Standard Monomial Theory for GL_n

1. Representations of groups; notions from linear algebra.

Representation theory can be thought of from a variety of perspectives. One perspective is, that representation theory describes vector spaces with a distinguished set of symmetries. This leads to constructions that parallel most of the constructions of linear and multilinear algebra.

1.1: Definition of a group representation. A *representation* ρ of a group G is a homomorphism from G into $GL_n(\mathbf{C})$ for some n. More abstractly, a *representation* of G on a complex vector space V is a homomorphism

$$\rho : G \to GL(V),$$

where $GL(V)$ denotes the group of all invertible linear transformations on V.

1.2 Example: Characters. Suppose that V is one-dimensional. Then for a non-zero element $\vec{v}_o$ of V, we must have $\rho(g)(v_o) = \gamma(g)v_o$, for some scalar $\gamma(g)$ depending on $g \in G$. The condition that ρ is a representation implies that γ is multiplicative in g:

$$\gamma(g_1g_2)\vec{v}_o = \rho(g_1g_2)(\vec{v}_o) = \rho(g_1)(\rho(g_2)(\vec{v}_o)) = \rho(g_1)(\gamma(g_2)\vec{v}_o) = \gamma(g_2)\rho(g_1)(\vec{v}_o) = \gamma(g_2)(\gamma(g_1)\vec{v}_o))$$

$$= (\gamma(g_1)\gamma(g_2))\vec{v}_o.$$

The fourth equation follows from linearity of $\rho(g_1)$, and the last equation appeals to commutativity of multiplication in $\mathbf{C}$. Comparing coefficients of $\vec{v}_o$, we see that

$$\gamma(g_1g_2) = \gamma(g_1)\gamma(g_2).$$

A function γ satisfying this multiplicative property is called a *character* of G. More specifically, we will call γ the *eigencharacter* attached to the vector $\vec{v}_o$. Evidently, having a character of a group G is equivalent to having a one-dimensional representation of G.

1.3: Equivalence of representations; intertwining operators. Given two representations $\rho : G \to GL(V)$ and $\rho' : G \to GL(V')$, a linear mapping $T : V \to V'$ is called an *intertwining operator* for ρ and ρ' (aka, a G-morphism from V to V') if it mediates between the two actions of G:

$$T\rho(g) = \rho'(g)T, \tag{1.3.1}$$

for all $g \in G$.

If there is an intertwining operator for ρ and ρ' that is a linear isomorphism between V and V', then ρ and ρ' are called *equivalent* representations.

A main task of representation theory is to describe the equivalence classes of representations of G, for a given group G.

1.4: Subrepresentations. Given a representation $\rho : G \to GL(V)$ of G on V, suppose that $U \subset V$ is a subspace of V. If $\rho(g)(U) \subset U$ for all g in G, we say that U is *invariant* under (the action of) G. If $U \subset V$ is a subspace of V that is invariant under G, then we can define a representation $\rho' = \rho_U$ by

$$\rho_U(g)(u) = \rho(g)(u),$$

for any g in G and any vector $u \in U$. The representation ρ' is called a *subrepresentation* of ρ. Precisely, it is the subrepresentation of ρ defined by (the G-invariant subspace) U.

1.5: Irreducible representations. A representation ρ of a group G is *irreducible* if the only subrepresentations are the zero subspace and the whole space.

The set of all isomorphism/equivalence classes of irreducible (finite dimensional) representations of G is denoted $\hat{G}$. Usually, an element of $\hat{G}$ is called an "irreducible representation" rather than "a class of equivalent irreducible representations".

1.6: Quotient representations. If $\rho : G \to GL(V)$ is a representation of G on V, and $U \subset V$ is a G-invariant subspace, then an element $\rho(g)$ will define a linear transformation $\overline{\rho}(g)$ on the quotient vector space V/U, by the recipe

$$\overline{\rho}(g)(v + U) = \rho(g)(v) + U.$$

It can be checked that $\overline{\rho} : G \to GL(V/U)$ is a representation of G on V/U. It is called the *quotient representation* of ρ defined by ρ_U.

1.6.1: Proposition. If $\rho : G \to GL(V)$ and $\sigma : G \to GL(U)$ are representations of G, and if $T : V \to U$ is an intertwining operator for ρ and σ, then the kernel $\ker T$ defines a subrepresentation of V, and the image $T(V)$ of V under T is a subrepresentation of U, and T factors to define an isomorphism $\overline{T} : V/\ker T \simeq T(V)$.

Proof: Most of the assertions of this are simply linear algebra. We should check that $\ker T$ is G-invariant. If $\vec{v}$ is in $\ker T$, then $T(\vec{v}) = 0$. For $g \in G$, since T is an intertwining operator, we have $T(\rho(g)(\vec{v})) = \sigma(g)(T(\vec{v})) = \sigma(g)(0) = 0$. Thus, $\rho(g)(\vec{v})$ is also in $\ker T$, as desired. Checking that T is an intertwining from $V/\ker T$ to $T(V)$ is similar.

1.6.2: Corollary. (Schur's Lemma) a) If ρ and σ are irreducible representations, and $T : U \to V$ is a non-zero intertwining operator, then T is an equivalence of representations.

b) Moreover, T is unique up to scalar multiples. In particular, if U is an irreducible representation of G, then any intertwining operator from U to itself is a multiple of the identity operator.

1.7: Composition series; multiplicities. Given a representation $\rho : G \to GL(V)$, with V finite dimensional, if V is not irreducible, then we can find a non-trivial G-invariant subspace $V_1 \subset V$. Among all such subspaces, we can look for one of minimal dimension. Such a V_1 then necessarily defines an irreducible subrepresentation of V. In particular, we can find irreducible subrepresentations of V.

In the context of the previous paragraph, consider the quotient representation on $\overline{V} = V/V_1$. We can similarly find a subspace $\overline{V}_2 \subset V/V_1$ such that $\overline{V}_2$ is an irreducible subrepresentation of V/V_1. Lifting $\overline{V}_2$ to a subspace V_2 with $V_1 \subset V_2 \subset V$, we see that V_2 will define a subrepresentation of V_1, with irreducible subrepresentation V_1, and an irreducible quotient representation V_2/V_1.

Continuing in this fashion, we can find a nested sequence of subspaces

$$V_1 \subset V_2 \subset V_3 \subset \ \ldots \ \subset V_m = V,$$

such that

i) each V_i is G-invariant; and

ii) the representation of G on V_i/V_{i-1} is irreducible.

Such a sequence of subspaces of V, or sometimes the corresponding sequence of irreducible representations, is called a *composition series* for ρ.

Given a composition series $\{V_j\}$ for the representation ρ, each subquotient representation $\overline{\rho}_i$ on V_i/V_{i-1} being irreducible, it defines a point in $\hat{G}$. The number of indices i such that $\overline{\rho}_i$ belongs to a given isomorphism class σ in $\hat{G}$ is called the *multiplicity* of σ in ρ. The Jordan-Hölder Theorem says that the multiplicity of a given (isomorphism class of) irreducible representation in ρ is the same for any two composition series for ρ. In other words, the multiplicities are invariants of the representation ρ.

If two representations have the same multiplicity for every irreducible representation, we will call them "numerically equivalent", or "Grothendieck equivalent".

1.7: Dual or contragredient representation. Let V be a vector space, and let V^* be its dual space. Recall that, for a linear operator $T \in End(V)$, the *adjoint* to T is the linear map $T^* : V^* \to V^*$ satisfying

$$T^*(\lambda)(\vec{v}) = \lambda(T(\vec{v}), \tag{1.7.1}$$

for $\lambda \in V^*$ and $v \in V$. The mapping $T \to T^*$ is antimultiplicative;

$$(TS)^* = S^*T^*, \tag{1.7.2}$$

for T and S in $End(V)$. This is an easy direct verification.

If $\rho : G \to GL(V)$ is a representation of G on V, then the *contragredient* or *dual* representation of ρ is the representation $\rho^* : G \to GL(V^*)$ defined by

$$\rho^*(g)(\lambda) = (\rho(g)^*)^{-1}(\lambda) \tag{1.7.3}$$

for $\lambda \in V^*$

Recall that, given a vector space V and a subspace U, then the subspace

$$U^\perp \subset V^* = \{\lambda \in V^* : \lambda(u) = 0, \ \text{all} u \in U,$$

is called the *annihilator* of U. We have the dimension formula

$$\dim U + \dim U^\perp = \dim V =\sim V^* \tag{1.7.4}$$

Moreover,

$$U^* \simeq V^*/U^\perp \qquad (1.7.5).$$

1.7.6: Proposition. If $\rho : G \to GL(V)$ is a representation of G, and $U \subset V$ is a subrepresentation, then $U^\perp$ is a subrepresentation of V^*. Moreover, the natural action of G on $V^*/U^\perp$ is the representation contragredient to ρ_U.

1.8: Tensor products of representations. If V and W are (complex) vector spaces, then we know how to form the tensor product $V \otimes W$. It is the universal target for bilinear maps on $V \times W$: there is a natural bilinear mapping

$$\beta : V \times W \to V \otimes W,$$

such that, if $B : V \times W \to U$ is a bilinear mapping from $V \times W$ to another vector space U, then there is a linear mapping

$$\tilde{B} : V \otimes W \to U \quad \text{such that} \quad \tilde{B} \circ \beta = B. \tag{1.8.1}$$

If $S \in End(V)$ and $T \in End(W)$ are linear transformations on V and W respectively, then $\beta \circ (S \times T) : (\vec{v}, \vec{w}) \to (S(\vec{v}) \otimes T(\vec{w})$ is a bilinear map from $V \times W$ to $V \otimes W$. Thus the universal property of $V \otimes W$ gives us a linear mapping

$$S \otimes T : V \otimes W \to V \otimes W. \tag{1.8.2}$$

The map $S \otimes T$ depends bilinearly on S and T, so it gives us a linear map

$$\beta : End(V) \otimes End(W) \to End(V \otimes W).$$

It is not hard to verify that this β is an isomorphism. In addition, the mapping β is multiplicative, in the sense that

$$\beta(S_1 \otimes T_1)\beta(S_2 \otimes T_2) = \beta(S_1 S_2 \otimes T_1 T_2). \tag{1.8.3}$$

Using this formula, the subspaces $\beta(End(V) \otimes I_W)$, and $\beta(I_V \otimes End(W))$ can be checked to be subalgebras of $End(V \otimes W$. Moreover, these subalgebras commute with each other.

1.8.4: Proposition. The subalgebras $\beta(End(V) \otimes I_W$, and $\beta(I_V \otimes End(W))$ in $End(V \otimes W)$ are mutual commutants: each subalgebra is the full subalgebra of $End(V \otimes W)$ commuting with the other.

Now suppose that $\rho : G \to GL(V)$ and $\sigma : H \to GL(W)$ are representations of two groups, G and H, on vector spaces V and W. We can define a representation $\rho \otimes \sigma$ of $G \times H$ on $V \otimes W$ by the recipe

$$\rho \otimes \sigma(g, h) = \rho(g) \otimes \sigma(h).$$

The representation $\rho \otimes \sigma$ is called the (outer) *tensor product* of the representations ρ and σ.

If $H = G$, we can consider the diagonal subgroup $\Delta(G) = \{(g, g)\}, g \in G$ of $G \times G$. The restriction $\rho \otimes \sigma_{|\Delta(G)}$ is a representation of G. It is called the (inner) tensor product of ρ and σ.

1.8.5: Theorem. If $\rho \in \hat{G}$ and $\sigma \in \hat{H}$, then $\rho \otimes \sigma$ is an irreducible representation of $G \times H$. Moreover the mapping

$$\tau : \hat{G} \times \hat{H} \to G \hat{\times} H$$

defined by $(\rho, \sigma) \to \rho \otimes \sigma$ is a bijection.

Proof: This follows from Wedderburn theory, described below.

1.9: Direct Sums. Given representations $\rho : G \to GL(V)$ and $\sigma : G \to GL(U)$, the *direct sum* of ρ and σ is the representation $(\rho \oplus \sigma) : G \to V \oplus U$ defined by $(\rho \oplus \sigma)(g) = \rho(g) \oplus \sigma(g)$, where $A \oplus B$ is the direct sum of operators $A \in End(V)$ and $B \in End(U)$. Clearly then, V and U are G invariant subspaces of $V \oplus U$, and evidently $V \oplus U$ is the direct sum of these two subspaces. The converse of this is also true.

1.9.1: Scholium. Suppose that $\rho : G \to GL(V)$ is a representation of G, and that U and U' are subspaces of V invariant under G. Suppose also that $V \simeq U \oplus U'$. Then ρ is equivalent to the direct sum of the subrepresentations ρ_U and $\rho_{U'}$:

$$\rho \simeq \rho_U \oplus \rho_{U'}.$$

1.9.3: Definition: Indecomposability. A representation $\rho : G \to GL(V)$ is *indecomposable* if V cannot be written as a direct sum of two G invariant subspaces.

1.10 Completely decomposable representations. If a representation is (equivalent to) a direct sum of irreducible representations, we call it *completely decomposable*.

1.10.1: Scholium. a) A completely decomposable representation is determined up to equivalence by the multiplicities of irreducible representations in it.

b) Any representation is numerically equivalent to a unique (up to equivalence) completely decomposable representation.

1.10.2: Proposition. TFAE:

a) A representation $\rho : G \to GL(V)$ is completely decomposable.

b) For any irreducible subrepresentation $U \subset V$, there is a subrepresentation U' such that $V \simeq U \oplus U'$.

c) For any subrepresentation $U \subset V$, there is a subrepresentation U' such that $V \simeq U \oplus U'$. (We call U' a *complement* to U in V.)

Proof: Suppose that V is completely decomposable. Write $V = \oplus j V_j$, where the V_j are irreducible. Let P_j be the projection of V to V_j, with kernel equal to the sum of the V_k, for $k \neq j$. Then each P_j is an intertwining map from V to itself (or, to V_j). Evidently $P_j P_k = 0$ if $j \neq k$, and $\sum_j P_j = I$, the identity map on V.

Let U be any irreducible subrepresentation of V. Choose j such that $P_j(U) \neq 0$. By the scholium just above, since U and V_j are irreducible, we can see that P_j must be an isomorphism from U to V_j.

Let $U_j' = (I - P_j)(V)$ be the sum complementary to V_j. We claim that U_j' is also complementary to U. To know this, it is enough to know that $U_j' \cap U = \{0\}$. But since $U_j' = \ker P_j$, and P_j is an isomorphism on U, this is clear. Thus, U is complemented in V.

Now let's assume that any irreducible subrepresentation is complemented in V, and let's show that any subrepresentation, irreducible or not, is complemented.

First, let us show that if $U \subset V$ is a subrepresentation, and if $W \subset U$ is an irreducible subrepresentation of U, then W is complemented in U.

We know that W is complemented in V. Let W' be a complement to W in V. We claim that $W' \cap U$ is a complement to W in U. Indeed, it is clear that $(W' \cap U) \cap W \subset W' \cap W = \{0\}$, so $W' \cap U$ intersects W trivially. But also, the dimension formula tells us that $\dim(W' \cap U) \geq \dim W' + \dim U - \dim V = (\dim V - \dim W) + \dim U - \dim V = \dim U - \dim W$, which means that it has the right dimension to be complementary to W. Thus, $W' \cap U$is a complement to W in U.

Let $U \subset V$ be a subrepresentation. Let W be an irreducible subrepresentation of U. Then W is also an irreducible subrepresentation of V, so W is complemented in V. Let W' be a subrepresentation of V complementary to W. Then $W' \cap U$ is a complement to W in U. Suppose that we can find a complement Y to $W' \cap U$ in W'. Then it is easy to see that Y is a complement to U in V. Hence it enough to show that

$U \cap W'$ is complemented in W'. But we have just seen that any submodule of V, in particular W', has the property that any irreducible submodule is complemented. Hence we may argue by induction on dimension that $U \cap W'$ is complemented in W'; then the above lets us conclude that, indeed, U is complemented in V.

Finally, it is easy to see that, if any submodule of V is complemented, then V is completely reducible. Indeed, suppose that we have found irreducible subrepresentations $V_1, V_2, \ . \ . \ . \ , V_\ell$ such that the span $U = \sum_{j=1}^{\ell} V_j$ is in fact a direct sum. Clearly U is a subrepresentation of V, so if U is not all of V, we can find a complement U' to U in V. Then let $V_{\ell+1}$ be an irreducible submodule of U'. Then it is clear that the span $\tilde{U} = \sum_{j=1}^{\ell+1}$ is again a direct sum. We can continue augmenting a direct sum of subrepresentations in this way until we exhaust V, so V itself is a direct sum of irreducible subrepresentations. This concludes the proof of the proposition.

1.10.3: Corollary. i) Any subrepresentation of a completely decomposable representation is completely decomposable.

ii) A direct sum of completely decomposable representations is completely decomposable.

1.10.4: Isotypic Components. If $\rho : G \to GL(V)$ is a completely decomposable representation of V, and $\sigma \in \hat{G}$ is an irreducible representation, let V_σ be the span of all subrepresentations of V that are equivalent to σ. We call V_σ the *σ-isotypic component* of V.

1.10.5: Lemma. if $W \subset V_\sigma$ is a G-subrepresentation, then $\rho_W \simeq \sigma$.

Proof: Indeed, if $V_\sigma \simeq \oplus_j U_j$, where each $U_j \simeq \sigma$, let P_j be projection to U_j, with kernel equal to the other U_k, $k \neq j$. Then $\sum_j P_j = I_{V_\sigma}$, so $P_j(W) \neq 0$ for some j. Then Schur's Lemma guarantees that $W \simeq U_j \simeq \sigma$.

1.10.6: Canonical isotypic decomposition. Let U be a complementary subrepresentation to V_σ. Then U is again completely reducible into irreducible components, and none of these components can be equivalent to σ, by definition of V_σ. Therefore, by induction on dimension, we may assert that U is the direct sum of its isotypic components, all of which will correspond to representations other than σ. These are therefore also isotypic components in V. For suppose τ is any representation not equivalent to σ, and W is a subspace of V equivalent to τ, and W is not contained in U. Let P be the projection onto V_σ with kernel U. Then P will define an embedding of W into V_σ, which entails that $W \simeq \sigma$, as seen above. Since W was assumed not to be equivalent to σ, we see therefore that W is required to be contained in U.

Therefore, adding V_σ to U, we see that V is the direct sum of its isotypic components. Note that this decomposition is unique, that is, it is dictated by the structure of V as a representation for G. There are no choices involved.

For each (equivalence class of) irreducible representation(s) σ of G, choose a standard representative, also denoted by σ. Then given a representation of G on a vector space V,, and a subrepresentation U that is isomorphic to σ, there is a G intertwining map $T_U : \sigma \to U \subset V$. Schur's Lemma tells us that this T_U is determined up to scalar multiples by U. Conversely, any non-zero G intertwining map $T : \sigma \to V$ will have as image a subrepresentation U_T of V isomorphic to σ. Thus, if we let $Hom_G(\sigma, V)$ denote the vector space of all G intertwining maps from σ to V, there is a natural mapping

$$\sigma \otimes Hom_G(\sigma, V) \to V,$$

whose image will be V_σ, the σ-isotypic component of V. Combining these maps from the various irreducible representations of V, we obtain a canonical description of V as a sum of its isotypic components, each of which is a tensor product, as just specified in formula ????:

$$V \simeq \oplus_{\sigma \in \hat{G}} V_\sigma \simeq \oplus_{\sigma \in \hat{G}} \ \sigma \otimes Hom_G(\sigma, V). \tag{1.10.7}$$

1.11: Wedderburn Theory.

If $\rho : G \to GL(V)$ is a representation of G, let A_ρ be the linear span of the operators $\rho(g)$ for $g \in G$. Then A_ρ is easily checked to be an algebra, the *enveloping algebra* of ρ. Evidently, a subspace $U \subset V$ is invariant under $\rho(G)$ if and only if it is invariant under A_ρ. In particular, V is completely reducible for G if and only if it is completely reducible for A_ρ.

For a subalgebra $A \subset End(V))$, let A' denote the *commutant* of A: the algebra of all operators that commute with all operators in A. It is easy to see that $A \subset A'' = (A')'$, the commutant of A', and that $A''' = A'$. When $A = A_\rho$, the commutant A' is also describable as the set of all G intertwining operators from V to itself.

Theorem: a) If ρ is irreducible, the $A_\rho = End(V)$ is the full algebra of linear transformations on V.
b) If ρ is completely reducible, then $(A'_\rho)' = A_\rho$; that is, A_ρ is its own double commutant.
c) If $V \simeq \oplus_{\sigma\in\hat{G}}\, \sigma \otimes Hom_G(\sigma, V)$ is the isotypic decomposition for ρ, then

$$A\rho \simeq \oplus_{\sigma\in\hat{G}; V_\sigma\neq\{0\}} End(\sigma), \qquad \text{and} \qquad A'_\rho \simeq \sum_{\sigma\in\hat{G}} End(Hom_G(\sigma, V)).$$

Moreover, the joint action of A_ρ and A'_ρ on each V_σ is irreducible.

Proof: See Lang, Algebra.

1.12: Regular representation. Let $\mathbf{C}(G)$ be the vector space of all complex-valued functions on G. Then G can act on $\mathbf{C}(G)$ by right translation and by left translation. Precisely, we define, for $f \in \mathbf{C}(G)$ and $g \in G$,

$$L(g)(f)(h) = f(g^{-1}h) \qquad \text{and} \qquad R(g)(f)(h) = f(gh),$$

for any $h \in G$.

It is easy to check that $g \to L(g)$ and $g \to R(g)$ are representations of G on $\mathbf{C}(G)$, that is $L(g)L(g') = L(gg')$ and $R(g)R(g') = R(gg')$. They are called the *left regular representation* and the *right regular representation* respectively. Of course, if G is not finite, $\mathbf{C}(G)$ is not finite-dimensional. Nevertheless, it is useful to consider the regular representations. In various circumstances, especially that of algebraic groups, it is also useful to consider various special classes of functions on G that define subrepresentations of L and R.

1.13: Matrix coefficients. Let $\rho : G \to GL(V)$ be a representation of G on V. If $\vec{v} \in V$ and $\lambda \in V^*$, the vector space dual of V, then a *matrix coefficient* of ρ is a function

$$\phi_{\lambda,\vec{v}}(g) = \lambda(\rho(g)(\vec{v})). \tag{1.13.1}$$

If $V = \mathbf{C}^n$ and $e_j : 1 \leq j \leq n$ is the standard basis for $\mathbf{C}^n$ and e_j^* is the standard basis for the dual $(\mathbf{C}^n)^*$ (which we may think of as the row vectors of length n), then the matrix coefficients $\phi_{e_j^*,e_k}$ are exactly the entries of the matrix $\rho(g)$.

It is easy to check that the matrix coefficient $\phi_{\lambda,\vec{v}}$ depends bilinearly on $\vec{v}$ and on λ:

$$\phi_{\lambda,\vec{v}+\vec{u}} = \phi_{\lambda,\vec{v}} + \phi_{\lambda,\vec{u}}, \tag{1.13.2a}$$

and

$$\phi_{\lambda+\mu,\vec{v}} - \phi_{\lambda,\vec{v}} + \phi_{\mu,vecv}, \tag{1.13.2b}$$

for vectors $\vec{v}$ and $\vec{u}$ in V, and linear functions λ and μ in V^*.

Matrix coefficients connect a general representation of G with the right and left regular representations. With notations as above, we can compute that

$$R(h)(\phi_{\lambda,\vec{v}})(g) = \phi_{\lambda,\vec{v}}(gh) = \lambda(\rho(gh)\vec{v}) = \lambda((\rho(g)\rho(h))\vec{v}) = \lambda(\rho(g)(\rho(h)\vec{v})) = \phi_{\lambda,\rho(h)\vec{v}}. \tag{1.13.3a}$$

and

$$\begin{aligned} L(h)(\phi_{\lambda,\vec{v}})(g) &= \phi_{\lambda,\vec{v}}(h^{-1}g) = \lambda(\rho(h^{-1}g)\vec{v}) = \lambda(\rho(h^{-1})\rho(g)\vec{v}) \\ &= \rho(h^{-1})^*(\lambda)(\rho(g)\vec{v}) = \rho^*(h)(\lambda)(\rho(g)\vec{v}) = \phi_{\rho^*(h)\lambda,\vec{v}}(g). \end{aligned} \tag{1.13.3b}$$

Thus, if we fix λ in V^*, and define $\Phi_\lambda : V \to \mathbf{C}(G)$ by the recipe

$$\Phi_\lambda(v)(h) = \lambda(\rho(h)(v)), \tag{1.13.4a}$$

the formulas of the preceding paragraph imply that Φ_λ is an intertwining operator:

$$R(g)(\Phi_\lambda(v)) = \Phi_\lambda(\rho(g)(v)). \tag{1.13.4b}$$

Conversely, suppose that $\Phi : V \to \mathbf{C}(G)$ intertwines ρ with R. For v in V, define $\lambda_\Phi(v) = \Phi(v)(1)$, where 1 indicates the identity element of G. Then $\lambda_\Phi = \lambda$ is a linear functional on V, and we can check that $\Phi_\lambda = \Phi$. Thus we see:

1.13.5: Proposition. The mapping $\lambda \to \Phi_\lambda$ defines an isomorphism from V^* to the space $Hom_G(V, \mathbf{C}(G))$, with G acting on $\mathbf{C}(G)$ by the right regular representation. In particular, every irreducible representation is equivalent to a subrepresentation of $\mathbf{C}(G)$.

As already mentioned, the matrix coefficient $\phi_{\lambda,\vec{v}}$ is bilinear in λ and $\vec{v}$. We can extend it to a linear map

$$\Phi : End(V) \to \mathbf{C}(G)$$

as follows. For $\vec{v}$ in V and λ in V^*, we can form the rank one operator

$$E_{\lambda,\vec{v}}(\vec{u}) = \lambda(\vec{u})\vec{v}. \tag{1.13.6a}$$

The mapping $(\lambda, \vec{v}) \to E_{\lambda,\vec{v}}$ is a bilinear map from $V^* \times V$ to the rank one operators on V. We have the relation

$$tr(E_{\lambda,\vec{v}}) = \lambda(\vec{v}), \tag{1.13.6b}$$

where $tr(T)$ indicates the trace of the linear operator T. (In other words, tr is the linear function on $End(V) \simeq V *^{\otimes} V$ that extends the natural bilinear pairing between V and V^*.) This says that we may write the matrix coefficient $\phi_{\lambda,\vec{v}}$ in the alternative form $\phi_{\lambda,\vec{v}}(g) = tr(\rho(g)E_{\lambda,\vec{v}})$. This allows us to extend the matrix coefficient mapping to a canonical linear map

$$\Phi : End(V) \to \mathbf{C}(G) \tag{1.13.7a}$$

by the recipe

$$\Phi(T)(g) = tr(\rho(g)T. \tag{1.13.7b}$$

The analogs of the intertwining formulas above for this extended mapping are

$$R(h)(\Phi(T)) = \Phi(\rho(h)T), \qquad \text{and} \qquad L(h)(\Phi(T)) = \Phi(T\rho(h^{-1}. \tag{1.3.7c}$$

The paraphrase of Proposition 1.13.5 for this extended matrix coefficient mapping is

1.13.8: Proposition. If ρ is an irreducible representation of G, the image $\Phi(End(V))$ in $\mathbf{C}(G)$ is the ρ -isotypic component of $\mathbf{C}(G)$ under the right regular representation R. In particular, ρ appears in $\mathbf{C}(G)$ with multiplicity equal to $\dim V$. The ρ isotypic component for R coincides with the ρ^* isotypic component for L.

We need a slight extension of matrix coefficients for some arguments. A given matrix coefficient mapping Φ_λ from a representation ρ on a space V will of course be an isomorphism if ρ is irreducible, by Schur's Lemma. However, if ρ is reducible, a given Φ_λ may have a kernel, and indeed, there may not exist a single λ such that the associated matrix coefficient map is faithful. However, if we select enough λs (for example, if we let λ run through a basis for V^*), then the intersections of the kernels of all the Φ_λ will be reduced to $\{0\}$. We can think of a collection λ_a of elements of V^* as defining an intertwining map from V to the direct sum of several copies of $\mathbf{C}(G)$, under the diagonal right regular representation. Hence we may assert:

1.13.9: Lemma. Any finite dimensional representation of G can be realized as a subrepresentation of a sum of copies of $\mathbf{C}(G)$.

2. Matrix groups.

2.1: Definition. A *matrix group* is a closed subgroup of $GL_n(\mathbf{R})$.

2.1.1: Examples. i) $GL_n(\mathbf{R})$ itself is a matrix group.

ii) The subset $B_n = B$ consisting of all (invertible) upper triangular matrices is a matrix group.

iii) The orthogonal group, the set of isometries of the standard Euclidean norm, equivalently, the linear transformations that preserve the inner product, is an algebraic group. If we write the inner product

$$\vec{x} \bullet \vec{y} = \sum_{j=1}^{n} x_j y_j,$$

then the condition that a linear transformation g preserve the inner product is that

$$g(\vec{x}) \bullet g(\vec{y}) = \vec{x} \bullet \vec{y}$$

for all $\vec{x}$ and $\vec{y}$ in $\mathbf{R}^n$. This is clearly a set of algebraic equations defining the orthogonal group. Since $\vec{x} \bullet \vec{y} = \vec{y}^t \vec{x}$, where $\vec{y}^t$ indicates the transpose of $\vec{y}$, the equations defining the orthogonal group can also be written in the form $\vec{y}^y \vec{x} = (g(\vec{y}))^t (g(\vec{x})) = \vec{y}^t g^t g \vec{x}$. If we let $\vec{x}$ and $\vec{y}$ vary arbitrarily in $\mathbf{R}^n$, we can see that this implies that $g^t g = I$, the identity operator. In other words, the orthogonal group can also be defined by the equations $G^t g = I$, or $g^t = g^{-1}$, or $(g^t)^{-1} = g$

iv) The invertible diagonal matrices

$$A_n = \left\{ \begin{bmatrix} a_1 & 0 & 0 & \dots & 0 \\ 0 & a_2 & 0 & \dots & 0 \\ 0 & 0 & a_3 & \dots & 0 \\ . & . & . & \dots & . \\ . & . & . & \dots & . \\ . & . & . & \dots & . \\ 0 & 0 & 0 & \dots & a_n \end{bmatrix} \right\},$$

with $a_i \neq 0$, form a matrix group.

v) The unipotent upper triangular matrices U_n, consisting of upper triangular matrices with 1s on the diagonal, is a matrix group. Both A_n and U_n are subgroups of B_n. The unipotent subgroup U_n is a normal subgroup of B_n, and B_n is the semidirect product of U_n and A_n.

vi) Let E_{ij} be the *matrix units*: the matrix E_{ij} has a 1 in the i-th row and j-th column, and zeroes everywhere else. In formulas, we can write $(E_{ij})_{ab} = \delta_{ia}\delta_{bj}$, where $(E_{ij})_{ab}$ denoted the (a, b) entry of E_{ij}, and δ_{ab} is Kronecker's delta. The matrix units satisfy the product relations

$$E_{ij} E_{k\ell} = \delta_{jk} E_{i\ell}.$$

If $i \neq j$, we can define

$$\Gamma_{ij}(s) = I + E_{ij}(s),$$

where I here denotes the identity matrix and s is a real number. Then the group $\Gamma_{ij} = \{\Gamma_{ij}(s) : s \in \mathbf{R}\}$ is a matrix group. Indeed, it is a *one-parameter group*, meaning that the map $s \to \Gamma_{ij}(s)$ is a group homomorphism from $\mathbf{R}$ to Γ_{ij}. Indeed, since $i \neq j$, we see that $E_{ij}^2 = 0$, so that

$$\Gamma_{ij}(s)\Gamma_{ij}(t) = (I + sE_{ij})(I + tE_{ij}) = I + (s + t)E_{ij} + stE_{ij}^2 = I + (s + t)E_{ij}.$$

The Γ_{ij} are called *root groups*. They are normalized by the diagonal matrices. If $\mathbf{a}$ is a diagonal matrix as in iv), then we have

$$\mathbf{a}\Gamma_{ij}(s)\mathbf{a}^{-1} = \Gamma_{ij}(a_i a_j^{-1} s).$$

2.2: Lie algebra of (real) matrices. A Lie algebra g of real matrices is a subspace of $M_n(\mathbf{R})$ that is closed under the bracket or bracket operation: if A, B belong to g, then the commutator $[A, B] = AB - BA$ also belongs to g.

2.3: Lie algebra of a matrix group

2.3.1: Theorem. ([ ]) A matrix group is a smooth submanifold inside $GL_n(\mathbf{R})$. In particular, it has a tangent space g at the identity (meaning, the identity matrix in $GL_n(\mathbf{R})$). The space g is a Lie algebra of matrices. It is called the *Lie algebra of G*.

2.3.2: Examples. i) The Lie algebra $gl_n(\mathbf{R})$ of $GL_n(\mathbf{R})$ is the full matrix algebra $M_{nn}(\mathbf{R})$ of square $n \times n$ matrices. Since this is closed under matrix multiplication, it is a fortiori closed under commutator.

ii) The Lie algebra b_n of the upper triangular group B_n is the space of upper triangular matrices. This is an associative subalgebra of $M_{nn}(\mathbf{R})$, so like the full algebra, it is clearly closed under commutator bracket.

iii) The Lie algebra o_n of the orthogonal group can be checked, following Theorem 3 below, to consist of the matrices A such that $(A\vec{x}) \bullet \vec{y} + \vec{x} \bullet (A\vec{y}) = 0$, or $\vec{y}^t A^t \vec{x} + \vec{y}^t A \vec{x} = 0$. Again, if we let $\vec{x}$ and $\vec{y}$ vary over all $\mathbf{R}^n$, we find that this implies that $A^t + A = 0$ or $A^t = -A$; in other words, A is skew-symmetric.

iv) The Lie algebra a_n of the group A_n of invertible diagonal matrices is again just the space of all diagonal matrices. As for gl_n and b_n, it is an associative algebra, so a fortiori a Lie algebra.

v) The Lie algebra of a root group Γ_{ij} just consists of the scalar multiples of the matrix unit E_{ij}.

2.4: One parameter groups; exponential map. A *one-parameter group* of matrices is a homomorphism $\mathbf{R} \to GL_n(\mathbf{R})$. (In other words, a one-parameter group is a representation of $\mathbf{R}$.

Given a matrix Z, the exponential of Z, denoted $\exp(sZ)$ is defined by

$$\exp(Z) = \sum_{k=0}^{\infty} \frac{Z^n}{n!} \tag{2.4.1}$$

$$= I + Z + \frac{1}{2}Z^2 + \frac{1}{6}Z^3 + \frac{1}{24}Z^4 + \ \ldots + \frac{1}{k!}Z^k + \ \ldots$$

2.4.2: Theorem. For any matrix Z, the map $t \to \exp(tZ)$ is a one-parameter group of matrices.

2.4.3: Theorem. Any one parameter group of matrices has the form $t \to \exp(tZ)$ for an appropriate matrix Z.

The matrix Z of Theorem 2.3.2 is called the *infinitesimal generator* of the one-parameter group under discussion.

2.4.4: Theorem. a) For a matrix group G, the Lie algebra g of G consists of all matrices Z such that $\exp(tZ)$ belongs to G for all t; that is, the one parameter group generated by Z belongs to G.

b) The exponential map $Z \to \exp Z$ defines a bijection between a neighborhood of 0 in g and a neighborhood of I in G.

2.5: Representations of Lie algebras. If g is a Lie algebra of matrices, and $R : g \to End(V)$ is a linear mapping, then R is called a *representation* of g iff R preserves the bracket operation:

$$[R(Z), R(Y)] = R([Z, Y]), \tag{2.5.1}$$

for any Z and Y in g.

2.5.2: Theorem. If $\rho : G \to GL(V)$ is a representation of the matrix group G, define $D\rho : g \to End(V)$ by the relation

$$\exp(tD\rho(Z)) = \rho(\exp(tZ), \tag{2.5.3a}$$

or alternatively,

$$D\rho(Z) = \frac{d}{dt}_{|t=0}(\rho(\exp tZ) \tag{2.5.3b}$$

for any Z in g. Then $D\rho$ is a representation of g on V.

Remark: $D\rho$ is defined by Theorem 2.5.2. The assertion is that it is linear, and that it preserves the bracket operation.

2.5.4: Definition. We call $D\rho$ the *derived representation* of ρ.

2.5.5: Examples. i) Consider the natural action ν of $GL_n(\mathbf{R})$ on $P(\mathbf{R}^n) = P_n = P$, the algebra of polynomial functions on $\mathbf{R}^n$, defined in the standard way:

$$\nu(g)(p)(\vec{x}) = p(g^{-1}(\vec{x})) \tag{2.5.5.1}$$

for $\vec{x} \in \mathbf{R}^n$. Although P is infinite dimensional, it has a natural grading by the subspaces of polynomials that are homogeneous of a given degree, and these are finite dimensional:

$$P(\mathbf{R}^n) \simeq \oplus_{d=0}^{\infty} P^d(\mathbf{R}^n).$$

The space P^d are defined by the action of the scalar matrices: a polynomial p belongs to P^d if $p(s\vec{x}) = s^d p(\vec{x})$ for all vectors $\vec{x}$ and all scalars s. In other words, P^d is an eigenspace for the action of the scalar matrices. Since the scalar matrices commute with the action of all of GL_n, as is easily checked, the spaces P^d are invariant under GL_n. The space is finite dimensional: it is spanned by the monomials $\vec{x}^{\vec{a}} = x_1^{a_1} x_2^{a_2} \ . \ . \ .x_n^{a_n}$. with $\sum_{j=1}^n a_j = d$. Here

$$\vec{a} = \begin{bmatrix} a_1 \\ a_2 \\ \cdot \\ \cdot \\ \cdot \\ a_n \end{bmatrix}, \tag{2.5.5.2}$$

where the a_j are non-negative integers.

We can compute the derived action $d\nu$ of gl_n on P. By definition

$$d\nu(Z)p(\vec{x}) = \frac{d}{dt}_{t=0}(\nu(exp(tZ))p(\vec{x}) = \lim_{h\to 0} \frac{p(exp(-hZ)\vec{x}) - p(\vec{x})}{h} = \lim_{h\to 0} \frac{p(\vec{x} - hZ\vec{x} + h^2 r(t,Z)\vec{x}) - p(\vec{x})}{h}$$
$$= -\partial_{Z\vec{x}} p(\vec{x}), \tag{2.5.5.3}$$

where $\partial_{\vec{u}} f$ indicates the directional derivative of a function f in the direction $\vec{u}$. In other words, $d\nu(Z)$ is differentiation with respect to the "linear coefficient vector field" $-Z(\vec{x})$.

Some special cases: a) If $Z = A$ is diagonal, with diagonal entries $a_{ii} = a_i$, then

$$d\nu(A) = -\sum_i a_i x_i \frac{\partial}{\partial x_i}.$$

In particular, if $A = I$, the identity matrix, then $d\nu(I) = -\Sigma_i x_i \frac{\partial}{\partial x_i} = -E$, where $E = \Sigma_i x_i \frac{\partial}{\partial x_i}$ is the "Euler degree operator", which acts on $P^d(\mathbf{C}^n)$ as the scalar d.

b) If $A = E_{ij}$ is a matrix unit, then

$$d\nu(E_{ij}) = -x_j \frac{\partial}{\partial x_i}.$$

Example: ii) As a continuation and variant of the first example, let $GL_n \times GL_m$ act on the $n \times m$ matrices M_{nm} by the recipe

$$\nu'(g, g')(T) = (g^t)^{-1} T g'^{-1} \tag{2.5.5.4}$$

for $G \in GL_n$, $g' \in GL_m$, and $T \in M_{nm}$.

Then a calculation analogous to those given above shows that

$$d\nu'(E_{ij}, 0) = \sum_b t_{ib} \frac{\partial}{\partial t_{jb}}, \qquad \text{and} \qquad d\nu'(0, E_{k\ell}) = \sum_a t_{ak} \frac{\partial}{\partial t_{a\ell}}. \tag{2.5.5.5}$$

Note that the replacement of g by $(g^t)^{-1}$ gets rid of the slightly disorienting negative transpose in the action $d\nu$. The operators $d\nu'(0, E_{k\ell})$ were known in the 19th century literature as the *Aronhold polarization operators* [ ].

3. Algebraic groups

3.1: Definition. A subgroup of $GL_n(\mathbf{C})$ that is defined by polynomial equations is called an *algebraic group*, or a complex algebraic group if we want to emphasize the field of scalars. One can also define real algebraic groups, or groups over $\mathbf{Q}$, or over number fields, etc. Because it is defined by polynomial equations, an algebraic group is an affine algebraic variety, and so has defined on it a privileged ring of functions, the restriction of polynomial functions on $GL)n\mathbf{C})$. These are sometimes called the *regular* functions and sometimes the *rational* functions. Because both words are overused, and in particular, we want to use "regular" in the sense of the regular representation, we will use the term *rregular* for these functions.

Since $GL_n(\mathbf{C})$ is not a closed subvariety of its ambient vector space $M_n(\mathbf{C})$ of $n \times n$ complex matrices, to make clear its structure as affine algebraic variety, we should embed it in $M_{n+1}(\mathbf{C})$ as $(n+1) \times (n+1)$ matrices of the form

$$\begin{bmatrix} g & 0 \\ 0 & \det g^{-1} \end{bmatrix},$$

where g is an invertible $n \times n$ matrix. This makes clear that the algebra of rregular functions on GL_n consists of the coordinate functions on M_n together with the function $\det^{-1}$.

3.1.1: Examples. All the examples i) to vi) of matrix groups given in the previous section are in fact algebraic groups.

3.2 Regular regular (= rregular) representation of A_n and B_n.

Let's study the regular representation of A_n and B_n, on the rregular functions. Let $\mathbf{C}_{rreg}(G)$ denote the algebra of regular functions on the algebraic group G

Since A_n is defined by the conditions that $T_{jk} = 0$, $j \neq k$, for T in A_n, the rregular functions on A_n are generated by the diagonal coordinates T_{jj}, and the reciprocal of determinant $\det^{(-1)} = \left(\prod_{j=1}^n a_j\right)^{-1}$.

It is easy to see that each a_j, and hence also $\det^{-1}$, is an eigenfunction for (left or right) translation by elements of A_n. Specifically, if T is a general diagonal matrix, and $\mathbf{a}$ is a given second diagonal matrix, with diagonal entries $\mathbf{a}_{jj} = a_j$, then $(T\mathbf{a})_{jj} = T_{jj}a_j$. Thus, T_{jj} is an eigenfunction for (right) multiplication by $\mathbf{a}$, with eigencharacter $\mathbf{a}_{jj} = a_j$. Since the T_{jj} together with $\det^{-1}$, generate the ring of rregular functions $\mathbf{C}_{rreg}(A_n)$, we see that $\mathbf{C}_{rreg}(A_n)$ is spanned by eigenfunctions for the right regular representation of A_n. More precisely,$\mathbf{C}_{rreg}(A_n)$ is spanned by the monomials

$$a^{\mathbf{m}} = \prod_i a_i^{m_i}, \tag{3.2.1}$$

where

$$\mathbf{m} = \begin{bmatrix} m_1 \\ m_2 \\ m_3 \\ . \\ . \\ . \\ . \\ m_n \end{bmatrix} \tag{3.2.2}$$

is an n-tuple of integers. The monomial $a^{\mathbf{m}}$ is an eigenvector for the right regular representation of A_n on $\mathbf{C}_{rreg}(A_n)$, with eigencharacter equal to itself. When we think of $a^{\mathbf{m}}$ as a character of A_n, we will denote it by $\psi^{\mathbf{m}}$.

A similar statement holds for the left regular representation.

The right regular representation of A_n, being spanned by eigenfunctions, is completely reducible. Hence, the direct sum of any number of copies of $\mathbf{C}_{rreg}(A_n)$ is also completely reducible. By the lemma at the end of the paragraph on matrix coefficients, it follows that any rregular representation of A_n is completely reducible, and indeed, a sum of eigenvectors. In particular, any irreducible representation of A_n is one-dimensional.

Like A_n, the upper triangular group B_n is defined by the vanishing of certain of the coordinate functions T_{jk}, specifically, those for which $j > k$. Therefore, the ring $\mathbf{C}_{rreg}(B_n)$ is generated by the functions T_{jk}, together with the reciprocal of determinant, which again is given by the formula $\det^{-1} = \left(\prod_{j=1}^{n} T_{jj}\right)^{-1}$.

Consider the element D_s, a diagonal matrix with diagonal entries $(D_s)_{jj} = s^{n-j}$. Then

$$(TD_s)_{jk} = s^{n-k}T_{jk}. \tag{3.2.3}$$

We also consider the derived action of the one-parameter groups Γ_{ab}. A computation shows that

$$(T\Gamma_{ab}(s))_{jk} = T_{jk} + s\delta_{kb}T_{ja}. \tag{3.2.4}$$

On the level of the infinitesimal action, this says that

$$dR(E_{ab})(T_{jk}) = \delta_{kb}T_{ja}.$$

We note that , from the formulas above, $R(D_s)(T_{jk} = s^{n-k}T_{jk}$, while

$$R(D_s)(dR(E_{ab})T_{jk}) = R(D_s)(\delta_{kb}T_{ja}) = \delta_{kb}s^{n-a}T_{ja} = s^{b-a}(s^{n-b}\delta_{kb}T_{ja}). \tag{3.2.5}$$

From this we see that

$$R(D_s)dR(T_{ab})R(D_s)^{-1} = s^{b-a}dR(E_{ab}).$$

It follows that E_{ab} maps the s^k-eigenspace for $R(D_s)$ to the s^{k+b-a} eigenspace. Since $a < b$ for all root groups in U_n, it follows that the action of U_n is unipotent upper triangular with respect to the eigenspace decomposition defined by D_s. Since we know that A_n acts diagonally (and will preserve the eigenspaces for D_s, it follows that the regular representation of B_n is upper triangularizable. Hence, by the lemma at the end of the paragraph on matrix coefficients, all finite dimensional rregular representations of B_n are upper triangularizable. Also, all irreducible rregular representations of B_n are one-dimensional, and U_n will act trivially on them. We conclude that

$$(\hat{B}_n)_{rreg} \simeq (\hat{A}_n)_{rreg}. \qquad 3.2.6)$$

That is, each irreducible rregular representations of B_n is identified by restriction to a character of A_n.

4: Highest weight theory.

4.1:. Theorem of the Highest Weight for GL_n.

Let $B_{GL_n} = B_n = B$ denote the group of upper triangular matrices in $GL_n(\mathbf{C})$. We have

$$B = A \cdot U,$$

where $A = A_n = A_{GL_n}$ is the group of diagonal matrices, and $U = U_n = U_{GL_n}$ is the group of upper triangular matrices with ones on the diagonal (the so-called *unipotent* matrices). The Theorem of the Highest Weight says:

4.1.1: Theorem: a) In any rregular representation of GL_n, the group B_n acts in an upper triangular fashion. More precisely, the group A_n acts in a diagonalizable fashion, and the group U_n acts in a unipotent triangularizable fashion. In particular, if V is the space of a representation of GL_n, then the subspace V^{U_n} of U_n-invariant vectors is always non-zero. It is invariant under A_n, which acts on it in a diagonalizable fashion. Thus, it is a direct sum of eigenspaces for B_n.

b) In any rregular irreducible representation ρ of GL_n on a vector space V, the subspace of V^{U_n} of U_n invariant vectors is one dimensional.

c) In the context of b), the A_n eigencharacter ψ_ρ defined by V^{U_n} determines the isomorphism class of ρ.

Terminology: Any B_n eigenvector in a representation of GL_n is called a *highest weight vector*. The corresponding eigencharacter of B_n is called a *highest weight*.

4.1.2: Dominant weights for GL_n; diagram notation.

Not all characters of A_n can be highest weights of (irreducible, finite dimensional rregular) representations of GL_n. A character of A_n that is a highest weight is called *dominant*. We will determine which characters of A_n are dominant. This entails studying:

4.1.3: Representations of sl_2.

The key to understanding what weights can be dominant is the representation theory of SL_2, which is most economically done by considering representations of the Lie algebra sl_2. This is three dimensional, and has a basis

$$e^+ = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}, \qquad e^- = \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix}, \qquad h = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}. \tag{4.1.3.1}$$

These elements satisfy the commutation relations

$$[h, e^+] = 2e^+, \qquad [h, e^-] = -2e^-, \qquad [e^=, e^-] - h.$$

A representation of sl_2 is a collection of 3 operators that satisfy these commutation relations. Typically, we will denote these operators by e^+, e^- and h, and not have any special notation for the representation.

When dealing with sl_2, it is useful to consider the *Casimir operator*

$$C = h^2 + 2(e^+e^- + e^-e^+) = h^2 + 2h + 4e^-e^+ = h^2 - 2h + 4e^+e^-.$$

It is easy to check that C commutes with $e^\pm$ and h. Therefore, any eigenspace for C will be invariant under sl_2.

Consider a representation of sl_2 on a finite dimensional vector space V. The commutation relation $[h, e^+] = 2e^+$ can be rewritten $he^+ = e^+h + 2e^+ = e^+(h+2)$. In other words, we can move h from the left to the right of e^+ by replacing it with $h+2$. In particular, if $\vec{v}$ is an eigenvector for h, with eigenvalue λ, then $e^+(\vec{v})$ is again an eigenvector for h, with eigenvalue $\lambda + 2$. Thus, if we start with an eigenvector $\vec{v}$ for h, and apply e^+ successively, we will get eigenvectors for h with eigenvalues λ, $lambda + 2$, $\lambda + 4$, $\lambda + 6$, etc. Since these eigenvalues are distinct, the eigenvectors, providing they are non-zero, must be linearly independent. Since V has finite dimension, eventually, one of the $(e^+)^k(\vec{v})$ will have to vanish. This means that we can always find a vector $\vec{v}_o$ that is an eigenvector for h and is annihilated by e^+. Renaming the eigenvalue if necessary, we will have:

$$h(\vec{v}_o) = \lambda\vec{v}_o, \qquad e^+(\vec{v}_o).$$

We can then compute that $\vec{v}_o$ is also an eigenvector for the Casimir operator:

$$C(\vec{v}_o) = (h^2 + 2h + 4e^-e^+)(\vec{v}_o) = (\lambda^2 + 2\lambda)\vec{v}_o = ((\lambda+1)^1 - 1)\vec{v}_o.$$

Starting with $\vec{v}_o$, apply e^- successively. Define $\vec{v}_k = (e^-)^k(\vec{v}_o) = e^-(\vec{v}_{k-1})$. Just as applying e^+ raises eigenvalues of h by 2, applying e^- lowers them by 2. Hence $h(\vec{v}_k) = (\lambda - 2k)\vec{v}_k$. Furthermore, we can compute that

$$e^+(\vec{v}_k) = e^+e^-(\vec{v}_{k-1}) = (\frac{1}{4})(C - ((h-1)^2 - 1)(\vec{v}_{k-1}) = (\frac{1}{4})((\lambda+1)^2 - 1 - ((\lambda - 2(k-1) - 1)^2 - 1))\vec{v}_{k-1}$$

$$= (\frac{1}{4})(\lambda + 1 + \lambda - 2k + 1)(\lambda + 1 - \lambda + 2k - 1)\vec{v}_{k-1} = (\frac{1}{4})(2(\lambda + 1 - k)(2k)\vec{v}_{k-1} = k(\lambda + 1 - k)\vec{v}_{k-1}.$$

This computation shows first, that the span of the $\vec{v}_k$ is invariant under all three operators e^+, e^- and h. That is, it forms a representation of sl_2. It is easy to check that this representation is irreducible.

Moreover, these formulas also tell us the precise structure of the representation generated by $\vec{v}_o$, and all the possibilities for irreducible representations of sl_2. Just as we concluded that successive applications of e^+ must eventually annihilate any vector, we must have $\vec{v}_k) = 0$ for sufficiently large k. If ℓ is the largest number such that $\vec{v}_\ell \neq 0$, then we will have $e^-(\vec{v}_\ell) = 0$, hence $0 = e^+e^-(\vec{v}_\ell) = 2(\ell+1)(\lambda-\ell)\vec{v}_\ell$. This can only happen if $\lambda = \ell$. That, is, λ should be a non-negative integer, and one less than the dimension of the span of the $\vec{v}_k$. The following statement summarizes this discussion.

4.1.3.2: Proposition. a) Up to isomorphism, there is exactly one irreducible representation V_ℓ of sl_2 of dimension ℓ for each positive integer ℓ.

b) The space V_ℓ has a basis $\vec{v}_k$ for $0 \leq k < \ell$, such that

$$h(\vec{v}_k) = (\ell - 1 - 2k)\vec{v}_k, \qquad e^-(\vec{v}_k) = \vec{v}_{k_1}, \qquad e^+(\vec{v}_k) = (k+1)(\ell-k)\vec{v}_{k-1}.$$

In particular, the highest weight vector in $V\ell$ has weight (= h eigenvalue) $\ell - 1$. Note that this is always non-negative.

c) The Casimir operator acts on V_ℓ by the eigenvalue $\ell^2 - 1$.

Remark: The representation V_ℓ can be realized by means of the natural action of SL_2 (restricted from GL_2 on the polynomials of degree $\ell - 1$ on $\mathbf{C}^2$. Thus

$$P(\mathbf{C}^2) \simeq \oplus_{d=0}^{\infty} P^d(\mathbf{C}^2) \simeq \oplus_{d=0}^{\infty} V_{d+1}.$$

4.1.4: Criterion for Dominance.

Now suppose that we have a representation ρ of GL_2, and let $\vec{u}$ be an eigenvector for the torus A_2, so that

$$\rho(\begin{bmatrix} a_1 & 0 \\ 0 & a_2 \end{bmatrix})(\vec{u}) = a_1^{m_1} a_2^{m_2} \vec{u}.$$

A computation of the associated infinitesimal representation of sl_2 shows that $d\rho(h)(\vec{u}) = (a_1 - a_2)\vec{u}$. If this is a highest weight vector for GL_2, then the Proposition 4.1.3.2 implies that $a_1 - a_2 \geq 0$, or $a_1 \geq a_2$.

We note that GL_n contains many copies of SL_2. In fact, each root group Γ_{ab} and its opposite Γ_{ba} generates a copy of SL_2 If we apply the highest weight criterion to each of these groups, we obtain the necessity of the following criterion.

Proposition: In order for a weight $\psi_{\mathbf{m}}$ of A_n, where $\mathbf{m} = \begin{bmatrix} m_1 \\ m_2 \\ m_3 \\ . \\ . \\ . \\ m_n \end{bmatrix}$ is an n-tuple of integers to be a highest weight of a representation of GL_n, it is necessary and sufficient that the m_j be decreasing: $m_{j+1} \leq m_j$.

In order to show the sufficiency of the dominance criterion, we need to produce representations of GL_n with highest weights of the stated form. One way to do this is to consider polynomial functions on matrices. Let $x_{ij} : 1 \leq i \leq n; 1 \leq j \leq m$ be the entries of a generic $n \times m$ matrix X. Consider the polynomials

$$\delta_{I,J} = \det \begin{bmatrix} x_{i_1j_1} & x_{i_1j_2} & x_{i_1j_3} & . & . & . & x_{i_1j_k} \\ x_{i_2j_1} & x_{i_2j_2} & x_{i_2j_3} & . & . & . & x_{i_2j_k} \\ x_{i_3j_1} & x_{i_3j_2} & x_{i_3j_3} & . & . & . & x_{i_3j_k} \\ . & . & . & . & . & . & . \\ . & . & . & . & . & . & . \\ . & . & . & . & . & . & . \\ x_{i_kj_1} & x_{i_kj_2} & x_{i_kj_3} & . & . & . & x_{i_kj_k} \end{bmatrix}, \tag{4.1.5a}$$

of $k \times k$ submatrices of the generic $n \times m$ matrix

$$T = \begin{bmatrix} x_{11} & x_{12} & x_{13} & . & . & . & x_{1m} \\ x_{21} & x_{22} & x_{23} & . & . & . & x_{2m} \\ x_{31} & x_{32} & x_{33} & . & . & . & x_{3m} \\ . & . & . & . & . & . & . \\ . & . & . & . & . & . & . \\ . & . & . & . & . & . & . \\ x_{n1} & x_{n2} & x_{n3} & . & . & . & x_{mn} \end{bmatrix}, \tag{4.1.5b}$$

defined by taking rows defined by a set $I = \{i_1 < i_2 < i_3 < \ . \ . \ . \ < i_k\}$ and $J = \{j_1 < j_2 < j_3 < \ . \ . \ . < j_k\}$ of indices between 1 and n (for I) or between 1 and m (for J).

Let GL_n act on the polynomials $P(M_{nm})$ on M_{nm} by multiplication on the right, as in formula (2.5.5.4) From the formulas (2.5.5.5), we can check that the infinitesimal action of a root group Γ_{ab} on one of the $\delta_{I,J}$ has the effect of a row operation, and can be expressed in terms of the index set I. If b is an index appearing in I, then the infinitesimal action of Γ_{ab} on $\delta_{I,J}$ is to replace b with a. If b does not belong to I, then Γ_{ab} leaves $\delta_{I,J}$ fixed. From this we see that if $I = K_o = \{1, 2, 3, 4, \ . \ . \ . k\}$ consists of the consecutive numbers from 1 to k, then $\delta_{K_o,J}$ is a highest weight vector for GL_n, with weight $\psi_{\mathbf{m}_k}$, where $\mathbf{m}_k$ is the n-tuple whose first k entries are 1, and the rest are 0.

Since the action ν' of GL_n on $P(M_{nm})$ is by algebra automorphisms, the product of highest weight vectors will be again a highest weight vector, with eigencharacter equal to the product of the eigencharacters of the factors. Hence, we see that by multipying the δ_{K_oJ} together, we can produce a highest weight vector that has any dominant character with non-negative exponents as its eigencharacter. Then multiplying by an appropriat power of $\det^{-1}$, we can achieve any dominant character as a highest weight.

4.1.5: Diagram notation. Let $\psi_{\mathbf{m}}$ be a character of A_n. Then $\psi_{\mathbf{m}}$ will extend to a polynomial on the space of all diagonal matrices if and only if all the entries of the n-tuple $\mathbf{m}$ are non-negative. Let us call such a character *polynomial*.

Suppose that $\psi_{\mathbf{m}}$ is both dominant and polynomial. Then the entries m_i of $\mathbf{m}$ are non-negative integers that (weakly) decrease as i increases. It is frequent practice to associate to such a character a it Young diagram (aka Ferrers diagram), which is an array of square boxes, all of the same size, in left justified rows, with m_1 boxes in the top row. m_2 boxes in the second row down, and so forth. The entries of $\mathbf{m}$ that are zero are ignored in this construction. Thus, the tuple $\begin{bmatrix} 5 \\ 3 \\ 3 \\ 2 \\ 0 \\ 0 \end{bmatrix}$ produces the diagram

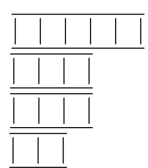

A first advantage of this notation is that is allows us to label a representation of GL_n by a diagram, for any n that is as large as the number of rows in the diagram. Thus, the same diagram can label representations of GL_n for different n. This will be quite useful in stating (GL_n, GL_m)-duality below. It also facilitates many combinatorial constructions that are useful for describing representations of GL_n.

4.2: The flag algebra.

Let $\rho : GL_n \to GL(V)$ be an irreducible representation of GL_n, and let ρ^* be the contragredient representation on V^*. Let λ_o be the B_n eigenvector in V^*. Consider the matrix coefficients $\phi_{\lambda_o,\vec{v}}$. For $g \in GL)n$ and $b \in B_n$, we can compute that

$$L_b(\phi_{\lambda_o,\vec{v}})(g) = \phi_{\lambda_o,\vec{v}}(b^{-1}g) = \lambda_o(\rho(b^{-1}g)(\vec{v})) = \lambda_o(\rho(b^{-1})\rho(g)\vec{v}) = \rho^*(b)(\lambda_o)(\rho(g)\vec{v}) \qquad (4.2.1)$$
$$= \psi_{\rho^*}(b)\lambda_o(\rho(g)\vec{v}) = \psi_{\rho^*}(b)\phi_{\lambda_o,\vec{v}}(g).$$

In particular, this shows that $\phi_{\lambda_o,\vec{v}}$ is invariant under translation on the left by elements of U_n. In other words, $\phi_{\lambda_o,\vec{v}}$ factors to define a function on the coset space U_n/GL_n. Conversely, suppose that we have a copy of ρ realized as functions on the coset space U_n/GL_n. Then evaluating at the identity coset gives us a linear functional on ρ that is invariant under U_n. By the uniqueness of λ_o, we conclude

4.2.2: Proposition. a) For any irreducible regular $\rho \in (\hat{GL}_n)_{rreg}$, there is a unique (up to multiples) embedding $\Phi_\rho : \rho \to \mathbf{C}_{rreg}(U_n/GL_n)$. Thus,

$$\mathbf{C}_{rreg}(U_n/GL_n) \simeq \oplus_\rho \in \hat{G}_{rreg} R_\rho$$

is a sum of one copy each of every regular irreducible representation of GL_n.

b) Moreover, under left translations by A_n, the space V_ρ consists of eigenvectors with eigencharacter ψ_{ρ^*}: $R_\rho = L_{\psi_{\rho^*}}$.

c) Since the actions of GL_n (on the right) and of A_n (on the left) are actions be automorphisms, we have $L_\psi \cdot L_{\psi'} = L_{\psi\psi'}$. Thus, the decomposition of $\mathbf{C}_{rreg}(U_n/GL_n)$ of part a) exhibits $\mathbf{C}_{rreg}(U_n/GL_n)$ as an $\hat{A}_n^+$-graded algebra that provides a home for each irreducible representation of GL_n in a unique way.

4.3: Multiplicity-free actions. Let G act on a vector space V. Consider the induced action on the polynomial functions: if the representation of G on V is denoted by ρ, then the action on functions is denoted $\hat{\rho}$, and is specified by

$$\hat{\rho}(g)(P(\vec{v})) = P(\rho(g)^{-1}(\vec{v})), \tag{4.3.1}$$

for $g \in G$, and P a polynomial on V. This is an action on the algebra $P(V)$ of polynomial functions on V by algebra automorphisms. In particular

$$\hat{\rho}(g)(P \cdot Q) = (\hat{\rho}(g)(P)) \cdot (\hat{\rho}(g)(Q)). \tag{4.3.2}$$

Here $P \cdot Q$ is indicating product of P and Q in $P(V)$.

Although the algebra $P(V)$ is infinite-dimensional, it is a direct sum of its homogeneous components

$$P(V) \simeq \oplus_{d \geq 0} P^d(V),$$

where $P^d(V)$ is the space of polynomials such that $p(t\vec{v}) = t^d p(\vec{v})$ for all points $\vec{v} \in V$. The spaces $P^d(V)$ are finite-dimensional, so that $P(V)$ is a sum of finite-dimensional representations.

We say that the action of GL_n on V (or on $P(V)$) is *multiplicity-free* if the multiplicity of any irreducible representation of GL_n in $P(V)$ is at most 1.

4.3.3: Proposition. a) A rregular action ρ of GL_n on a vector space V is multiplicity-free if the subgroup B_n has a dense (and open) orbit.

b) Suppose that $\vec{x}_o$ is a point such that $B_n(\vec{x}_o)$ is dense in V. Let $H \subset B_n$ be the stabilizer of $\vec{x}_o$ in B_n. Let p be a polynomial on V that is a B_n eigenvector, and let ψ_p be the associated eigencharacter of B_n. Then ψ_p is trivial on H.

Proof: a) By the theorem of the highest weight, to show that ρ is multiplicity-free, it suffices to show that there is at most one B_n eigenvector for any character ψ of B_n.

Let x_o be a point such that the B_n orbit of x_o is dense in V. Let p be a polynomial function on V that is a B_n eigenvector, with eigencharacter ψ. Then for $b \in B_n$, we have

$$p(\rho(b)(x_o)) = \rho * (b^{-1})(p)(x_o) = \psi(b)^{-1} p(x_o).$$

This formula shows that p is determined up to scalar multiples on the B_n orbit $B_n(x_o)$ by the character ψ. Since the B_n orbit of x_o is dense, p is determined everywhere by its restriction to $B_n(x_o)$. Hence, p is determined everywhere by ψ. This shows that there is at most one ψ eigenvector for every character of B_n, so statement a) follows. Note that the above formula also shows that, if $p(x_o) = 0$, then p must be identically zero.

b) Again supposing p to be the eigenvector with eigencharacter ψ, let h be an element of the stabilizer in B_n of x_o. Then we can compute that

$$\psi(h)p(x_o) = \rho^*(h)(p)(x_o) = p(\rho(h^{-1}(x_o)) = p(x_o),$$

since $\rho(h)(x_o) = x_o$. Since $p(x_o) \neq 0$, by the remark at the end of part a), we conclude that $\psi(h) = 1$, which is statement b).

4.3.4: Theorem: a) Suppose that the action of GL_n on V is multiplicity-free. Consider the subalgebra $P(V)^{U_n}$ of U_n-invariants. Let p be a B_n eigenvector in $P(V)^{U_n}$. Let

$$p = \prod_{j=1}^{k} q_j$$

be the factorization of p into prime polynomials in $P(V)$. Then each q_j is in fact a B_n eigenvector.

b) The collection of eigencharacters ψ for B_n-eigenvectors in $P(V)^{U_n}$ is a free semigroup generated by the eigencharactors of the prime B_n eigenvectors.

c) The ring $P(V)^{U_n}$ is a polynomial ring on the prime B_n eigenfunctions as generators.

Proof: a) Consider a factorization as in statement a). Take $b \in B_n$ and apply it to both sides of the equation. Then we have

$$\psi(b)(p) = \rho^*(b)(p) = \rho^*(b)(\prod_{j=1}^{k} q_j) = \prod_{j=1}^{k} \rho^*(b)(q_j),$$

or $p = \prod_{j=1}^{k} \psi(b)^{-1}\rho^*(b)(q_j)$ is another factorization of p into prime factors. By the uniqueness of prime factorization, it follows that, up to order and scalar multiples, the $\rho^*(b)(q_j)$ are same as the q_j. In other words, $\rho^*(B)$ permutes the q_j up to multiples. Hence, if we raise b to a suitable power b^s, we will have $\rho^*(b^s)(q_j) = \alpha_j(b^s)(q_j)$ for all j. Since there are a finite number, say e of the q_j, we could take $s = e!$. In particular, we can take s independent of b and of q_j. However, since B_n is a connected group, every element of B_n has the form b^s for some other $b \in B_n$. Thus, the $q_j$are eigenvectors for all b in b_n, as desired.

b) Let $\{q_a : 1 \leq a \leq d\}$ be prime B_n eigenfunctions, with associated eigencharacters ψ_a. We claim that the ψ_a generate a free semigroup. Suppose not. Suppose that there is some relation $\prod_a \psi_a^{r_a} = \prod_a \psi^{s_a} = \phi$ for suitable non-negative integers r_a and s_a. Then $p_1 = \prod_a q_a^{r_a}$ and $p_2 = \prod q_a^{s_a}$ will both be eigenfunctions for B_n with eigencharacter ϕ. Since ρ is multiplicity-free, we have $p_2 = cp_1$ for some scalar c. By uniqueness of prime factorization, this implies that $s_a = r_a$, so that the proposed relation is trivial. This proves statement b).

c) Now take the full collection $\{q_a\}$ of prime B_n eigenfunctions. If we form a monomial $\prod_a \psi_a^{r_a}$ in the q_a, we know that it will be a B_n eigenfunction, with eigencharacter $\prod_a \psi_a^{r_a}$, From b) we know that all these eigencharacters are distinct. If follows that the monomials in the q_a are all linearly independent. This in turn implies that the algebra generated by the q_a is a polynomial ring. Also, this ring includes all B_n eigenvectors, by part b). This finishes the proof of c).

4.4: (GL_n, GL_m) **duality**: Consider $GL_n \times GL_m$ acting on the space M_{nm} of $n \times m$ matrices, by the recipe:

$$(g, g')(T) = (g_1^t)^{-1} T (g')^{-1}.$$

(The purpose of the use of $(g^t)^{-1}$ rather than simply g in the action on the left is to make the representations that occur be polynomial representations, in the sense that the action of A_n acts by polynomial characters.)

4.4.1: Theorem: a) The action of $GL_n \times GL_m$ on M_{nm} defined above is multiplicity-free.

b) More precisely, we have the decomposition

$$P(M_{nm} \simeq \sum_D \rho_n^D \otimes \rho_m^D, \qquad (4.4.2)$$

where D runs over all diagrams with at most $\min(n, m)$ rows.

Remark: the explicit decomposition (4.4.2) is much stronger than multiplicity-freeness for the joint action: it says that the ρ_n^D isotypic component of $P(M_{nm})$ under the action of GL_n is the same as the ρ_m^D isotypic component of $P(M_{nm})$ under the GL_m action. In other words, the isotypic component for one group determines the isotypic component of the other group. For this reason, we call this decomposition $(GL_n.GL_m)$ *duality.*

Proof: a) First, we should note that we are here working with the action of a product $GL_n \times GL_m$ rather than just GL_n. Because irreducible representations of product groups are tensor products of representations of the factors, the theorem of the highest weight is valid also for the product $GL_n \times GL_m$, and the set of dominant characters consists of products of dominant characters for GL_n and GL_m. Also, to show that the action of $GL_n \times GL_m$ is multiplicity-free, it suffices to show that the product $B_n \times B_m$ has a dense orbit.

The action of B_n on M_{nm} is by top-to-bottom row operations, and the action of B_m is by left-to-right column operations. Let J_{nm} be the $n \times m$matrix with ones on the main diagonal and zeroes elsewhere. The theory of Gaussian elimination tells us that the orbit of J_{nm} is dense in M_{nm}. Therefore, Proposition 4.3.3 guarantees us that the $GL_n \times GL_m$ action on M_{nm} is multiplicity-free.

b) By the theorem of the highest weight, determining the $GL_n \times GL_m$ decomposition of $P(M_{nm})$ is equivalent to knowing the algebra $P(M_{nm})^{U_n \times U_m}$ of $GL_n \times GL_m$ highest weight vectors. Since the action of $GL_n \times GL_m$ on M_{nm} is multiplicity-free, part b) of Proposition ??? tells us that the ring $P(M_{nm})^{U_n \times U_m}$ of $GL_n \times GL_n$ highest weight vectors is a polynomial ring. Determining the $GL_n \times GL_m$ decomposition of $P(M_{nm})$ is equivalent to knowing the prime joint $GL_n \times GL_m$ highest weight vectors.

Above, in determining the set of dominant characters of the torus A_n of GL_n, we displayed some highest weight vectors for GL_n in $P(M_{nm})$. These were the determinants $\delta K_o j$, where K_o is the set of integers from 1 to k, and J is any k element subset of the first m integers.

In that discussion, we considered the more general class of determinants δ_{IJ}. We saw there that the infinitesimal action of the root groups of GL_n on the δ_{IJ} has the effect of row operations. A very similar calculation shows that the infinitesimal action of the root groups of GL_m on the δ_{IJ} has the effect of column operations. We can conclude that the determinant

$$\delta_k = \delta_{K_o K_o} \tag{4.4.4}$$

is a simultaneous $GL_n \times GL_m$ highest weight vector. We claim that the δ_k for $1 \leq k \leq \min(n, m)$ are the prime generators for $P(M_{nm})^{U_n \times U_m}$.

It is well-known that in fact, determinants are prime polynomials in their entries. However, we do not need to invoke that result here. It suffices to consider the highest weight of δ_k, which is $\psi_{\mathbf{m}_k} \times \psi'_{\mathbf{m'}_k}$, where $\psi_{\mathbf{m}_k}$ is the A_n character parametrized by the n-tuple $\mathbf{m}_k$, and $\psi'_{\mathbf{m}_k}$ is the analogous character for $\hat{A}_m$. The characters $\psi_{\mathbf{m}_k}$ and $\psi'_{\mathbf{m'}_k}$ are primitive elements of the cone of polynomial dominant weights in $\hat{A}_n$, or $\hat{A}_m$ respectively. Hence, the only way to factor $\psi_{\mathbf{m}_k} \times \psi'_{\mathbf{m'}_k}$ in the polynomial dominant cone in $\hat{A}_n \times \hat{A}_m$ is as $(\psi_{\mathbf{m}_k} \times 1) \cdot (1 \times \psi'_{\mathbf{m'}_k}$, where 1 here denotes the trivial character of the relevant torus. However, since the centers of GL_n and GL_m coincide in the scalar operators on M_{nm}, we see that any joint eigencharacter must have the same degree in both tori. This precludes the factorization involving a product of a trivial character with a non-trivial character. Thus, the characters of the δ_k must be primitive in the cone of possible characters; correspondingly, the δ_k are prime joint highest weight vectors.

To show that the δ_k exhaust all prime joint highest weight vectors. we can simply count how many there should be. For simplicity, assume that $m \leq n$. In the reverse situation, the argument is similar. We note that the stabilizer of J_{nm} in $A_n \times A_m$ includes the matrices

$$\begin{bmatrix} a_1 & 0 & 0 & . & . & . & . & . & . & 0 \\ 0 & a_2 & 0 & . & . & . & . & . & . & 0 \\ 0 & 0 & a_3 & . & . & . & . & . & . & 0 \\ . & . & . & . & . & . & . & . & . & . \\ . & . & . & . & . & . & . & . & . & . \\ . & . & . & . & . & . & . & . & . & . \\ . & . & . & . & . & . & . & . & . & . \\ . & . & . & . & . & . & . & . & . & . \\ 0 & 0 & 0 & . & . & . & . & . & . & a_n \end{bmatrix} \times \begin{bmatrix} a_1^{-1} & 0 & 0 & . & . & . & 0 \\ 0 & a_2^{-1} & 0 & . & . & . & 0 \\ 0 & 0 & a_3^{-1} & . & . & . & 0 \\ . & . & . & . & . & . & . \\ . & . & . & . & . & . & . \\ . & . & . & . & . & . & . \\ 0 & 0 & 0 & . & . & . & a_m^{-1} \end{bmatrix}. \tag{4.4.1.1}$$

This is a subgroup of $A_n \times A_m$ of codimension m which equals $\min(n, m)$ under our assumption. Hence, there can be at most m prime generators for $P(M_{nm})^{U_n \times U_m}$. Since the δ_k already give us that number, they must exhaust the prime generators.

Finally, we note that if we use diagram notation, then δ_k is the highest weight of the representation $\rho_n^{C_k} \otimes \rho_m^{C_k}$, where C_k is the diagram consisting of a column of length k. In diagram notation, taking products of highest weight vectors corresponds to justaposing diagrams row by row. It follows that the diagrams attached to any joint $GL_n \times GL_m$ highest weight vector will be the same. Furthermore, since we have column diagrams of any length up to $\min(n,m)$, we can construct all diagrams with at most $\min(n,m)$ rows. This proves part b) of the theorem.

5. Standard Monomial Theory.

5.1: Hodge's Theory.

From (GL_n, GL_m)-duality (Theorem 4.4.1)

$$P(M_{nm}) \simeq \oplus_D \ \rho_n^D \otimes \rho_m^D,$$

and the theorem of the highest weight, we find that the ring of GL_m highest weight vectors

$$P(M_{nm})^{U_m} \simeq (\oplus_D \ \rho_n^D \otimes \rho_m^D)^{U_m} \simeq \oplus_D \ \rho_n^D \otimes (\rho_m^D)^{U_m}, \tag{5.1.1}$$

is a sum of one copy of each irreducible representation of GL_n corresponding to a Young diagram of at most $\min(n,m)$ rows. Also, $P(M_{nm})^{U_m}$ is multi-graded by the highest weights of GL_m. In fact this ring is isomorphic to the subring of the flag algebra of GL_n defined by the same collection of representations.

We want to describe this ring. We will give a basis for this ring, that will be compatible with the multigraded structure, and which therefore will provide a basis for each ρ_n^D. This basis was originally constructed by Hodge [ ].

When D is a column of length k, the space $\rho_n^D \otimes \rho_m^D$ in $P(M_{nm})$ is spanned by the determinants

$$\delta_{I,J} = \det \begin{bmatrix} x_{i_1j_1} & x_{i_1j_2} & x_{i_1j_3} & . & . & . & x_{i_1j_k} \\ x_{i_2j_1} & x_{i_2j_2} & x_{i_2j_3} & . & . & . & x_{i_2j_k} \\ x_{i_3j_1} & x_{i_3j_2} & x_{i_3j_3} & . & . & . & x_{i_3j_k} \\ . & . & . & . & . & . & . \\ . & . & . & . & . & . & . \\ . & . & . & . & . & . & . \\ x_{i_kj_1} & x_{i_kj_2} & x_{i_kj_3} & . & . & . & x_{i_kj_k} \end{bmatrix}, \tag{5.1.2}$$

of $k \times k$ submatrices of the generic $n \times m$ matrix

$$T = \begin{bmatrix} x_{11} & x_{12} & x_{13} & . & . & . & x_{1m} \\ x_{21} & x_{22} & x_{23} & . & . & . & x_{2m} \\ x_{31} & x_{32} & x_{33} & . & . & . & x_{3m} \\ . & . & . & . & . & . & . \\ . & . & . & . & . & . & . \\ . & . & . & . & . & . & . \\ x_{n1} & x_{n2} & x_{n3} & . & . & . & x_{mn} \end{bmatrix},$$

defined by taking rows defined by a set $I = \{i_1 < 1_2 < 1_3 < \ . \ . \ . \ < i_k\}$ and $J = \{j_1 < j_2 < j_3 < \ . \ . \ . < j_k\}$ of indices between 1 and n (for I) or between 1 and m (for J).

Among these, the GL_m-highest weight vectors are the

$$\delta_I = \delta_{IK_o}, \tag{5.1.3}$$

where $K_o = \{1, 2, 3, \ . \ . \ . \ , k\}$ consists of just the first k positive integers.

The determinants δ_I will generate the algebra $P(M_{nm})^{U_m}$. This is easy to see since the span of the δ_I is GL_n invariant, and the monomials in the $\delta_k = \delta_{K_o}$ are the GL_n highest weight vectors of the spaces $\rho_n^D \otimes (\rho_m^D)^{U_m}$. Since the δ_I are geneators for the algebra, any element of the algebra is expressible as a sum of monomials in the δ_I – that is, the monomials in the δ_I span $P(M_{nm})^{U_m}$. However, they will do so in a redundant manner, and many linear combinations of monomials in the δ_I will be zero. We would like to eliminate this redundancy by selecting from the monomials in the δ_I a basis for $P(M_{nm})^{U_m}$.

5.1.4: Tableau order.

In order to do this, we put a partial ordering on the generators δ_I, as follows. We form the sets I into *column tableaux*

$$T_I = \begin{bmatrix} i_1 \\ i_2 \\ i_3 \\ . \\ . \\ . \\ i_k \end{bmatrix}, \tag{5.1.4}$$

which are just column vectors with the elements of I listed in increasing order from top to bottom. We then define the *tableau order* on the I or T_I by declaring: $I \leq J$, or $T_I \leq T_J$ if and only if

i) $^{\#}(I) \geq^{\#} (J)$; and

ii) $i_a \leq j_a$ for $1 \leq a \leq^{\#} (J)$.

That is, the length (or depth) of column T_I is at least as large as the length/depth of T_J, and each entry of T_I is less than or equal to the corresponding entry of T_J.

We now declare a monomial $\prod_{a=1}^{r} \delta_{I_a}$ in the generators δ_I to be *standard* if all members of the collection T_{I_a} are comparable; that is, if the collection of column tableaux T_{I_a} form a totally ordered set or *chain* with respect to the tableau order.

If the T_{I_a} do form a chain, there is no harm in supposing that the T_{I_a} are listed in increasing order: $T_{I_a} \leq T_{I_{a+1}}$. Henceforward we do make this convention.

With this terminology, we can now state a first version of standard monomial theory.

5.1.5: SMT Theorem 1. (Hodge) The standard monomials in the δ_I form a basis for $P(M_{nm})^{U_m}$.

Remarks: i) The condition for standardness of a monomial refers only to the generators that appear to positive powers, and does not depend on how large these powers are. If we have a chain $\{T_{I_a} : 1 \leq a \leq r\}$ of column tableaux, then any monomial in the corresponding generators δ_{I_a} is standard. Such monomials are in particular linearly independent. This means that the whole polynomial ring generated by the δ_{I_a} injects into the flag algebra. Furthermore, since any standard monomial will belong to one of these polynomial rings (namely, the one generated by the factors of that monomial), the flag algebra is spanned by the polynomial rings generated by chains of generators. Since every chain of tableaux is clearly contained in a largest possible such chain (a *maximal chain*), it is enough to look at the polynomial rings on the maximal chains. Thus, a reformulation of the above theorem can be rephrased as follows:

5.1.6: SMT Theorem 2. a) Let C denote a maximal chain of column tableaux. Then the corresponding set of generators $\{\delta_I : T_I \in C\}$ generates a polynomial subring R_C of $P(M_{nm})^{U_m}$.

b) The poynomial rings R_C span $P(M_{nm})^{U_m}$.

c) All linear dependences among elements of the subrings R_C of $P(M_{nm})^{U_m}$ results from the intersections of two maximal chains C and C'.

We refer to the situation described in SMT Theorem 2 by saying that $P(M_{nm})^{U_m}$ is an *almost direct sum* of the polynomial subrings R_C.

5.1.7: Semistandard tableux.

ii) A convenient combinatorial scheme for parametrizing standard monomials is provided by the construction of tableaux. We have attached to each generator δ_I the column tableau T_I, which is just the column vector with entries equal to the elements of I, listed in increasing order down the column. We can imagine these entries occupying boxes arranged in a vertical stack of height $^{\#}(I)$. Then, to describe a product of δ_I, we can just juxtapose the stacks corresponding to the factors in order from left to right. If we are dealing with a standard monomial, we list the columns in increasing order from left to right. This will result in a Young diagram whose boxes are filled with numbers from 1 to n. Such a filling of the boxes of a diagram with numbers is called a *tableau*. The tableaux that result from standard monomials will have the following properties:

i) The numbers in each column are strictly increasing from top to bottom.

ii) The numbers in each row are weakly increasing from left to right.

Conversely, it is easy to convince oneself that these two conditions characterize the tableaux arising form standard monomials. The condition of strict increase on the entries in a column just say that this column is indeed of the form T_I for the set I consisting of the column entries. The condition that the columns form a diagram, meaning that their lengths decrease from left to right, is one part of the condition that the columns increase in the tableau order from left to right. The other part of the tableau order condition then just says that the entries in any row should increase from left to right, which is the second part of the semistandardness criterion.

In summary, then we see that there is a natural bijection between standard monomials and semistandard tableaux. A pleasant aspect of this bijection is that it is consistent with the multigrading on $P(M_{nm})^{U_m}$ induced by the decomposition into joint (GL_n, A_m) modules. More precisely, a standard monomial μ will belong to the summand $\rho_n^D \otimes (\rho_m^D)^{U_m}$ exactly when the diagram of the tableau T_μ corresponding to μ is D. We leave it to the reader to check this consistency.

5.2: Proof of Standard Monomial Theory I - independence.

We will establish the SMT Theorems in two parts. In the first part, we will show that the standard monomials are linearly independent. In the second part, we will show that they span $P(M_{nm})^{U_m}$.

To show that standard monomials are linearly independent, we use a term order on the monomials in the usual coördinates x_{ij} on M_{nm}. Arguments like this were first given by Sturmfels [ ], [ ].

Consider the polynomial ring in some variables z_a. The notion of a term order on monomials in the Z_a comes from the theory of Gröbner bases [ ]. A *term order*, or *monomial ordering* is a well-ordering (a total ordering in which any subset has a least element) $\geq_z$ that is compatible with multiplication in the following sense.

For monomials μ, μ' and ν, if $\mu \geq_z \mu'$, then also $\mu\nu \geq_z \mu'\nu$.

Suppose that we have a term order $\geq_z$ on monomials in the z_a. Consider any polynomial

$$p(z_1, z_2, \ . \ . \ . \ , z_\ell) = \sum_\mu c_\mu \mu,$$

where the μ are monomials in the z_a, and c_μ is the coefficient of μ in p. For a given polynomial p coefficients c_μ will be non-zero for only finitely many μ. If $c_\mu \neq 0$, we say that μ *occurs* in p. We then define the *leading term* $\ell t(p)$ of p to be the largest monomial that occurs in p:

$$\ell t(p) = \max\{\mu : \mu \text{ occurs in } p\}. \tag{5.2.1}$$

Important properties of leading terms that we will use are formulated in the following lemmas.

5.2.2: Lemma. The leading terms are multiplicative: If p and q are polynomials in the z_a, then

$$\ell t(pq) = \ell t(p)\ \ell t(q).$$

Proof: Left to the reader.

5.2.3: Lemma. Distinct leading terms implies linear independence: If $\{p_a\}$ are a collection of polynomials, and $\ell t(p_a) \neq \ell t(p_b)$ for distinct indices a and b, then the collection $\{p_a\}$ is linearly independent.

Proof: Since a monomial ordering is a total ordering, the largest monomial among the leading terms of all the p_a will be the largest monomial among all the monomials appearing in all the p_a. By assumption, this monomial will appear in only one of the p_a, say p_{a_o}. If $\sum_a \gamma_a p_a = 0$ is a purported linear dependence among the p_a, then if we expand each p_a in monomials, the corresponding sum of coefficients must cancel out for each monomial μ. Thus, if $p_a = \sum_\mu c_{a\mu}\mu$ is the expression of p_a as a linear combination of monomials, we must have $\sum_a \gamma_a c_{a\mu} = 0$ for every μ. For $\mu = \ell t(p_{a_o})$, this sum contains only one term, namely $\gamma_{a_o} c_{a_o\mu}$.

Since $c_{a_o\mu} \neq 0$ by definition, we conclude that $\gamma_{a_o p} = 0$. That is, the polynomial p_{a_o} does not actually occur in the linear dependence, or the linear dependence is actually a dependence among the smaller set of polynomials gotten by deleting p_{a_o}. This smaller set will obviously satisfy the hypothesis of the lemma. Hence by induction on the size of the collection of polynomials, the conclusion follows.

5.2.4: Term order for SMT.

We will use the following term order on the monomials in the usual coördinates x_{ij} on M_{nm}. First we define an ordering on the x_{ij} themselves:

$$x_{ab} > x_{cd} \quad \Leftrightarrow \quad b < d, \quad \text{or} \quad b = d \text{ and } a < c. \tag{5.2.4.1}$$

This is just the lexicographic order on the pairs (i, j), with the second element being taken first. We extend this order to a monomial ordering by the following criterion. Given two monomials

$$\mu = \prod_{ab} x_{ab}^{m_{ab}}, \qquad \text{and} \qquad \mu' = \prod_{ab} x_{ab}^{m'_{ab}}, \tag{5.2.4.2}$$

we say that $\mu \geq_T \mu'$ provided that

i) $\deg(\mu) > \deg(\mu')$, or

ii) $\deg(\mu) = \deg(\mu')$, but $m_{11} > m'_{11}$, or

iii) For some pair of indices (a_o, b_o), we have $m_{ab} = m'_{ab}$ for $b > b_o$ and for $b = b_o$, but $a > a_o$, but $m_{a_o b_o} = m'_{a_o b_o}$.

This ordering is called the *graded lexicographic oder* for the given ordering of the variables x_{ab}.

It is easy to compute the leading term of a generator δ_I with respect to the term order $\geq_T$.

5.2.5: Lemma: Given a set $I = \{i_1 < i_2 < \ . \ . \ .i_k\}$ of indices, then with respect to the term order $\geq_T$ defined above, the leading term of δ_I is the product of the diagonal elements of the matrix defining δ_I:

$$\ell t(\delta_I) = \prod_{a=1}^{k)} x_{i_a a}. \tag{5.2.6}$$

Proof: The proof is straightforward. The expansion of δ_I in monomials is given by the standard expression of a determinant as the alternating sum of products of entries, one from each row and column. According to the definition of the ordering $\geq_T$, the largest monomial among all these products must contain the largest available entry, which is $x_{i_1 1}$. Again appealing to our knowledge of determinants, we may say that the monomials in the expansion that contain $x_{i_1 1}$ as a factor are the product of $x_{i_1 1}$ with the determinant of the $(k-1) \times (k-1)$ submatrix obtained by eliminating the row and column containing $x_{i_1 1}$ (that is, eliminating the first row and first column). Among all these terms, the largest must contain the largest possible of the entries of this submatrix, which by inspection is $x_{i_1 2}$. Continuing in this fashion, we conclude that the stated monomial is indeed the leading term of δ_I.

We can also readily extend this lemma to a description of the leading term of any monomial in the δ_I.

5.2.7: Lemma. Let $Q = \prod_{b=1}^{r} \delta_{I_b}$ be a product of generators δ_I of the algebra $P(M_{nm})^{U_m}$. Let T_Q be the tableau corresponding to Q. That is, T_Q is the array of boxes obtained by juxtaposing left to right the column tableaux defined by the I_b. (Without loss of generality, we may assume that I_b decrease in length as b increases, so that the shape of T is a standard Young diiagram.) Then

$$\ell t(Q) = \prod_{ij} x_{ij}^{e_{ij}},$$

where e_{ij} is the number of times the index i appears in row j of the tableaux T.

Proof: Indeed, we can restate the formula of the lemma 5.2.5 for the leading term of δ_{I_b} so that it becomes the statement of this lemma for the case of a single column. That is, in the product $\ell t(\delta_I)$ the coördinate x_{ij} appears if and only if i is the entry of I in row j. Since the tableau T_Q is obtained by juxtaposing the column tableaux of the δ_{I_b}, and since the leading term of a product is just the product of leading terms, the lemma follows.

5.2.8: Corollary. The leading term of a standard monomial determines the tableau of that monomial. Hence, a standard monomial is determined by its leading term.

Proof: Since the entries in a row of a semistandard tableau T increase from right to left, each row of such a T is determined by its content – that is, by the number of entries of each index a in it. According to the lemma, it is exactly this information that is supplied by $\ell t(\delta_T)$.

Now we may assert the main goal of this discussion:

5.2.9: Corollary. The standard monomials are linearly independent.

Proof: Indeed, we have just seen that any standard monomial is determined by its leading term. Combining this with Lemma 5.2.3, this implies the corollary.

Here is a supplementary result. It is not really required, but is nice to know.

5.2.10: Corollary. Any monomial in the generators δ_I of $P(M_{nm})^{U_m}$ has highest term equal to the highest term of a standard monomial.

Proof: Given any tableau T obtained by juxtaposing column tableau, we may rearrange the elements in each row to put them in weakly increasing order. Let the resulting tableau be $\tilde{T}$. From Lemma , we know that $\tilde{T}$ has the same leading term as does T. We claim that $\tilde{T}$ is semistandard. By what we have done so far, it is clear that the shape of $\tilde{T}$ is a diagram, and that the entries in each row of $\tilde{T}$ are weakly increasing. Thus, to know that $\tilde{T}$ is semistandard, we only need to be sure that the entries in each column of $\tilde{T}$ are strictly increasing. So consider two entries of $\tilde{T}$ in column c of $\tilde{T}$, in rows s and $s+1$. Let a be the entry in row s, and let b be the entry in row $s+1$. We want to know that $a < b$.

The fact that b appears in column c means that there are c entries in row $s+1$ of $\tilde{T}$, of size b or less, since the rows of $\tilde{T}$ are weakly increasing. The content of each row of $\tilde{T}$ is the same as the content of that row in T. Thus, there are at least c entries in row $s+1$ of T that have size c or less. The columns of T are strictly increasing. Therefore, in row s of T, there are at least c entries of size $b-1$ or less. This means that, in $\tilde{T}$, the entry in column c of row s is at most $b-1$; that is $a \leq b-1$. This implies that $\tilde{T}$ is semistandard, as desired.

5.3: Proof of Standard Monomial Theory II: spanning.

We now turn to showing that the standard monomials in $P(M_{nm})^{U_m}$ span the algebra. Since we already know that they are linearly independent, this will imply that they are a basis. In other words, it will finish proving Standard Monomial Theory.

To know that the standard monomials span, we will describe a different basis for each representation ρ_n^D. These bases come with a combinatorial description attached, and it turns out that this description can be translated easily into the semistandard tableau enumeration of standard monomials. This will show that there are the correct number of standard monomials, and so will finish the proof.

The second basis for ρ_n^D is obtained by considering restrictions of representations from GL_n to GL_{n-1} or more precisely, to $GL_{n-1} \times GL_1$ embedded as block diagonal matrices in GL_n. It turns out that this restriction has multiplicity one as a representation of $GL_{n-1} \times GL_1$, and that we can describe explicitly which representations of $GL_{n-1} \times GL_1$ occur.

This result follows by refining the study of $P(M_{nm})$ as a $GL_n \times GL_m$ module. It turns out that, if we restrict the action of $GL_n \times GL_m$ on M_{nm} to $(GL_{n-1} \times GL_1) \times GL_m$, this action is still multiplicity free.

5.3.1: Theorem. (Branching from GL_n to GL_{n-1}):

a) The action of $(GL_{n-1} \times GL_1) \times GL_m$ on M_{nm}, obtained by restriction from the action of $GL_n \times GL_m$, is multiplicity free.

b) The ring $P(M_{nm})^{U_{n-1} \times U_m}$ of $(GL_{n-1} \times GL_1) \times GL_m$ highest weight vectors is a polynomial ring generated by the elements

$$\delta_k =, \qquad \text{and} \qquad \delta'_k = . \tag{5.3.1.1}$$

where k ranges from 1 to $\min(n, m)$. (Note that, in case $n \leq m$, we have $\delta'_n = \delta_n$, so there are only $2n-1$ generators.)

c) We have the branching rule

$$(\rho_n^D)_{|Gl_{n-1}\times GL_1} \simeq \oplus_E \ \rho_{n-1}^E \otimes \psi^e, \tag{5.3.1.2}$$

where: i) E ranges over all diagrams such that
a) $E \subset D$ and
b) $D - E$ is a skew row.
ii) The exponent $e = e(D, E)$ is equal to $^\#(D) -^\# (E)$.

Remarks: i) (Statement b) means that E is obtained by erasing at most one box from each column of D).
ii) The quantity $^\#(D) -^\# (E)$ is equal to the number of boxes removed from D to produce E.

Proof: For convenience, we will assume in this argument that $m < n$. The cases when $m \geq n$ can be handled with some modifications of what we do below.

We again use the sufficiency condition 4.3.3 a): the action of $(GL_{n-1}\times GL_1)\times GL_m$ will be multiplicity-free if the subgroup $B_{n-1} \times GL_1 \times B_m$ has a dense orbit. Let J_{nm} here denote the $n \times m$ matrix with ones down the main diagonal (i.e., in the (i, i) entries, and zeroes elswhere. Let E_n denote the matrix all of whose entries are zero, except for 1s in the last row. We claim that the $B_{n-1} \times GL_1 \times B_m$ orbit of $J_{nm} + E_n$ is dense in M_{nm}.

We can think of $M_{nm} \simeq M_{(n-1)m} \oplus M_{1m}$, where $M_{(n-1)m}$ consists of the first $n-1$ rows of M_{nm}, and M_{1m} makes up the last row. Note that, because of our restriction to the case $m < n$, the matrix J_{nm} belongs to $M_{(n-1)m}$. As in the proof of (GL_n, GL_m)-duality, we know that the $B_{n-1} \times B_m$ orbit of J_{nm} is dense in $M_{(n-1)m}$. Moreover, as calculated there, we know that the stabilizer of J_{nm} in $B_{n-1} \times B_m$ includes the group S_J consisting of matrices

$$\begin{bmatrix} a_1 & 0 & 0 & . & . & . & . & . & . & 0 \\ 0 & a_2 & 0 & . & . & . & . & . & . & 0 \\ 0 & 0 & a_3 & . & . & . & . & . & . & 0 \\ . & . & . & . & . & . & . & . & . & . \\ . & . & . & . & . & . & . & . & . & . \\ . & . & . & . & . & . & . & . & . & . \\ . & . & . & . & . & . & . & . & . & . \\ . & . & . & . & . & . & . & . & . & . \\ 0 & 0 & 0 & . & . & . & . & . & . & a_n \end{bmatrix} \times \begin{bmatrix} a_1^{-1} & 0 & 0 & . & . & . & 0 \\ 0 & a_2^{-1} & 0 & . & . & . & 0 \\ 0 & 0 & a_3^{-1} & . & . & . & 0 \\ . & . & . & . & . & . & . \\ . & . & . & . & . & . & . \\ . & . & . & . & . & . & . \\ 0 & 0 & 0 & . & . & . & a_m^{-1} \end{bmatrix}. \tag{5.3.1.4}$$

This group can easily be checked to act on E_n to produce a dense orbit in M_{1m}. Combining this with what we already know, we can conclude that $J_{nm} + E_m$ generates a dense $B_{n-1} \times GL_1 \times B_m$ orbit, as claimed. Thus, statement a) of the theorem is proved.

b) It is easy to check that the δ_k and δ_k' are $(B_{n-1} \times GL_1) \times B_m$ eigenvectors. According to Theorem 4.3.4 c), to completely verify b, we need to check that the δ_k and δ_k' are must be prime $B_{n-1} \times GL_1 \times B_m$ eigenvectors, and that they exhaust the list of prime eigenvectors.

We already know from our discussion of (GL_n, GL_m) duality, that the δ_k are prime. To show that the δ_k' must be prime $B_{n-1} \times GL_1 \times B_m$ eigenvectors, consider their eigencharacters. The characters restricted to B_m are the fundamental dominant weights for GL_m. If we were to have a factorization $\delta_k' = f_1 f_2$ into $B_{n-1} \times GL_1 \times B_m$ eigenvectors, then the B_m weight of f_1 and f_2 would have to be either ψ_k, the k-th fundamental dominant weight of GL_m, or just 1. But according to $(GL_n, GL)n)$ duality and the highest weight theory, the only B_m fixed vectors in $P(M_{nm})$ are the constant functions. Therefore the factorization of δ_k' has to be trivial; in other words, δ_k' is also prime.

We also want to know that the δ_k and the δ_k' are the full list of generators. This can be verified by counting. It is easy to check that the stabilizer of $J_{nm} + E_n$ inside the product $A_n \times A_m$ of diagonal tori is the subgroup $\tilde{S}$ consisting of matrices

$$\begin{bmatrix} a_n & 0 & 0 & \cdots & & & 0 \\ 0 & a_n & 0 & \cdots & & & 0 \\ 0 & 0 & a_n & \cdots & & & 0 \\ \cdot & \cdot & \cdot & \cdots & & & \cdot \\ \cdot & \cdot & \cdot & \cdots & & & \cdot \\ \cdot & \cdot & \cdot & \cdots & & & \cdot \\ 0 & \cdots & & a_{m+1} & \cdots & & 0 \\ 0 & \cdots & & & a_{m+2} & \cdots & 0 \\ \cdot & \cdot & \cdot & \cdots & & & \cdot \\ \cdot & \cdot & \cdot & \cdots & & & \cdot \\ 0 & 0 & 0 & \cdots & & & a_n \end{bmatrix} \times \begin{bmatrix} a_n^{-1} & 0 & 0 & \cdots & 0 \\ 0 & a_n^{-1} & 0 & \cdots & 0 \\ 0 & 0 & a_n^{-1} & \cdots & 0 \\ \cdot & \cdot & \cdot & \cdots & \cdot \\ \cdot & \cdot & \cdot & \cdots & \cdot \\ \cdot & \cdot & \cdot & \cdots & \cdot \\ 0 & 0 & 0 & \cdots & a_n^{-1} \end{bmatrix}. \qquad (5.3.1.5)$$

This subgroup has dimension $n - m$, or codimension $2m$ in the group $A_n \times A_m$ of diagonal matrices in $GL_n \times GL_m$. Thus, there can be at most 2m independent $B_{n-1} \times GL_1 \times B_m$ eigencharacters that appear in $P(M_{nm})^{U_{n-1} \times U_m}$. This number is already accounted for by the proposed generators, so they must exhaust the generating set of prime $B_{n-1} \times GL_1 \times B_m$ eigenvectors. This finishes the proof of statement b)

c) To turn b) into the explicit branching rule of part c), note that both δ_k and δ'_k have the same B_m weights, but different $B_{n-1} \times GL_1$ weights. If we want to produce $B_{n-1} \times GL_1 \times B_m$ eigenvectors with a given B_m weight, we can start by using only the elements δ_k, and then we can replace various of the δ_k with δ'_k, one step at a time. Doing this has the effect of taking one box off of one column of the diagram corresponding to the B_{n-1} weight, and adding one to the GL_1 weight. Thus, this process of replacement will produce exactly the diagrams E specified in part c). This concludes the proof of the branching rule from GL_n to GL_{n-1}.

We can use this branching rule to describe an A_n -eigenbasis for an irreducible representation ρ_n^D of GL_n. We know from part c) of the Theorem that under the action of $GL_{n-1} \times GL_1$, ρ_n^D breaks up into invariant and irreducible subspaces, ρ_{n-1}^E, one for each diagram E such that $E \subset D$ and $D - E$ is a skew row. In turn, each ρ_{n-1}^E breaks up into $GL_{n-2} \times (GL_1)^2$ invariant and irreducible subspaces ρ_{n-2}^F, one for each diagram F such that $F \subset E$, and $E - F$ is a skew row. Continuing in this fashion, we will end up with a collection of subspaces that are invariant and irreducible under $(GL_1)^n = A_n$. Since all irreducible representations of A_n are one-dimensional, this final decomposition gives an A_n eigenbasis for ρ_{n-1}. A given basis vector (more correctly, the line it generates) will be parametrized by a nested sequence of diagrams

$$D = D_n \supset D_{n-1} \supset D_{n-2} \supset \quad . \; . \; . \supset D_1,$$

such that $D_k - D_{k-1}$ is a skew row, and also, D_k has depth at most k (i.e., at most k rows). In particular, this shows that we can count the dimension of ρ_n^D by the collection of such nested sequences of diagrams. The basis constructed here was first described by Gelfand and Tsetlin [ ], and is called the *Gelfand-Tsetlin basis.*

We now observe that we can easily translate collections of diagrams as described above into semistandard skew tableaux, and vice versa. More precisely, given the nested pair of diagrams $E \subset D$, we can determine E either by telling what it is, or by telling what it isn't. That is, we can determine E by starting with D and specifying $D - E$, the boxes to be removed to obtain E. For a nested sequence sequence such as ???, we can use this idea as follows. Starting with $D = D_n$, label all the boxes of $D + n - D_{n-1}$ with the number n. Then label all the boxes of $D_{n-1} - D_{n-2}$ with the number $n - 1$. Continue in this fashion, labeling the boxes of $D_k - D_{k-1}$ with the number k. When we are done, we will have filled the boxes of D with the integers from 1 to n, that is, we will have constructed a tableau $T = T(\{D_k\})$.

We claim that this tableau is semistandard. It is clear that the rows will be weakly increasing, since the set of numbers equal to k or less fill the diagram D_k - that is, fill each row from the left up to a certain place, with no gaps. The columns are weakly increasing for the same reason. The condition that the columns be strictly increasing is then seen to be equivalent to the statement that a give number, say k, occurs at most once in each column, which is the same as to say, that $D_k - D_{k-1}$ is a skew row.

Conversely, we can argue in reverse, and show that, it T is a semistandard tableau in the diagram D, and if we set $D_k(T)$ to be the set of boxes of D that contain entries equal to k or less, then the $D_k(T)$ form a nested sequence of subdiagrams of D, with the property that $D_k(T)$ has at most k rows, and $D_K(T) - D_{k-1}(T)$ is a skew row.

This discussion has established the following statement¿

5.3.2: Proposition. There is a bijection, as described just above, between the set of semistandard tableaux filling a diagram D and nested sequences $\{D_k\}$ of subdiagrams of D, satisfying the conditions of Theorem 5.3.1c).

5.3.3: Corollary. The number of semistandard tableaux filling D with integers from 1 to n is the dimension $\dim \rho_n^D$.

Since we know from our discussion of linear independence of the standard monomials that the standard monomials contained in ρ_n^D are counted by semistandard tableau in D, this corollary implies that the semistandard monomials span $\rho_n^D \otimes (\rho_m^D)^{U_m}$. This establishes Hodge's theorem, that the standard monomials are a basis for $P(M_{nm})^{U_m}$.

5.3.4: The Gelfand-Tsetlin Cone.

A key part of our argument for standard monomial theory depended on computing leading terms for elements in $P(M_{nm})^{U_m}$. These leading terms themselves have an interesting structure. Let $A \subset P(\mathbf{C}^n)$ be a subalgebra of a polynomial ring, and suppose that we have a term order on the monomials. Consider the leading terms of non-zero elements of A. Because of the multiplicativity of leading terms (Lemma 5.2.2), the leading term of elements of A will be a semigroup under multiplication. More precisely, the semigroup of monomials on $\mathbf{C}^n$ is just the free abelian semigroup $(\mathbf{Z}^+)^n$, and the leading terms of elements of A will form a subsemigroup $\ell t(A)$. It is obviously of interest to describe $\ell t(A)$. In fact, in the case when $\ell t(A)$ is finitely generated, the semigroup ring of $\mathbf{C}(\ell t(A))$ will be a Noetherian ring (even an integral domain), and it is shown in [ ] that $\mathbf{C}(\ell t(A))$ strongly resembles A, in the sense that it is a flat deformation of A.

It turns out that the Gelfand-Tsetlin basis suggests a nice description of $\ell t(P(M_{nm})^{U_m})$. This is based on the fact that a diagram is described by a decreasing sequence of positive integers – the lengths of the rows of the diagram. Thus given a nested sequence D_k of diagrams as in Theorem 5.3.1 c), we can describe this sequence by specifying the numbers r_{km}, the length of the m-th row of D_k. These numbers r_{km} will not be arbitrary – they must satisfy some relations imposed by their origin.

For one, the fact that D_k is a diagram with at most k rows is expressed by the inequalities

$$r_{km} \geq r_{k(m+1)}, \qquad \text{and} \qquad r_{km} = 0 \quad \text{when } m > r.$$

Second, the fact that $D_{k-1} \subset D_k$ is expressed by the inequalities

$$r_{km} \geq r_{(k-1)m}.$$

Finally, the fact that $D_k - D_{k-1}$ is a skew row is expressed by the inequalities

$$r_{k(m+1)} \leq r_{(k-1)m}.$$

These inequalities can be summarized by the *interlacing conditions*

$$r_{k1} \geq r_{(k-1)1} \geq r_{k2} \geq r_{(k-1)2} \geq r_{k3} \geq \quad . \quad . \quad . \quad . \quad \geq r_{k(k-1)} \geq r_{(k-1)(k-1)} \geq r_{kk}. \qquad (5.3.4.1)$$

The full structure of these inequalities can be pleasantly visualized by arranging the r_{km} in a triangular array, with the r_{km} arranged left to right (i.e., decreasing order), with the k-th row just above the $(k+1)$-th row. We illustrate this below.

$$\begin{array}{ccccccccccccc} r_{n1} & & r_{n2} & & r_{n3} & & \cdots & & & & \cdots & & r_{nn} \\ & r_{(n-1)1} & & r_{(n-1)2} & & r_{(n-1)3} & & \cdots & & & & r_{(n-1)(n-1)} & \\ & & r_{(n-2)1} & & r_{(n-2)2} & & r_{(n-2)3} & & \cdots & & r_{(n-2)(n-2)} & & \end{array}$$

$$
\begin{array}{ccccccccccccc}
\cdot & & \cdot & & \cdot & & \cdot & & \cdot & & \cdot & & \cdot \\
& \cdot & & \cdot & & \cdot & & \cdot & & \cdot & & \cdot & \\
& & \cdot & & \cdot & & \cdot & & \cdot & & \cdot & & \\
& & & r_{21} & & r_{32} & & r_{33} & & & & & \\
& & & & r_{21} & & r_{22} & & & & & & \\
& & & & & r_{11} & & & & & & &
\end{array}
\qquad (5.3.4.2)
$$

This array is called a *Gelfand-Tsetlin pattern* or GT pattern for short. It is characterized by the conditions that each number in it is (weakly) larger than the (one or two) numbers just above and below it to the right (and hence, weakly smaller than the one or two numbers just above it and below it to the left).

It is easy to see from the above description that if we have two GT patterns, then the pattern formed by adding the entries in each position will again satisfy the conditions defining GT patterns, and hence will be one. Thus, under addition of entries, the GT patterns form a semigroup.

Furthermore, we can think of this semigroup as being a *lattice cone*: the intersection of a convex cone (in $\mathbf{R}^{\frac{n(n+1)}{2}}$)with the integer lattice $\mathbf{Z}^{\frac{n(n+1)}{2}}$. Let the coördinates in $\mathbf{R}^{\frac{n(n+1)}{2}}$ be x_{ab} for $1 \leq b \leq a \leq n$. Consider the inequalities

$$x_{ab} \geq 0, \qquad x_{ab} \geq x_{(a-1)b}, \qquad \text{and} \qquad x_{ab} \geq x_{(a+1)(b+1)}. \tag{5.3.4.3}$$

These define a convex cone in $\mathbf{R}^{\frac{n(n+1)}{2}}$. The GT patterns can be thought of as the integer vectors in this cone. Thus, we will call this lattice cone the *Gelfand-Tsetlin (or GT) cone.* (It is also referred to as the Gelfand-Tsetlin polytope).

Tracing through the several identifications we have made above (GT patterns with nested sequences of Young diagrams with semistandard tableaux with leading terms of standard monomials), we see that we can make the following statement:

5.3.5: SMT Theorem 3. Under the term order defined in 5.2.4, the leading terms of elements of the algebra $P(M_{nm})^{U_m}$ form a semigroup isomorphic to the Gelfand-Tsetin cone.

More or less this result can be found in [ ].

By the machinery of [ ], this implies

SMT Theorem 4: The algebra $P(M_{nm})^{U_m}$ is a flat deformation of the semigroup ring of the Gelfand-Tsetlin cone.

This result is mostly due to Gonciulea and Lakshmibai [ ], with some addenda [ ], [ ].

5.4: The Gelfand-Tsetlin poset and Hibi Rings.

The SMT Theorem 4 is quite elegant, but does not quite translate the full import of of SMT Theorem 2. What has happened to the column tableaux, and what has happened to the polynomial subrings that make up $P(M_{nm})^{U_m}$? A further reformulation of SMT Theorem 4 can restore these features of the original theory.

We have described GT patterns as certain arrays of numbers, some of which must be weakly larger or smaller than others. The relevant inequalities can be described effectively described by arranging the numbers in a certain pattern (the triangular array (5.3.4.2)), and requiring that the number in a given position be greater than the numbers in certain other positions. An alternate way of describing this situations is to consider the positions themselves as elements of a set, and to describe the relevant inequalities as (partial) order relations on this set. A GT pattern would then be interpreted as a function on this set, and all the inequalities required of GT patterns can be summarized by saying that the function it defines should be order-preserving. We will call the resulting partially ordered set the *Gelfand-Tsetlin (GT) poset.* Then the Gelfand-Tsetlin cone can be described as the semigroup of all non-negative integer-valued order preserving functions on the GT poset. The GT poset is simply the triangular array of formula ???, but with only the positions indicated, and no numbers inserted. The GT poset for GL_6 is shown below.

O O O O O O

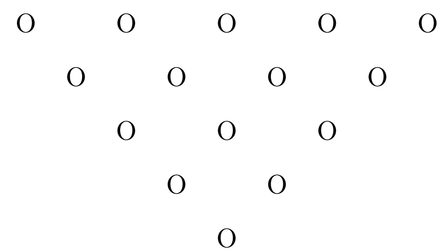

(5.4.1)

The order relations are given by requiring that a given element is greater than either element to its left in the row just above or just below, and then requring the standard condition of an order relation, namely that it be transitive. This implies that a given element of the GT poset dominates all elements that can be reached from it by steps to the northeast or southeast.

This description of the GT cone, as order-preserving functions on the GT poset, can be formulated abstractly, for any partially ordered set. Thus, let X be a poset, and let $\mathbf{R}^X$ denote the real vector space of all real-valued functions on X. Denote by $(\mathbf{R}^+)^X$ be the positive orthant in $\mathbf{R}^X$ – the cone of non-negative valued functions on Γ. Let $(\mathbf{R}^+)^X_{\geq}$ denote the subcone of non-negative real-valued order-preserving functions.

We can define similar objects for integer-valued functions. Let $\mathbf{Z}^X$ be the group of integer-valued functions on X. Then $\mathbf{Z}^X$ is a lattice in $\mathbf{R}^X$ – a discrete, spanning subgroup. We can further consider $(\mathbf{Z}^+)^X$, the semigroup of non-negative integer valued functions on X, and $(\mathbf{Z}^+)^X_{\geq}$, the subsemigroup of order preserving non-negative integer valued function on X. We obviously have the relations

$$\mathbf{R}^X \quad \supset \quad (\mathbf{R}^+)^X \qquad\qquad\qquad \supset \quad (\mathbf{R}^+)^X_{\geq}, \tag{5.4.2a}$$

and

$$\mathbf{Z}^X \quad \supset \quad (\mathbf{Z}^+)^X = (\mathbf{R}^+)^X \cap \mathbf{Z}^X \quad \supset \quad (\mathbf{Z}^+)^X_{\geq} = (\mathbf{R}^+)^X_{\geq} \cap \mathbf{Z}^X, \tag{5.4.2b}$$

In particular, $(\mathbf{Z}^+)^X$ is the lattice cone of integer points in the positive orthant, and $(\mathbf{Z}^+)^X_{\geq}$ is a subcone of this. SMT Theorem 4 invokes $(\mathbf{Z}^+)^X_{\geq}$ when X is the GT poset.

We would like to understand the structure of $(\mathbf{Z}^+)^X_{\geq}$. The main observation to make is that, for a finite set X, the generic real-valued function on X will be one-to-one. Since $\mathbf{R}$ is totally ordered, a generic real-valued function implicitly defines a total ordering $\geq_f$ on X. In order for f to be order-preserving, the total order $\geq_f$ should be compatible with the given partial ordering, which we will denote by $\geq_X$ when we need to distinguish it from other orderings under discussion; that is, for any pair of elements x and y of X, we must have that

$$x \geq_X y \Rightarrow x \geq_f y.$$

Conversely, if this condition is satisfied, f will be order-preserving.

Given any set X, a total ordering of X can be produced simply by numbering the elements of X:

$$X = \{x_1, x_2, x_3, \quad . \quad . \quad . \quad . \quad , x_s\}.$$

Given one such total order $\geq_o$, any other total order is produced by permuting the elements of X. It has the form $\geq_\nu = \geq_o \circ \nu$, where ν is a permutation of X. The set

$$(\mathbf{R}^+)^X_o = \{f : X \to \mathbf{R} : f(x_1) \geq f(x_2) \geq \quad . \quad . \quad . \quad \geq f(x_s) \geq 0\} \tag{5.4.3}$$

is a closed convex cone in $\mathbf{R}^X$. It will be a subcone of $(\mathbf{R}^+)^X_{\geq_X}$ exactly when $\geq_o$ is compatible with $\geq_X$.

Independent of his, $(\mathbf{R}^+)^X_o$ has some remarkable properties in its own right. In order to state them, we need some terminology. Let W_C be the subgroup of $GL(\mathbf{R}^X)$ generated by all permutations of X and by sign changes (multiplications by functions taking the values ± 1). It is well-known to be generated by reflections. A set of generators consists of the pairwise exhanges $x_1 \leftrightarrow x_{i+1}$, and the sign change $f(x_n) \to -f(x_n)$ (with all other values of f being fixed).

A subset Y of a poset X is called *increasing* if it contains everything greater than any of its elements. That is, Y is increasing if and only if $y \in Y$ and $y' \geq_X y$ imply that $y' \in Y$. We note that if Y and Z are increasing subsets of X, then also the intersection $Y \cap Z$ is increasing, and so is the union $Y \cup Z$.

For a total order, the increasing subsets are just those consisting of the largest so many elements. Thus, for the total order $\geq_o$, the increasing subsets are exactly the $E_k = \{x_1, x_2, x_3, \ . \ . \ . \ , x_k\}$. There is one of each cardinality $k \leq s =^{\#} (X)$, and they form an increasing chain of subsets: $E_1 \subset E_2 \subset E_3 \subset \ . \ . \ . \ \subset E_s = X$. This chain is *complete* in the sense that it cannot be enlarged: there is exactly one subset of each cardinality up to $^{\#}(X)$. Conversely, such a complete increasing chain of subsets defines a total order on X, by declaring these to be the increasing subsets of X. The resulting total order will be compatible with $\geq_X$ if and only if each of the subsets in the chain is increasing with respect to $\geq_X$.

5.4.4: Proposition. a) The cone $(\mathbf{R}^+)^X_o$ is a fundamental domain for W_C in $\mathbf{R}^X$, in the sense that any $f \in \mathbf{R}^X$ is equivalent to a unique element of $(\mathbf{R}^+)^X_o$ by means of some $w \in W_C$.

b) The cone $(\mathbf{R}^+)^X_o$ is an *integral simplicial cone*, in the sense that $(\mathbf{Z}^+)^X_{\geq_o} = \mathbf{Z}^X \cap (\mathbf{R}^+)^X_{\geq_o}$ is a free abelian semigroup, on the basis consisting of the characteristic functions of the subsets of X that are increasing with respect to $\geq_o$.

Proof: The proofs of these statements are straightforward and are left to the reader. For part b), we mention that if f is non-negative function on X that is order-preserving for $\geq_o$, and if the E_j are the increasing subsets with respect to $\geq_o$, as above, then we can write

$$f = f(x_n)\chi_X + \sum_{i=1}^{s-1}(f(x_i) - f(x_{i+1}))\chi_{E_i},$$

where the χ_{E_i} are the characteristic functions of the set E_i. This formula exhibits f as a non-negative linear combination of the χ_{E_i}. If f takes values in $\mathbf{Z}$, then this is an integral combination of the χ_{E_i}, so that these also form an integral basis for $(\mathbf{Z}^+)^X_{\geq_o}$.

Using the cones $(\mathbf{Z}^+)^X_{\geq_t}$, for total orders $\geq_t$ on X, we can give a description of $(\mathbf{Z}^+)^X_{\geq_X}$ that is reminiscent of the original version of standard monomial theory.

5.4.5: Proposition. (Abstract SMT): a) The cone $(\mathbf{Z}^+)^X_{\geq_X}$ is generated by the characteristic functions χ_Y of all increasing subsets $Y \subset X$. More precisely, any f in $(\mathbf{Z}^+)^X_{\geq_X}$ may be written as a positive sum of the characteristic functions of a nested family of increasing subsets.

b) The relations $\chi_Y + \chi_Z = \chi_{Y\cap Z} + \chi_{Y\cup Z}$, for Y and Z increasing subsets of X, is a set of defining relations for the semigroup $(\mathbf{Z}^+)^X_{\geq_X}$ and the generating set of χ_Y.

c) The cone $(\mathbf{Z}^+)^X_{\geq_X}$ is the non-overlapping union of the integrally simplicial subcones $(\mathbf{Z}^+)^X_{\geq_t}$, where $\geq_t$ runs over all total orderings of X compatible with the given partial order $\geq_X$.

d) The set of total orderings of X compatible with $\geq_X$ is in bijection with the set of complete chains of increasing subsets of X.

Remarks: i) If one looks at the semigroup ring $\mathbf{C}(\mathbf{Z}^+)^X_{\geq_X}$, then the subrings $\mathbf{C}(\mathbf{Z}^+)^X_{\geq_t}$ for compatible total orders $\geq_t$ are polynomial rings, by part b) of Proposition ??? Thus, part c) of Proposition ??? tells us that $\mathbf{C}(\mathbf{Z}^+)^X_{\geq_X}$ is the almost direct sum of the polynomial rings $\mathbf{C}(\mathbf{Z}^+)^X_{\geq_t}$. For the case when X is the Gelfand-Tsetlin poset, we will see that this exactly gives us the decomposition of $P(M_{nm})^{U_m}$ described in SMT Theorem 2. This is the main justification for calling this result "Abstract Standard Monomial Theory". However, the other parts of the Proposition also are germane for the form of SMT.

ii) Although we have so far not discussed the relations in $P(M_{nm})^{U_m}$, in fact, they are quadratic relations that resemble the relations of part b) of the Proposition. The exact form of the relations in $P(M_{nm})^{U_m}$ was instrumental in the early proofs of SMT, but we have not needed them. However, below we will show that SMT has strong implications for the form of the relations in $P(M_{nm})^{U_m}$.

iii) We will also show that statement d) of the Proposition has an interpretation in terms of SMT, and in that context, gives us an interpretation or description of the maximal chains of column tableaux.

Proof: a) Given an order preserving $f : X \to \mathbf{R}^+$, let $t_1 < t_2 < t_3 < \ . \ . \ . \ t_p$ be the values of f, and set $Y_i = \{y \in X : f(y) \geq t_i$. Then we may write

$$f = t_1\chi_X + \sum_{i>1}(t_i - t_{i-1})\chi_{Y_i}. \tag{5.4.5.1}$$

Since f is order-preserving, each Y_i must be an increasing set. Moreover, clearly $Y_{i+1} \subset Y_i$. Thus, this formula exhibits f as a linear combination of characteristic functions of increasing subsets. If f takes non-negative integer values, then the coefficients of the expansion are also non-negative integers. This establishes statement a).

b) Suppose Y_i and Z_j are nested families of increasing subsets, strictly ordered by inclusion (that is, $Y_{i+1} \subset Y_i$, and similarly for the Z_j. Consider two positive integer linear combinations $f = \sum_i a_i \chi_{Y_i}$, and $g = \sum_j b_j Z_j$. Then $f + g$ is also in $(\mathbf{Z}^+)^X_{\geq X}$, and so can be written in the form (5.4.5.1) above. However, the collection of Y_i and Z_j together might not be totally ordered by inclusion. We want to show that we can rewrite f as a sum of characteristic functions of a totally ordered sequence of increasing subsets, using only the relations of statement b).

We will argue by induction on the ℓ^1 norm of $f + g$, that is, on

$$||f + g||_1 = \sum_{x \in X} (f(x) + g(x)).$$

Note that, for non-negative functions f and g, we have $||f + g||_1 = ||f||_1 + ||g||_1$.

Suppose without loss of generality that $a_1 \leq b_1$. Set $\tilde{f} = \sum_{i \geq 2} a_i \chi_{Y_i}$, and define $\tilde{g}$ similarly. Then we have

$$f+g = a_1\chi_{Y_1}+\tilde{f}+b_1\chi_{Z_1}+\tilde{g} = a_1(\chi_{Y_1}+\chi_{Z_1})+(b_1-a_1)\chi_{Z_1}+\tilde{f}+\tilde{g} = a_1(\chi_{Y_1 \cup Z_1}+\chi_{Y_1 \cap Z_1})+(b_1-a_1)\chi_{Z_1}+\tilde{f}+\tilde{g}$$

$$= a_1\chi_{Y_1 \cup Z_1} + (a_1\chi_{Y_1 \cap Z_1} + (b_1 - a_1)\chi_{Z_1} + \tilde{f} + \tilde{g}) = a_1\chi_{Y_1 \cup Z_1} + h,$$

where $h = a_1\chi_{Y_1 \cap Z_1} + (b_1 - a_1)\chi_{Z_1} + \tilde{f} + \tilde{g}$. Since the ℓ^1 norm of h is less than that of $f + g$, we may assume by induction that we can rewrite h as a sum of the desired form using only the relations of statement b). (It might be objected that h is a sum of more than two functions, so the inductive hypothesis does not apply. However, we can successively rewrite sums of pairs of summands of h in the kosher way. After each rewriting, the number of summands needed to express h goes down by one, so finally h itself will be rewritten using the stated relations. This is consistent with induction on the ℓ^1 norm, since rewriting does not change that, and the ℓ^1 norm of a subsum is smaller than the ℓ^1 norm of the whole sum.) Since the support of h is evidently contained inside $Y_1 \cup Z_1$, the sum $a_1\chi_{Y_1 \cup Z_1} + h$ will also be in standard form, so $f + g$ will have been rewritten in standard form, using the relations of b). Hence, statement b) is proved.

We take statement d) as self-evident.

c) We have already shown how to write f in $(\mathbf{Z}^+)^X_{\geq X}$ as a linear combination of characteristic functions of a nested family of increasing subsets. to show c), it is enough to show that any such family can be augmented to a complete family of increasing subsets. From statement d), a complete family of nested subsets of X defines a total order on X, and the total order is compatible with the given order on X if and only if the subsets are all increasing.

To know we can augment a given nested family $\{Y_i\}$ of increasing subsets to a maximal one, it is enough to know that, if $\{Y_i\}$ is not complete, then we can add one more set to the collection to make a larger nested family of increasing subsets. If $\{Y_i\}$ is not complete, then there must be an index i such that $Y_i - Y_{i+1}$ has at least two elements. Select a maximal (that is, non-dominated in $Y_i - Y_{i+1}$) element y' of $Y_i - Y_{i+1}$. Then $Y_{i+.5} = Y_{i+1} \cup \{y'\}$ will be an increasing subset contained strictly between Y_i and Y_{i+1}, so we can add it to the collection to get a larger one. This proves statement c).

5.4.6: Fundamental Theorem of Distributive Lattices and Standard Monomial Theory.

It remains to connect the Abstract SMT with our results on $P(M_{nm})^{U_m}$. This involves elucidating the connection between column tableaux with the tableau order and the GT poset. The column tableaux are of course special semistandard tableaux - they are the basic constituents from which semistandard tableaux are made. As semistandard tableaux, they correspond to certain Gelfand-Tsetlin patterns. According to the recipe described above in the proof of Theorem 5.3.1 c), the i-th row of the Gelfand-Tsetlin pattern corresponding to a semistandard tableau T records the lengths of the rows of the diagrams consisting of the boxes of T filled with entries not exceeding $n + 1 - i$. If T is a column tableau, then all the subdiagrams of T will have rows of length 1. Therefore, the corresponding GT pattern P_T will consist entirely of ones and zeros, Since it defines an increasing function on the GT poset Γ, it is the characteristic function of an increasing subset of Γ. Conversely, every characteristic function of an increasing subset of Γ must come from a column tableau. From the recipe of correspondence, one sees that the larger the entries in a column tableau, the fewer 1s are in the corresponding GT pattern. The following result summarizes this discussion

5.4.6.1: Proposition: (S. Kim) a) There is a natural bijection between column tableaux and the collection of increasing subsets of the GT poset. In this bijection, the tableau order is translated to (reverse) inclusion.

b) In particular, maximal chains of column tableaux correspond bijectively to complete nested sequences of increasing subsets of Γ, equivalently to total orderings of Γ compatible with the given partial order.

c) Hence the polynomial subrings of $P(M_{nm})^{U_m}$ corresponding to maximal chains of column tableaux correspond to the semigroup rings of the subcones of $(\mathbf{Z}^+)^{\Gamma}_{\geq \Gamma}$ defined by compatible total orders on Γ.

This result is related to the theory of lattices, especially distributive lattices. Although the above account stands on its own, these connections help to put the above account of SMT in perspective, so we will discuss them here.

Let X be a poset. For an element x of X, consider the sets

$$L(x) = \{z \in X : z \leq x\} \qquad \text{and} \qquad G(x) = \{z \in X : z \geq x\}.$$

If y is another element in X, then the set $L(x) \cap L(y)$ is the subset of elements of X that are less then both x and y, and$G(x) \cap G(y)$ is the set of elements that are greater than both x and y. The poset X is called a *lattice* if for any pair x and y in X, there is an element z such that $L(x) \cap L(y) = L(z)$, and an element w such that $G(x) \cap G(y) = G(w)$. In other words, there is a largest element (namely z) that is smaller than both x and y, and there is a smallest element (namely w) that is larger than both x and y. If X is a lattice, then the element z is denoted by $x \wedge y$, and the element w is denoted $x \vee y$.

The operations $\wedge$ and $\vee$ satisfy some obvious identities. For example

$$x \wedge x = x \qquad x \wedge y = y \wedge x, \qquad (x \wedge y) \wedge z) = x \wedge (y \wedge z),$$

and analogous identities involving $\vee$ instead of $\wedge$. If X is the collection of subsets of a set P, ordered by inclusion, then $\wedge$ corresponds to set intersection, and $\vee$ corresponds to union. Union and intersection satisfy some identities that involve both of $\wedge$ and $\vee$. For example, for subsets of a set P, we have $R \cup (S \cap T) = (R \cup S) \cap (R \cup T)$, and likewise $R \cap (S \cup T) = (R \cap S) \cup (R \cap T)$. If X is a lattice such that the analogous statements hold for $\wedge$ and $\vee$:

$$x \vee (y \wedge z) = (x \vee y) \wedge (x \vee z), \qquad \text{and} \qquad x \wedge (y \vee z) = (x \wedge y) \vee (x \wedge z),$$

then X is called a *distributive lattice*. (Note that these identities can be thought of as saying that $\wedge$ distributes over $\vee$, and vice versa).

G. Birkhoff showed that, in fact, all distributive lattices were naturally lattices of subsets of some set. Here is the relevant version of his result, dubbed "The fundamental theorem of distributive lattices." An element x of a lattice L is called *meet irreducible* if it is not $y \wedge z$ for two elements of L, both not equal to x. We denote by $J(L)$ the subset of L consisting of meet irreducible elements.

5.4.6.2: Theorem. (Birkhoff) Let L be a distributive lattice. Then L is anti-isomorphic to the lattice of increasing subsets of a unique poset, namely the poset $J(L)$. Here "anti-isomorphic" means that, if $\alpha(x)$ is the increasing subset of $J(L)$ corresponding to $x \in L$, then $x \geq y \Rightarrow \alpha(x) \subseteq \alpha(y)$.

Proof: See for example [ ]

Remark: This result is usually formulated in terms of decreasing subsets. However, in our context, it is more appropriate to work with increasing subsets. Increasing subsets can be exchanged with decreasing subsets by reversing the sense of the order relation.

This result is related to SMT by virtue of the following facts.

5.4.6.3: Proposition. a) The collection of (semistandard) column tableaux with the tableau order is a distributive lattice. Precisely, given column tableaux

$$T_I = \begin{bmatrix} i_1 \\ i_2 \\ i_3 \\ . \\ . \\ . \\ i_k \end{bmatrix} \qquad \text{and} \qquad T_J = \begin{bmatrix} j_1 \\ j_2 \\ j_3 \\ . \\ . \\ . \\ j_\ell \end{bmatrix},$$

with $\ell \leq k$, we have

$$T_{I \wedge J} = \begin{bmatrix} \min(i_1, j_1) \\ \min(i_2, j_2) \\ \min(i_3, j_3) \\ . \\ . \\ . \\ \min(i_k, j_k) \end{bmatrix} \qquad \text{and} \qquad T_{I \vee J} = \begin{bmatrix} \max(i_1, j_1) \\ \max(i_2, j_2) \\ \max(i_3, j_3) \\ . \\ . \\ . \\ \max(i_\ell, j_\ell) \end{bmatrix},$$

where we make the agreement that $\min(i_b, j_b) = i_b$ if $b > \ell$.

b) The set of meet-irreducible column tableaux are precisely the tableaux with consecutive entries:

$$T_{[i,j]} = \begin{bmatrix} n+1-i \\ i+1 \\ i+2 \\ . \\ . \\ . \\ n-i+j \end{bmatrix},$$

for $1 \leq a \leq b \leq n$. As a poset, this is isomorphic to the GT poset.

Proof: Left as a pleasant exercise for the reader.

5.4.6.3: Concluding remark. The semigroup rings of the lattice cones are known as *Hibi rings* [ ]. They were in fact introduced in terms of lattices. Hibi considered the question: for what lattices L is the ring generated by elements of L, subject to relations $x \cdot y = (x \wedge y) \cdot (x \vee y)$ an integral domain? Hibi's answer was: when L is a distributive lattice. In that case, from the Fundamental Theorem for Distributive Lattices and Proposition 5.4.5, we conclude that the resulting algebra is the semigroup ring of $(\mathbf{Z}^+)^{J(L)}_{\geq J(L)}$.