \neq 0 \text{ with } (\lambda I_n - A) u = 0$$

$$\iff \ker(\lambda I_n - A) \neq 0$$

↓

$\iff \lambda I_n - A$ is not injective

$\circledast$ $\iff \lambda I_n - A$ is not invertible

$\iff \det(\lambda I_n - A) = 0$

$\iff \chi_A(\lambda) = 0.$

[Note that the $\Leftarrow$ direction of $\circledast$ uses the fact that $\dim(V/W) = \dim V - \dim W$ for vector spaces, which is sometimes called the "rank-nullity theorem".].

Then as a consequence of the Theorem on Determinants we have for all invertible matrices $P \in GL_n(K)$ that

$$\begin{aligned}
\chi_{PAP^{-1}}(x) &= \det(xI_n - PAP^{-1}) \\
&= \det(xPI_nP^{-1} - PAP^{-1}) \\
&= \det(P(xI_n - A)P^{-1}) \\
&= \det(P)\det(xI_n - A)\det(P^{-1}) \\
&= \cancel{\det(P)}\det(xI_n - A)\cancel{\det(P)}^{-1} \\
&= \det(xI_n - A) \\
&= \chi_A(x).
\end{aligned}$$

From this we conclude that:

- The matrix $A \in \mathrm{Mat}_n(K)$ has at most n distinct eigenvalues.
- The eigenvalues are invariant under change of basis, so they are really a property of K-linear endomorphisms.

Now we have two polynomials associated to the matrix A,

- $m_A(x)$ related to linear relations among the powers of A
- $\chi_A(x)$ related to eigenvalues of A,

and might wonder how these polynomials are related to each other. Here is the answer.

☆ Cayley-Hamilton Theorem:

$$m_A(x) \mid \chi_A(x).$$

There are many ways to prove this; we'll prove it by applying the Rational Canonical Form. But first let's discuss a few consequences.

- Since $\chi_A(x) \neq 0$ the theorem tells us that $m_A(x) \neq 0$. In other words, there is no such thing as a "transcendental matrix"; all matrices are "algebraic".

- The roots of $m_A(x)$ (if any exist) are eigenvalues of A. Indeed we have $\chi_A(x) = m_A(x) f(x)$ for some $f(x) \in K[x]$. If α is a root of $m_A(x)$ then

$$\chi_A(\alpha) = m_A(\alpha) f(\alpha) = 0 \cdot f(\alpha) = 0,$$

hence α is an eigenvalue of A.

- Evaluating $\chi_A(x)$ at A gives

$$\chi_A(A) = m_A(A) \cdot f(A) = 0 \cdot f(A) = 0,$$

thus A is annihilated by its own characteristic polynomial. [In fact, this property is equivalent to the statement of the theorem.]

Now the proof.

☆ Proof of Cayley-Hamilton:

This will follow directly from the RCF after a few observations.

1. Let V be the $K[x]$-module defined by

$$\begin{aligned}\lambda : K[x] &\longrightarrow \mathrm{End}_K(V)\\ x &\longmapsto A\end{aligned}$$

and let $f_1(x) \mid f_2(x) \mid \cdots \mid f_d(x)$ be the associated monic invariant factors. We showed previously that $(f_d(x))$ equals the annilator ideal $\mathrm{Ann}_{K[x]}(V) \subseteq K[x]$. But note that we also have

$$\begin{aligned}\mathrm{Ann}_{K[x]}(V) &= \{ f(x) \in K[x] : \lambda_{f(x)}(u) = 0 \ \forall u \in V \}\\ &= \{ f(x) \in K[x] : f(A)u = 0 \ \forall u \in V \}\\ &= \{ f(x) \in K[x] : f(A) = 0 \}\\ &= \ker \varphi_A = (m_A(x)).\end{aligned}$$

Since $(f_d(x)) = (m_A(x))$ then since both polynomials are monic we conclude that $m_A(x) = f_d(x)$.

2. Consider any monic polynomial

$$f(x) = r_0 + r_1 x + \cdots + r_{m-1} x^{m-1} + x^m \in K[x]$$

with corresponding $m \times m$ companion matrix

$$C_{f(x)} = \begin{pmatrix} 0 & & & & & -r_0 \\ 1 & 0 & & & & -r_1 \\ & 1 & 0 & & & \vdots \\ & & & \ddots & & \vdots \\ & & & & 0 & -r_{m-2} \\ & & & & 1 & -r_{m-1} \end{pmatrix}.$$

Then one can check by induction that the characteristic polynomial is

$$\chi_{C_{f(x)}}(t) = \det(t I_m - C_{f(x)}) = f(t).$$

3. By the RCF there exists an invertible matrix $P \in GL_n(K)$ such that

$$PAP^{-1} = \begin{pmatrix} C_{f_1(x)} & & & O \\ & C_{f_2(x)} & & \\ & & \ddots & \\ O & & & C_{f_l(x)} \end{pmatrix}$$

Then since the determinant of a block diagonal matrix is equal to the product of the determinants of the blocks we conclude that

$$\begin{aligned}\chi_A(t) &= \chi_{PAP^{-1}}(t) \\ &= \prod_{i=1}^{d} \chi_{C_{f_i(x)}}(t) \\ &= \prod_{i=1}^{d} f_i(t) = m_A(t) \prod_{i=1}^{d-1} f_i(t).\end{aligned}$$

QED.

It is also true that any eigenvalue of A is a root of the minimal polynomial $m_A(x)$.

Proof: Let $\lambda \in K$ be an eigenvalue with eigenvector $u \in K^{\oplus n}$. Then we have $A^k u = \lambda^k u$ for all $k \geq 0$. It follows that for any polynomial $f(x) \in K[x]$ we have $f(A)u = f(\lambda)u$. In particular, taking $f(x) = m_A(x)$ gives

$$0 = 0u = m_A(A)u = m_A(\lambda)u$$

Since $u \neq 0$ this implies that $m_A(\lambda) = 0$. ///

You will explore the relationship between $m_A(x)$ & $\chi_A(x)$ further on HW4.

Here is the most general thing that can be said about the topic. I'll state it without proof.

☆ Minimal vs. Characteristic Polynomial:

Let $A \in \mathrm{Mat}_n(K)$ over a general field K and suppose the minimal polynomial has the unique prime factorization

$$m_A(x) = p_1(x)^{m_1} p_2(x)^{m_2} \cdots p_d(x)^{m_d}.$$

where each $p_i(x)$ is irreducible of degree r_i. If we define the dimensions

$$n_i := \dim_K \ker\left(p_i(A)^{m_i}\right)$$

Then $r_i \mid n_i$ for all i, and we have

$$\chi_A(x) = p_1(x)^{n_1/r_1} p_2(x)^{n_2/r_2} \cdots p_d(x)^{n_d/r_d}.$$

4/19/16

HW4 due next Wed Apr 27
[see the new version online (yes, my computer got fixed)]
The Final Exam is Wed May 4 at 11:00 – 1:30 in Memorial 217 (maybe I can get the room changed).

This Week: Epilogue

I was hoping to talk more about non-commutative rings & algebras this semester. In particular, I wanted to discuss the "Artin-Wedderburn Theorem" and then use it to derive the Fundamental Theorem of Galois Theory (via a nice proof by Mitya Boyarchenko). Alas, it was not to be.

Algebra is just too big.

Instead I will spend the final two lectures on the concept of "base change". This will allow us to give the correct proof of uniqueness in the FTFGMPID.

Let R be a ring (either commutative or non-commutative). Our philosophy in MTH 762 was to replace the study of R by the study of the abelian category R-Mod. There is a deep theorem underlying this philosophy.

☆ Definition: We say that rings R & S are Morita equivalent if there exists an equivalence of categories $R\text{-Mod} \approx S\text{-Mod}$, i.e., a pair of functors

$$(\circledast) \qquad F : R\text{-Mod} \rightleftarrows S\text{-Mod} : G$$

with natural isomorphisms

$$G \circ F \cong id_{R\text{-Mod}} \quad \& \quad F \circ G \cong id_{S\text{-Mod}}.$$

In this case the functors F, G will automatically be additive so that $(\circledast)$ is an equivalence of abelian categories. The equivalence $(\circledast)$ also implies the existence of an equivalence of categories of right R & S modules $\text{Mod-}R \approx \text{Mod-}S$, so there is no need to distinguish between "left" and "right" Morita equivalence. ///.

Now here's the theorem.

☆ Theorem on Morita Equivalence:

Let R & S be <u>commutative</u> rings. Then

$$\underset{\substack{\uparrow\\ \text{isomorphism}\\ \text{of rings}}}{R \approx S} \quad \Longleftrightarrow \quad \underset{\substack{\uparrow\\ \text{equivalence of}\\ \text{categories.}}}{R\text{-Mod} \approx S\text{-Mod}}$$

In other words, the commutative ring R is determined up to isomorphism by its category of modules. ///

Remarks:

- This is another "Cayley"-type theorem telling us that the definition of the category R-Mod is not "too general".

- The theorem is false for non-commutative rings [in fact for <u>any</u> ring R and for any $n \geq 1$ we have an equivalence

$$R\text{-Mod} \approx \text{Mat}_n(R)\text{-Mod} \;].$$

↓

Thus to determine R from the category R-Mod would require more than just the abelian category structure. The subject of "noncommutative geometry" tries to discover this extra structure.

The theorem on Morita equivalence suggests that homomorphisms of rings should somehow correspond to functors between their categories. To be specific:

Consider any ring homomorphism

$$\varphi : R \longrightarrow S$$

This determines a fairly obvious functor

$$\varphi_* : S\text{-Mod} \to R\text{-Mod}$$

called "restriction of scalars". To define it consider any S-module M defined by a ring homomorphism

$$\lambda : S \longrightarrow \mathrm{End}_{Ab}(M)$$

Then pre-composing with φ gives a new ring homomorphism

$$\lambda \circ \varphi : R \longrightarrow \mathrm{End}_{Ab}(M)$$

which defines an R-module structure on the abelian group M. In the special case that $i : R \hookrightarrow S$ is an inclusion, the functor i_* is literally "restriction of scalars".

OK, but that was too easy. We really want a way to "extend scalars" from R to S. Since restriction trivially preserves limits and colimits (we say that φ_* is an "exact functor") we suspect that there might be two functors $R\text{-Mod} \to S\text{-Mod}$, one right adjoint and one left-adjoint to φ_*.

One of these is easier to construct, so we'll do it first. Let N be an R-module defined by a ring homomorphism

$$\mu : R \longrightarrow \mathrm{End}_{Ab}(N)$$

and recall that the ring homomorphism $\varphi : R \to S$ also defines an R-module structure on the abelian group $|S|$ by

$$R \longrightarrow \mathrm{End}_{Ab}(|S|)$$
$$r \longmapsto (s \longmapsto \varphi(r)s).$$

Now consider the abelian group $\mathrm{Hom}_R(S,N)$ of R-linear maps between these two R-modules. We can define a natural S-module structure on $\mathrm{Hom}_R(S,N)$ as follows:

Given $\alpha : S \to N$ & $s \in S$ we define the map $s\alpha : S \to N$ by setting

$$(s\alpha)(s') := \alpha(ss') \quad \forall\ s' \in S.$$

Aluffi calls the resulting S-module

$$\varphi^!(N) := \mathrm{Hom}_R(S,N).$$

[and I guess this is pronounced as "shriek"].

By chasing definitions (!) one can show that this defines a functor $\varphi^! : R\text{-Mod} \to S\text{-Mod}$ and that for all $N \in R\text{-Mod}$ & $M \in S\text{-Mod}$ we have a natural isomorphism

$$\mathrm{Hom}_R(\varphi_*(M), N) \cong \mathrm{Hom}_S(M, \varphi^!(N)).$$

In particular, this tells us that $\varphi^!$ is a right-adjoint functor, so it preserves limits.

That was the easy one. Next, suppose we have some hypothetical functor φ^* from R-Mod to S-Mod satisfying the natural isomorphism

$$\mathrm{Hom}_R(N, \varphi_*(M)) \cong \mathrm{Hom}_S(\varphi^*(N), M).$$

What would this mean?

First note that we have an isomorphism of S-modules

$$\begin{aligned}\mathrm{Hom}_S(S, M) &\approx M \\ \alpha &\longmapsto \alpha(1)\end{aligned}$$

If we think of this as an R-module by restriction then the above isomorphism becomes

$$(*) \quad \mathrm{Hom}_R(N, \mathrm{Hom}_S(S, M)) \cong \mathrm{Hom}_S(\varphi^*(N), M).$$

Where have we seen this before? It was hinted at on HW3 Problem 2. The isomorphism $\circledast$ tells us two things.

1. Universal Property of $\varphi^*(N)$:

From HW3.2(a) we have

$$\begin{array}{c} \mathrm{Hom}_R(N, \mathrm{Hom}_S(S,M)) \subseteq \mathrm{Hom}_{set}(N, \mathrm{Hom}_{set}(S,M)) \\ \hspace{6cm}\updownarrow \\ \hspace{6cm}\mathrm{Hom}_{set}(N\times S, M). \end{array}$$

Thus we can think of each element of the left group as a function $N\times S \to M$. In fact, the functions on the left have the property of being "R-bilinear". Then $\circledast$ tells us that for each R-bilinear function $f: N\times S \to M$ there should be a unique S-$\underline{\text{uni}}$linear function $\bar{f}: \varphi^*(N) \to M$.

2. How to construct it:

From $\circledast$ and from HW3.2(b) we have

$$\Big\downarrow$$

$$\begin{aligned} \mathrm{Hom}_S(\varphi^*(N), M) &\cong \mathrm{Hom}_R(N, \mathrm{Hom}_S(S, M)) \\ &\subseteq \mathrm{Hom}_{set}(N, \mathrm{Hom}_S(S, M)) \\ &\cong \mathrm{Hom}_S(S^{\oplus N}, M). \end{aligned}$$

which suggests that maybe we can construct $\varphi^*(N)$ as a quotient of the (infinitely generated) free S-module $S^{\oplus N}$.

Indeed, this is how the module $\varphi^*(N)$ is usually constructed: By starting with the obvious S-module

$$S^{\oplus N} = \left\{ \sum_i s_i n_i \ : \ n_i \in N \right\}$$

and then dividing by all the relations necessary to achieve the desired universal property.

I'll spare you the details and just tell you the name of this thing:

$$\varphi^*(N) = N \otimes_R S .$$

We call this the "tensor product" of N with S over the base ring R.

Basically, a "tensor product" is a machine for turning bilinear functions into unilinear functions. I'm sad that I didn't have more time to motivate this. For now I'll just mention a few properties.

- From the natural isomorphism

$$\mathrm{Hom}_R(N, \mathrm{Hom}_S(S,M)) \cong \mathrm{Hom}_S(N \otimes_R S, M)$$

we see that $(-) \otimes_R S : R\text{-Mod} \to S\text{-Mod}$ is a left-adjoint functor, so it commutes with colimits.

- Plugging $N = R$ into the isomorphism gives

$$\begin{aligned} \mathrm{Hom}_R(R, \mathrm{Hom}_S(S,M)) &\cong \mathrm{Hom}_S(R \otimes_R S, M) \\ \mathrm{Hom}_S(S,M) &\cong \mathrm{Hom}_S(R \otimes_R S, M) \end{aligned}$$

and then from Yoneda we conclude that

$$S \approx R \otimes_R S$$

as S-modules.

- Finally, since the construction of the free module $R^{\oplus A}$ is a colimit we conclude that for any set A we have

$$R^{\oplus A} \otimes_R S \approx (R \otimes_R S)^{\oplus A} \approx S^{\oplus A}.$$

This is a very useful property. We will use it next time to complete the proof of uniqueness in the FTFGMPID.

4/21/16

HW4 due next Wed Apr 27.
Final Exam on Wed May 4.
I plan to hold a review session before
that so let me know your schedules.

Last Day of Class.

Last time we discussed "restriction and extension of scalars" for commutative rings. The same discussion can be generalized to non-commutative rings but we didn't have time to go into the details.

Recall: Let R & S be commutative rings and let $\varphi: R \to S$ be any ring homomorphism. Then we have an adjoint pair of functors between the abelian categories of modules:

$$\begin{tikzcd} R\text{-Mod} \arrow[r, shift left, "\varphi^*"] & S\text{-Mod} \arrow[l, shift left, "\varphi_*"] \end{tikzcd}$$

The functor φ_* is called "restriction of scalars" and it is easy to define:

Given any S-Module N defined by a ring hom

$$\mu : S \to \mathrm{End}_{Ab}(N).$$

we let $\varphi_*(N)$ be the R-module structure on the abelian group N defined by the ring hom

$$\mu \circ \varphi : R \to \mathrm{End}_{Ab}(N). \quad ///$$

The functor $\varphi^* : R\text{-Mod} \to S\text{-Mod}$ is called "extension of scalars". It is quite a bit trickier to construct, but if you are willing to just believe in its <u>existence</u> then it is uinquely characterized by being left-adjoint to φ_*. That is, for all modules $M \in R\text{-Mod}$ & $N \in S\text{-Mod}$ we have a "natural" isomorphism of abelian groups

$$\mathrm{Hom}_R(M, \varphi_*(N)) \cong \mathrm{Hom}_S(\varphi^*(M), N).$$

[Note that this is yet another example of a "free-forgetful adjunction", where the forgetful functor is obvious and the free functor is not.]

The word "natural" means that the isomorphism of abelian groups is induced from a "natural isomorphism" of bifunctors

$$\mathrm{Hom}_R(-, \varphi_*(-)) \cong \mathrm{Hom}_S(\varphi^*(-), -).$$

You can imagine this might be hard to check; luckily we don't have to check it because we're using it as the definition of φ^*.

///

More often you will see restriction and extension of scalars expressed in the jargon of the "hom-tensor adjunction": Given a module $N \in S\text{-Mod}$ recall that we have a natural isomorphism $N \cong \mathrm{Hom}_S(S, N)$. Then we will abuse terminology to write

$$\varphi_*(N) = N = \mathrm{Hom}_S(S, N).$$

For any R-Module M we will define the notation

$$M \otimes_R S := \varphi^*(M)$$

[Yes, the symbol $\otimes$ has some extra meaning but for now we can just take this as a funny notation for extension of scalars.]

Then our adjunction becomes

(*) $$\mathrm{Hom}_R(M, \mathrm{Hom}_S(S, N)) \cong \mathrm{Hom}_S(M \otimes_R S, N).$$

The reason we might want to write it this way is because (*) is a special case of the more general hom-tensor adjunction

$$\mathrm{Hom}_R(M, \mathrm{Hom}_S(P, N)) \cong \mathrm{Hom}_S(M \otimes_R P, N),$$

where P is an arbitrary S-Module.
[Mnemonic: Just move P across the comma.]

[Remark: There is a version of the hom-tensor adjunction for non-commutative rings but there we have to be careful about "bimodules" and left vs. right actions. The most important example of the non-commutative version is "restriction and induction" of group representations.]

///

To show that this abstract stuff is really not so hard, let's prove that $R \otimes_R S \cong S$. Plugging $M = R$ into ⊛ gives

$$\begin{aligned} \mathrm{Hom}_S(R \otimes_R S, N) &\cong \mathrm{Hom}_R(R, \mathrm{Hom}_S(S,N)) \\ &\cong \mathrm{Hom}_S(S,N), \end{aligned}$$

and then "by Yoneda" we conclude that

$$R \otimes_R S \cong S.$$

[Note that the homomorphism $\varphi : R \to S$ was implicit here; that's one weakness of the notation.] Furthermore, since $(-)\otimes_R S$ is a left-adjoint functor by definition, we know from the "Meta-Theorem" that it commutes with colimits. For example, if A is any set then we obtain an isomorphism of S-modules

$$R^{\oplus A} \otimes_R S \cong (R \otimes_R S)^{\oplus A} \cong S^{\oplus A}.$$

[I think this is where category theory starts to spit out more than we put into it.]

That was enough theory. Let's finally apply this (relatively speaking, of course) to the theory of modules over an integral domain.

So let R be an integral domain and consider an R-Module M. Let $A \subseteq M$ be a maximal R-linearly independent set and consider the corresponding free submodule

$$R^{\oplus A} \subseteq M.$$

By maximality, the quotient $M/R^{\oplus A}$ is a torsion module. [Proof: If not then there exists $m \notin R^{\oplus A}$ with $rm \notin R^{\oplus A}$ for all $r \in R$. But then $A \cup \{m\}$ is R-linearly independent.]

Now consider the short exact sequence

$$0 \longrightarrow R^{\oplus A} \longrightarrow M \longrightarrow M/R^{\oplus A} \longrightarrow 0.$$

Since a cokernel is a colimit it will be preserved under extension of scalars. In particular, consider the inclusion of R in its field of fractions $R \rightarrow K := \mathrm{Frac}(R)$.

Then we obtain an exact sequence of K-vector spaces

$$K^{\oplus A} \longrightarrow M \otimes_R K \rightarrow (M/R^{\oplus A}) \otimes_R K \longrightarrow 0 .$$

In general the left arrow may no longer be injective but here it will be [because "field of fractions" is an example of localization which is always exact]. Also, since $M/R^{\oplus A}$ is torsion over R it's definitely torsion over K, which implies that $(M/R^{\oplus A}) \otimes_R K = 0$. Thus we obtain an exact sequence

$$0 \longrightarrow K^{\oplus A} \longrightarrow M \otimes_R K \longrightarrow 0 ,$$

which tells us that $M \otimes_R K \cong K^{\oplus A}$. By Steinitz Exchange over the field K we conclude that the cardinality of A is well-defined. We call it the <u>rank</u> of M.

[Yes, we already proved this on HW2 but don't you think this proof is better?]

///

Next assume that R is a PID and that M is finitely generated so we have an expansion in terms of elementary divisors,

$$(*) \qquad M \cong R^{\oplus k} \bigoplus_{i,j} R/(p_i^{\alpha_{ij}}) .$$

We just saw that the rank k satisfies

$$k = \dim_{\mathrm{Frac}(R)} (M \otimes_R \mathrm{Frac}(R)),$$

and so is well-defined. More generally, if $P \subseteq R$ is any prime ideal then R/P is a domain and we can consider its field of fractions. We define the <u>P-rank</u> of M as the (well-defined) number

$$\dim_{\mathrm{Frac}(R/P)} (M \otimes_R \mathrm{Frac}(R/P)).$$

We already know that the ideals (p_i) from $(*)$ are well-defined (they are the prime ideals containing the annihilator ideal $\mathrm{Ann}_R(M)$). Also, since R is a PID we know that each (p_i) is a <u>maximal</u> ideal, so that

$$\mathrm{Frac}(R/(p_i)) = R/(p_i) .$$

By tensoring with the fields $R/(p_i)$ we hope to recover all of the numbers α_{ij}.

First I will state without proof that for all elements $a \in R$ we have

$$\frac{R}{(a)} \otimes_R \frac{R}{(p_i)} \cong \begin{cases} R/(p_i) & \text{if } a = p^{\alpha} \text{ for some } \alpha \geq 1 \\ 0 & \text{otherwise} \end{cases}$$

This implies that the "(p_i)-rank" of M

$$\dim_{R/(p_i)} \left(M \otimes_R R/(p_i) \right)$$

equals $\# \{ j : \alpha_{ij} \neq 0 \}$. So these numbers are well-defined.

Finally, I will state without proof that

$$\frac{p_i^{\beta} R}{(p_i^{\alpha})} \otimes_R \frac{R}{(p_i)} \cong \begin{cases} R/(p_i) & \text{if } \beta < \alpha \\ 0 & \text{if } \beta \geq \alpha \end{cases}$$

Now consider the (well-defined) "(p_i)-torsion submodule" of M,

$$T_{(p_i)} = \bigoplus_j R/(p_i^{\alpha_{ij}}).$$

By applying p_i^β on both sides and then tensoring with the field $R/(p_i)$ we obtain

$$\dim_{R/(p_i)} \left(p_i^\beta T_{(p_i)} \otimes_R R/(p_i) \right).$$

$= \# \{ j : \alpha_{ij} > \beta \}$. Then by letting β run over all natural numbers this gives us enough invariants to determine the numbers α_{ij} uniquely.

This completes the proof of the Fundamental Theorem of Finitely Generated Modules over a Principal Ideal Domain.

QED.

In MTH 762 I have succeeded in convincing myself that the categorical view of linear algebra is useful, interesting, and intelligible to ordinary human beings. Hopefully I convinced a few of you as well.