Notes II: Representation theory.

2.1. Group representation basics.

A (complex) *matrix representation* of a group G is a homomorphism

$$\rho \colon G \to GL_n(\mathbb{C})$$

from G to the group of $n \times n$ invertible complex matrices. A (complex) *linear representation* of G is a homomorphism

$$\rho \colon G \to \mathrm{End}(V)$$

where V is a finite-dimensional vector space over $\mathbb{C}$. Of course every matrix representation induces a linear representation on $V = \mathbb{C}^n$, and conversely, choosing a basis in a linear representation on V induces a matrix representation.

Matrix representations ρ, ρ' are *similar* if there is a matrix A such that $\rho'(g) = A^{-1}\rho(g)A$ for all $g \in G$. Similar matrix representation induce isomorphic linear representations; conversely, isomorphic linear representations with arbitrary bases, or the same linear representation with two choices of basis, induce similar matrix representations.

The space V of a linear representation, with its G-action given by ρ, is also called a G-module. When ρ is understood from context we often write gv instead of $\rho(g)v$, for $g \in G$ and $v \in V$.

A *submodule* of V is a subspace $W \subseteq V$ such that $gW \subseteq W$ for all $g \in G$. We say that V is *irreducible* if its only submodules are 0 and V itself. We say that V is *completely reducible* if every submodule of V is a direct summand of V, which implies V is a direct sum of irreducible modules.

When the group G is $GL_n(\mathbb{C})$ or a subgroup thereof, we will chiefly be concerned with *polynomial representations*, in which the entries of the representing matrix $\rho(g)$ are polynomials in the entries of the matrix $g \in GL_n$. As an algebraic variety, GL_n is by definition the open set in the affine n^2 dimensional space of $n \times n$ matrices defined by the non-vanishing of the polynomial $\det g$. Therefore the globally defined *regular functions* on GL_n are generated by the polynomials in the entries of g and the multiplicative inverse of the determinant $(\det g)^{-1}$. If the entries of $\rho(g)$ are regular functions of the entries of g, we say that ρ is a *rational representation*.

Note that $\det \colon g \mapsto \det(g)$ is itself a 1-dimensional matrix representation of GL_n, and it follows that if ρ is a rational representation so is $(\det)^k \otimes \rho$, where the latter is defined by $g \mapsto (\det g)^k \rho(g)$, for any integer k. Since the regular functions on g have denominator a power of $\det g$ it follows that for every rational representation ρ, the representation $(\det)^k \otimes \rho$ is polynomial, for k sufficiently large. In other words, every rational representation has the form $(\det)^k \otimes \rho$ for some polynomial representation ρ and integer k. Thus the theories of polynomial representations and of rational representations are essentially interchangeable.

Example: The symmetric group S_n acts on $V = \mathbb{C}^n$ by permuting the basis vectors $e_1, \ldots, e_n$. The vector $v = e_1 + \cdots + e_n$ is invariant, so its span $W = \mathbb{C}v$ is a 1-dimensional submodule of V, on which S_n acts by the *trivial* representation. If we set $W' = \{\sum_i a_i e_i : a_1 + \cdots + a_n = 0\}$, then W' is also a submodule, spanned by the vectors $e_i - e_{i+1}$, and

$V = W \oplus W'$. This exhibits V as a direct sum of two irreducible submodules. *Exercise:* prove W' is irreducible.

Example: Let $B \subseteq GL_2$ be the subgroup of upper triangular 2×2 matrics. We have the *defining representation* of B on $V = \mathbb{C}^2$, in which each matrix in B is represented by itself. Since

$$\begin{bmatrix} x_1 & y \\ 0 & x_2 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \end{bmatrix} = \begin{bmatrix} x_1 \\ 0 \end{bmatrix},$$

the subspace W spanned by the basis vector $e_1 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$ is a submodule, on which B acts through the homomorphism sending $\begin{bmatrix} x_1 & y \\ 0 & x_2 \end{bmatrix}$ to $[x_1]$. There is a quotient module V/W, on which B acts via the homomorphism sending $\begin{bmatrix} x_1 & y \\ 0 & x_2 \end{bmatrix}$ to $[x_2]$, but there is no submodule W' such that $V = W \oplus W'$, so V is not completely reducible.

Theorem: Every complex representation of a finite group G is completely reducible.

This theorem also holds over any field whose characteristic does not divide the order of G. *Exercise:* find an example of a non-completely reducible representation of a finite group G over a field of prime characteristic dividing $|G|$.

Theorem: Every rational representation of $G = GL_n$ is completely reducible.

This theorem holds more generally for rational representations of semi-simple algebraic groups and for continuous representations of semi-simple Lie groups. For more on this, consult a textbook on Lie groups and Lie algebras.

Example: The defining representation of GL_n on $V = \mathbb{C}^n$ is obviously a polynomial representation. It is irreducible since for any non-zero vector v, we can find group elements g_i such that the vectors $g_i v$ span V.

Example: As we have seen, the homomorphism sending $g \in GL_n$ to $[\det g]$ defines a 1-dimensional matrix representation of GL_n, the *determinant* representation. Clearly it is a polynomial representation, and irreducible because 1-dimensional.

Example: $V \otimes V$ is a polynomial representation. A basis of $V \otimes V$ is given by the elements $e_i \otimes e_j$, where $e_1, \ldots, e_n$ is the usual basis of V. Then $g(e_i \otimes e_j) = ge_i \otimes ge_j = \sum_{i',j'} g_{i'i} g_{j'j} e_{i'} \otimes e_{j'}$, which shows that the entries of the matrix representing g on $V \otimes V$ are products $g_{i'i} g_{j'j}$, so this is a polynomial representation. *Exercise:* show that $V \otimes V = \bigwedge^2 V \oplus S^2 V$ is a decomposition of $V \oplus V$ into irreducibles, where $\bigwedge^2 V$ and $S^2 V$ are the second exterior and symmetric powers of V respectively (this will be easy later on).

Example: GL_n acts on the dual space V^* via the homorphism sending g to $(g^{-1})^T$. This is a rational but not a polynomial representation. The representation $(\det V) \otimes V^*$ is polynomial, since the entries of $(\det g)(g^{-1})^T$ are polynomials in the entries of g.

2.2. Characters.

If ρ and ρ' are similar matrix representations of G, $\rho'(g) = A^{-1} \rho(g) A$, for some A, then clearly we have $\operatorname{tr} \rho'(g) = \operatorname{tr} \rho(g)$. In particular, given a linear representation $\rho \colon G \to GL(V)$, the trace of $\rho(g)$ in a corresponding matrix representation is independent of choice of basis, and equal to its value in any isomorphic representation W.

Definition: The function $\chi^V : G \to \mathbb{C}$ sending g to the trace of its representing matrix is the *character* of the representation V.

Proposition: Characters are *class functions*, that is, they are constant on conjugacy classes of G.

Proof: If $g' = h^{-1}gh$ then $\rho(g') = \rho(h)^{-1}\rho(g)\rho(h)$, so the matrices $\rho(g')$ and $\rho(g)$ have the same trace.

The following two properties of characters are easy to verify by analyzing the structure of the matrix of a direct sum or tensor product of two linear maps.

Proposition: We have

$$\chi^{V\oplus W} = \chi^V + \chi^W$$

(more generally, $\chi^V = \chi^{V/W} + \chi^W$ for any submodule $W \subseteq V$, even if not a direct summand), and

$$\chi^{V\otimes W} = \chi^V \chi^W.$$

For a finite group G, the *group algebra* $\mathbb{C}G$ is the vector space of formal linear combinations of elements of G (with basis G), and multiplication defined by extending the multiplication in G linearly. Then G acts on $\mathbb{C}G$ by left multiplication, yielding a representation of G, called the *regular representation.* The basic facts about finite group representations are summarized in the following theorem.

Theorem: Let G be a finite group. Then every irreducible representation V_α of G occurs as a submodule of the regular representation $\mathbb{C}G$, and in fact $\mathbb{C}G$ is the direct sum over all α of $\chi^\alpha(1) = \dim V_\alpha$ copies of V_α. In particular there are only finitely many isomorphism classes of irreducible G modules. Moreover their characters form a basis of the space of class functions, so the number of non-isomorphic irreducibles is equal to the number of conjugacy classes in G.

The character of a representation of an infinite group G is well-defined, in particular, we can speak of the character of a polynomial representation of GL_n. Note that the character is itself a polynomial in the entries of g. The diagonalizable elements of GL_n form a dense set, so χ^V is determined by its values on diagonalizable elements g, and hence, since it is a class function, by its values on the diagonal elements. Denoting by $\tau(\mathbf{x})$ the diagonal matrix with diagonal entries $x_1, \ldots, x_n$, $\chi^V(\tau(\mathbf{x}))$ is a polynomial in $\mathbf{x}$. We will denote it simply by $\chi^V(\mathbf{x})$ and refer to this polynomial as the character of V.

Note that for any permutation $w \in S_n$, $\tau(x_{w(1)}, \ldots, x_{w(n)})$ is conjugate to $\tau(\mathbf{x})$ by a permutation matrix in GL_n. This implies that $\chi^V(\mathbf{x})$ is a *symmetric* polynomial in $\mathbf{x}$.

Note also that the character of $\det g$ is $x_1 \cdots x_n = e_n(\mathbf{x})$. The character of the rational representation $(\det g)^k$ is $(x_1 \cdots x_n)^k$, where k may be any integer. Thus the characters of the rational representations of GL_n are rational functions of the form $(x_1 \cdots x_n)^k f(\mathbf{x})$, where f is a symmetric polynomial. In other words they are symmetric *Laurent polynomials* in $\mathbf{x}$. These characters make sense literally, not just formally, since the x_i's are the eigenvalues of g, and these are non-zero as g is invertible.

2.3. Schur's Lemma.

Lemma: If V and W are irreducible G-modules, and $\phi\colon V \to W$ is a G-module homorphism (a linear map that commutes with the G actions), then either $\phi = 0$ or ϕ is an isomorphism.

Proof: Assume $\phi \neq 0$. Then the kernel of ϕ is a submodule of V, not equal to V, and since V is irreducible, $\ker(\phi) = 0$. Similarly the image of ϕ is a submodule of W, not equal to 0, and since W is irreducible, $\mathrm{im}(\phi) = W$. This shows ϕ is injective and surjective, hence an isomorphism.

Corollary: If V is irreducible then the space of G-module homomorphisms $\mathrm{Hom}_G(V, V)$ is 1-dimensional, *i.e.*, every $\phi\colon V \to V$ is a scalar multiple of the identity map.

Proof: Given $\phi\colon V \to V$ let α be an eigenvalue of ϕ. Then $\phi - \alpha 1_V$ is singular, hence not an isomorphism, hence zero by Schur's Lemma. Note the corollary uses the fact that $\mathbb{C}$ is algebraically closed, so that ϕ must have an eigenvalue $\alpha \in \mathbb{C}$.

2.4. Polynomial representations of GL_n.

For the moment we fix $G = GL_n$. We denote by B the *Borel subgroup* of upper triangular matrices, by N the *unipotent radical* of B consisting of upper triangular matrices with 1's on the diagonal, and by T the *torus* of diagonal matrices in G. Note that $T \cong (\mathbb{C}^*)^n$ is indeed an algebraic torus group. We also write B^- and N^- for the lower-triangular analogs of B and N, the *opposite* Borel subgroup and its unipotent radical.

The elements of T will be denoted $\tau(x_1, \dots, x_n)$, as in 2.2. Thus our convention for representing a character of g as a function $\chi(x_1, \dots, x_n)$ amounts to evaluating χ on the torus T.

The torus group T is itself semi-simple (it's a product of copies of GL_1), so any G module V is completely reducible as a T module. One can show that the polynomial representations of T are simply the one-dimensional representations $\mathbb{C}_\lambda$, in which $\tau(\mathbf{x})$ acts as multiplication by $\mathbf{x}^\lambda$. Here λ is a sequence of non-negative integers $(\lambda_1, \dots, \lambda_n)$. The rational representation are the same except that the λ_i are any integers, possibly negative. Complete reducibility as a T module then means that V decomposes as a direct sum of *weight spaces* V_λ, whose elements $v \in V_\lambda$ are simultaneous eigenvectors for all $\tau(\mathbf{x}) \in T$, with eigenvalue $\mathbf{x}^\lambda$. Weights for which λ is a partition are called *dominant*.

Proposition: The character $\chi^V(\mathbf{x})$ is equal to

$$\chi^V(\mathbf{x}) = \sum_\lambda (\dim V_\lambda)\mathbf{x}^\lambda.$$

In particular the character of any representation of GL_n has non-negative integer coefficients.

Proof: The trace of $\tau(\mathbf{x})$ on V is the sum of its eigenvalues, counted with multiplicities. Since the weight spaces are by definition the eigenspaces, χ^V contains the term $\mathbf{x}^\lambda$ with coefficient equal to the dimension of the corresponding weight space V_λ.

Note that the Proposition implies that when $\lambda' = w(\lambda)$ is a permutation of λ, we have $\dim V_{\lambda'} = \dim V_\lambda$. This can also be seen directly, as the permutation matrix of w carries V_λ into $V_{\lambda'}$.

Since GL_n is an open subset of the affine space of all $n \times n$ matrices $M_n(\mathbb{C})$, its tangent space at 1 is just that of $M_n(\mathbb{C})$, and can itself be identified with $M_n(\mathbb{C})$. In this context it is usually denoted $\mathfrak{gl}_n$ rather than $M_n(\mathbb{C})$. The group structure of GL_n is reflected in a *Lie algebra* structure on $\mathfrak{gl}_n$, which will not concern us for now.

Given a polynomial representation

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

since $\rho(1) = 1_V$, the differential of ρ is a linear map of tangent spaces

$$d\rho\colon \mathfrak{gl}_n \to \mathfrak{gl}(V).$$

We can compute it explicitly from its definition

$$d\rho(a) = \lim_{\epsilon\to 0} \frac{1}{\epsilon}(\rho(1+\epsilon a) - \rho(1)),$$

or from either of the equivalent and more convenient formulas

$$d\rho(a) = \frac{d}{dt}\rho(1+ta)|_{t=0} = \frac{d}{dt}\rho\exp(ta)|_{t=0}.$$

Lemma (A):

$$\exp(t\ d\rho(a)) = \rho(\exp(ta)).$$

Proof: Let $g_t = \exp(ta)$. Note that $g_t g_{t'} = g_{t+t'}$. Therefore,

$$\frac{d}{dt}\rho(g_t)|_{t=t_0} = \frac{d}{dt}\rho(g_{t_0})\rho(g_{t-t_0})|_{t=t_0} = \rho(g_{t_0})\frac{d}{dt}\rho(g_t)|_{t=0} = \rho(g_{t_0})d\rho(a).$$

This shows that $\rho(g_t)$ satisfies the first-order differential equation and initial condition

$$\frac{d}{dt}\rho(g_t) = \rho(g_t)d\rho(a); \quad \rho(g_0) = 1,$$

which characterize $\exp(t\ d\rho(a))$.

Lemma (B): If $v \in V_\lambda$ is a weight vector, then $d\rho(E_{ij})v \in V_{R_{ij}\lambda}$ is a weight vector of weight $R_{ij}\lambda$, where E_{ij} is the unit matrix with a 1 in row i and column j, and zeroes elsewhere.

Proof: We have

$$\begin{aligned}\rho(\tau(\mathbf{x}))d\rho(E_{ij})v &= \frac{d}{dt}\rho(\tau(\mathbf{x})\exp(tE_{ij}))v|_{t=0}\\ &= \frac{d}{dt}\rho(\exp(t\frac{x_i}{x_j}E_{ij})\tau(\mathbf{x}))v|_{t=0},\end{aligned}$$

since $\tau(\mathbf{x})E_{ij}\tau(\mathbf{x})^{-1} = (x_i/x_j)E_{ij}$. Since v has weight λ, the above is equal to

$$\frac{d}{dt}\rho(\exp(t\frac{x_i}{x_j}E_{ij}))\mathbf{x}^\lambda v|_{t=0} = \mathbf{x}^{R_{ij}\lambda}\frac{d}{dt}\rho(\exp(tE_{ij}))v|_{t=0} = \mathbf{x}^{R_{ij}\lambda}d\rho(E_{ij})v.$$

Corollary: If λ is a partition, maximal in dominance order among the weights of V, then every weight vector $v \in V_\lambda$ is N-invariant, *i.e.*, $V_\lambda \subseteq V^N$.

Proof: Matrices of the form $\exp(tE_{ij})$ generate N. By Lemma 2.4 (A), we have $\rho(\exp(tE_{ij}))v = \exp(td\rho(E_{ij}))v = v$, since $d\rho(E_{ij})v = 0$, as the weight space $V_{R_{ij}\lambda}$ is zero by maximality.

Let us remark that since T normalizes N, the space of N invariants V^N is a T submodule of V, and hence is itself the direct sum of its weight spaces V^N_λ.

Corollary: Every non-zero representation of G has $V^N \neq 0$. Hence $\dim V^N = 1$ implies V is irreducible.

Proof: The first part is clear, since there is some maximal non-zero weight space. If $V = W \oplus W'$ were a direct sum of non-zero submodules, then each would contribute to V^N, giving $\dim V^N > 1$.

2.5. The flag variety.

Definition: A (complete) *flag* F in $V = \mathbb{C}^n$ is a sequence of subspaces $0 \subseteq F_1 \subseteq F_2 \subseteq \cdots \subseteq F_{n-1} \subseteq V$, with $\dim F_i = i$.

Recall that the i-dimensional subspaces $W \subseteq V$ form a projective variety, the *Grassmann variety* $G_i(V)$. (Sometimes we will denote it instead by $G^{n-i}(V)$, when we want to view it as parametrizing the $n-i$ dimensional quotient spaces V/W of V.)

For $k < l$, the locus in $G_k \times G_l$ consisting of pairs (W_k, W_l) such that $W_k \subseteq W_l$ is a closed subvariety, *i.e.*, containment is a *closed condition*. It follows that the complete flags in V can be represented as a closed subvariety X of the product of Grassmann varieties,

$$X \subseteq G_1(V) \times \cdots \times G_{n-1}(V).$$

Now $G = GL(V)$ acts (algebraically) on X, with $g \in G$ mapping the flag F to the flag gF defined by $(gF)_i = g(F_i)$. Consider the *standard flag* E:

$$E_i = \{e_1, \dots, e_i\},$$

where $e_1, \dots, e_n$ is the usual basis of unit coordinate vectors in $V = \mathbb{C}^n$. We have $gE_i = E_i$ if and only if the first i columns of the matrix g have non-zero entries only in the first i rows, that is, iff g is block upper-triangular, with diagonal blocks of sizes $i \times i$ and $(n-i)\times(n-i)$. We have $gE = E$ iff this holds for all i, that is, the matrix g is upper-triangular. This shows that the stabilizer of E is B.

Given any flag F, we can choose a compatible basis $f_1, \dots, f_n$ such that $\{f_1, \dots, f_i\}$ spans F_i for all i, by simply taking f_i to be any element of F_i not in F_{i-1}. Then the element $g \in G$ carring e_i to f_i for all i has $gE = F$. Thus the G orbit of the standard flag E is the whole flag variety X. In particular, X is irreducible, since it is the image of the irreducible variety G by the morphism mapping g to gE.

Since G acts transitively on X and B is the stabilizer of the flag $E \in X$, we can identify X with the space of cosets G/B, where a coset gB corresponds to the flag gE, a correspondence that does not depend on the choice of coset representative. In this way G/B is given the structure of a projective variety.

(For any closed subgroup B of an algebraic group G, there is a canonical way to make G/B into an algebraic variety, but usually it will not be a projective variety. Note that G itself here is an affine, not a projective variety.)

We will have something to say later about the construction of homogeneous coordinates for G/B. For now, we work out local coordinates in a neighborhood of the standard flag E, or the coset $1B$. Recall the following result of matrix algebra.

Proposition: Every $n \times n$ matrix can be written in the form $LPDU$, where $U \in N$ is upper uni-triangular, $L \in N^-$ is lower uni-triangular, P is a permutation matrix, and D is a diagonal matrix.

In the language of G/B this says that every coset gB is contained in $N^- wB$ for some permutation w, and thus every flag belongs to $N^- wE$ for some w. Consider the morphism

$$\phi \colon N^- \to G/B; \quad \phi(g) = gE.$$

This morphism is obviously equivariant with respect to left multiplication by elements of N^-; in particular its fibers over all points of its image are isomorphic. We have $\phi^{-1}(\{E\}) = N^- \cap B = \{1\}$, so ϕ is injective, and therefore, since G/B is non-singular, ϕ is an isomorphism onto its image $N^- E$. Thus $N^- E$ is an affine cell of dimension $\dim N^- = \binom{n}{2}$.

To make this explicit, observe that for $n \in N^-$, nE is the flag F with F_i spanned by the vectors $ne_1, \dots, ne_i$, that is, by the first i columns of the matrix

$$n = \begin{bmatrix} 1 & & \dots & 0 \\ a_{21} & 1 & & \vdots \\ a_{31} & a_{32} & 1 & \\ & & \dots & \\ a_{n1} & a_{n2} & \dots & 1 \end{bmatrix}.$$

The entries of this matrix thus serve as local coordinates on $N^- E$. Using this description, moreover, one can show (exercise) that a flag F belongs to $N^- E$ if and only if, for all i, with $W_i = \mathbb{C}\{e_{i+1}, \dots, e_n\}$, we have $F_i + W_i = \mathbb{C}^n$, or equivalently, $F_i \cap W_i = 0$. Now for any fixed space W, the condition $F_i \cap W \neq 0$ is a closed condition on the Grassmannian $G_i(\mathbb{C}^n)$, so our conditions on F are open conditions, and $N^- E$ is an *open* affine cell in the flag variety.

By symmetry, for any $g \in G$, the set $gN^- E$ is an open affine neighborhood of gE, so this provides us with a covering of the flag variety by open affines, which are actually affine cells.

2.6. The Borel-Weil construction.

Let W be a rational representation of the Borel subgroup $B \subseteq G = GL_n$. We are going to describe a canonical procedure for constructing an algebraic vector bundle $G \times_B W$ on G/B, with an equivariant G action, whose fiber at the standard flag E (corresponding to the coset $1B \in G/B$) is W.

First we review terminology. A *vector bundle* of rank r over a space X is a space E, together with a map $\pi \colon E \to X$, and a structure of r-dimensional vector space on each fiber of π. In particular, $\mathbb{C}^r \times X$, with π the projection on X, is a vector bundle, called the *trivial* bundle over X. In general we require every vector bundle to be *locally trivial.* This means that every point of X has an open neighborhood U such that $\pi \colon \pi^{-1}(U) \to U$ is isomorphic to a trivial bundle. Here *isomorphism* has the meaning you would expect: $\pi \colon E \to X$ and $\pi' \colon E' \to X$ are isomorphic if there is a map $\alpha \colon E \to E'$ such that $\pi' \circ \alpha = \pi$, and α induces a linear isomorphism on each fiber.

The meaning of the term "vector bundle" depends on the category in which we work. Thus if we require that E should be a topological space, and that the map π and the trivializing isomorphisms should be continuous, we have defined *topological* vector bundles. For our purposes we will work with *algebraic* vector bundles. Thus X should be a complex algebraic variety, as should E, and we regard $\mathbb{C}^r$ as affine r space. The map π and the trivializing isomorphisms should be *morphisms*, that is, they have to be given by regular functions in local coordinates, and of course the trivializing isomorphisms should be isomorphisms of algebraic varieties.

I should point out that the concept of algebraic vector bundle is distinct from and more rigid than that of topological vector bundle. There exist topological vector bundles which cannot be given an algebraic structure, and topological isomorphisms between algebraically non-isomorphic vector bundles.

Definition: Given a B-module W, the space $G\times_B W$ is the orbit space $(G\times W)/B$, where B acts on $G\times W$ by $b(g,w) = (gb^{-1}, bw)$. There is a canonical map $\pi\colon G\times_B W \to G/B$ given by $(g,w)\mapsto gB$.

Proposition: For any $g\in G$, the map $\phi_g\colon W\to G\times_B W$ sending w to $B(g,w)$ is a bijection from W onto the fiber of π over the point $gB\in G/B$. Moreover, for any other $g' = gb\in gB$, the map $\phi_g^{-1}\circ\phi_{g'}\colon W\to W$ is $\rho(b)$, where ρ is the given representation of B on W.

Proof: Since $\pi(g,w) = gB$, ϕ_g does map W to the fiber over gB, and indeed surjectively onto it, since w is arbitrary. Now let $g' = gb$. I claim that for all w, $w' = b^{-1}w$ is the unique element such that $\phi_{g'}(w') = \phi_g(w)$. This establishes both that ϕ_g is injective (take $b = 1$) and the "moreover" part of the proposition.

For the claim, suppose $B(g',w') = B(g,w)$. The element b such that $g' = gb$ is unique, so the only representation of (g',w') as element of $B(g,w)$ is as $(gb, b^{-1}w)$. This says exactly that $w' = b^{-1}w$.

Now we assign the fibers of $\pi\colon G\times_B W\to G/B$ vector space structures via their bijections ϕ_g with W. By the Proposition, since $\phi_g^{-1}\circ\phi_{g'}$ is linear, this assignment is independent of the choice of g.

We want to verify this structure makes $G\times_B W$ into an algebraic vector bundle, that is, we have algebraic local trivializations. Since everything is G-equivariant it is enough to do this on the open set $N^-E\in G/B$. The space $\pi^{-1}(N^-E)$ consists of the orbits $B(n,w)$, where $n\in N^-$ and $w\in W$. Since ϕ is injective, the pairs (n,w) are a system of orbit representatives, *i.e.*, their orbits are distinct. This gives an isomorphism of $\phi^{-1}(N^-E)$ with $N^-\times W\cong N^-E\times W$, which preserves the linear structure we assigned to the factors. This shows that $G\times_B W$ is an algebraic vector bundle, and that it trivializes over N^-E (and hence also over gN^-E for every $g\in G$).

The Borel-Weil construction actually sets up an equivalence between B-modules and G-equivariant vector bundles over G/B, as given by the following proposition, whose proof I leave as an exercise.

Proposition: Let A be a G-equivariant algebraic vector bundle over G/B, that is, a vector bundle with an algebraic action of G that commutes with the map $\pi\colon A\to G/B$, and such that each g carries the fiber over a flag F linearly onto the fiber over gF. Then $A\cong G\times_B W$, where W is the fiber of A over the B-fixed point $E = 1B$, viewed as a B-module.

Exercise: Show that if W is originally a G-module, with the G action restricted to B, then $G\times_B W$ is isomorphic to the trivial bundle $(G/B)\times W$, with the obvious equivariant G action on the latter.

To avoid confusion we'll use a different notation for the algebraic vector bundle $G \times_B W$ itself (viewed as an algebraic variety), and its sheaf of sections, a locally free sheaf of $\mathcal{O}$-modules on G/B. The latter we denote by $E(W)$.

Proposition: The space of global sections $H^0(G/B, E(W))$ is a rational representation of G.

Proof: It's finite-dimensional because G/B is projective. Since $E(W)$ is an algebraic vector bundle, G acts algebraically on its space of sections, so it's a rational representation.

Let us remark that the higher sheaf cohomology groups of $E(W)$ are also rational G-modules, and that H^0 can perfectly well be zero, even if W is not.

2.7. Line bundles.

A rank-1 vector bundle is called a *line bundle.* By the correspondence in 2.7, the equivariant line bundles on G/B correspond to 1-dimensional representations of B.

Proposition: Every 1-dimensional rational representation of B is of the form $\mathbb{C}_\lambda$, where λ is a weight, and B acts on $\mathbb{C}_\lambda$ through the homomorphism $B \to T$ with kernel N, mapping each upper-triangular matrix to its diagonal part, and T acts by $\tau(\mathbf{x})v = \mathbf{x}^\lambda v$.

Proof: If W is 1-dimensional then it has only one weight space $W = W_\lambda$. Hence N acts trivially on W, by the corollary to Lemma 2.4(B), while T acts according to the weight, by definition. As B is the semi-direct product of T and N, the result follows.

For reasons that will appear shortly, we denote by L_λ the line bundle

$$L_\lambda = G \times_B \mathbb{C}_{w_0\lambda},$$

where $w_0 \in S_n$ is the permutation reversing the indices: $w_0(i) = n + 1 - i$. The sheaf of sections of L_λ is the locally free sheaf

$$E(\mathbb{C}_{w_0\lambda}).$$

Now let us recall how the Grassmann variety $G_k(V)$ is embedded in projective space. A point of $G_k(V)$ is a k-dimensional subspace $W \subseteq V$, and the k-th exterior power of the latter, $\bigwedge^k W$, is a 1-dimensional subspace of $\bigwedge^k V$. By definition the projective space $\mathbb{P}(\bigwedge^k V)$ is the variety $G_1(\bigwedge^k V)$ of such 1-dimensional subspaces, so the map $W \mapsto \bigwedge^k W$ defines a map $G_k(V) \mapsto \mathbb{P}(\bigwedge^k V)$. This is the *Plücker embedding.*

In concrete terms the Plücker embedding assigns W the homogeneous coordinates given by the $k \times k$ minors of any $k \times n$ matrix with row space W (the ratios of these minors being independent of the choice of basis for W). To see this note that, if $e_1, \ldots, e_n$ is a basis of V, and $v_1, \ldots, v_k$ is a basis of W, then the coordinates of $v_1 \wedge \cdots \wedge v_k$ in terms of the basis $\{e_{i_1} \wedge \cdots \wedge e_{i_k} : i_1 < \cdots < i_k\}$ are given precisely by the minors of the matrix with rows v_i.

Recall that the standard ample sheaf $\mathcal{O}(1)$ on projective space $\mathbb{P}(U)$ is *dual* to the sheaf of sections of the tautological line bundle L, whose fiber over a 1-dimensional subspace $S \in \mathbb{P}(U)$, $S \subseteq U$, is S itself. Each linear coordinate function x on U gives rise by restriction to a linear functional on each fiber S, and so to a global section of the dual bundle to L, or a global section of $\mathcal{O}(1)$. By multiplication, polynomials homogeneous of degree d in the coordinates on U give rise to global sections of the d-th tensor power $\mathcal{O}(d)$ of $\mathcal{O}(1)$. This gives the ring homomorphism, which is actually an isomorphism, identifying the coordinate ring of U, which is also the homogeneous coordinate ring of $\mathbb{P}(U)$, with the graded ring of global sections

$$\bigoplus_{d\geq 0} H^0(\mathbb{P}(U), \mathcal{O}(d)).$$

In particular, the sheaves $\mathcal{O}(d)$ have plenty of non-zero global sections—enough that their vanishing loci can distinguish between any distinct closed subschemes of $\mathbb{P}(U)$, whence the term "ample."

Restricting $\mathcal{O}(1)$ from $\mathbb{P}(\bigwedge^k(V))$ to $G_k(V)$, as embedded by the Plücker embedding, it becomes the dual to the sheaf of sections of the highest exterior power $\bigwedge^k W$ of the tautological bundle W on $G_k(V)$, whose fiber at a point W is W itself. We can summarize this as follows:

Proposition: The ample bundle $\mathcal{O}(1)$ induced on $G_k(V)$ by the Plücker embedding is the dual of the k-th exterior power of the tautological bundle.

For our purposes, it will be more natural to formulate this in terms of $G^d(V) = G_k(V)$, where $d = n - k$. The relevant bundle on $G^d(V)$ is the tautological *quotient bundle* with fiber V/W at W, a rank d vector bundle. Exterior multiplication $\bigwedge^k W \otimes \bigwedge^d(V/W) \to \bigwedge^n V \cong \mathbb{C}$ identifies $\bigwedge^d(V/W)$ with the dual space of $\bigwedge^k W$. This identification is canonical except for the choice of a non-zero basis vector in $\bigwedge^n V$, which does not depend on W. Thus $\bigwedge^d(V/W)$ is isomorphic to the dual bundle of $\bigwedge^k W$.

Corollary: The ample bundle $\mathcal{O}(1)$ induced on $G^d(V)$ by the Plücker embedding is isomorphic to the d-th exterior power of the tautological quotient bundle.

Corollary: Let Q_d be the tautological quotient bundle on the flag variety G/B whose fiber at a flag F is V/F_{n-d}. Then its highest exterior power $M_d = \bigwedge^d Q_d$ is isomorphic to $\mathcal{O}(1)$, pulled back from the Grassmann variety G^d, through the embedding of G/B in the product of Grassmann varieties. In particular, M_d has non-zero global sections, and any tensor product $M_1^{e_1} \otimes \cdots \otimes M_{n-1}^{e_{n-1}}$ with all $e_d > 0$ is very ample on G/B, *i.e.*, it is the restriction of $\mathcal{O}(1)$ for some projective embedding of G/B.

Note that we could have also constructed M_d as the highest exterior power of the dual of the tautological subspace bundle S_{n-d} with fiber F_{n-d} at F. This gives the same line bundle with a different G action, because G does not act trivially on $\bigwedge^n V$. *Exercise:* show that $\bigwedge^{n-d}(S_{n-d})^*$ is equivariantly isomorphic to $(\det)^{-1} \otimes M_d$, where det is the trivial line bundle $G/B \times \mathbb{C}$, with G acting by the determinant representation on $\mathbb{C}$, or $\det = L_{(-1,-1,\dots,-1)}$ in our notation for equivariant line bundles.

As Q_d is clearly G-equivariant, it must be $L_\lambda = G \times_B \mathbb{C}_{w_0\lambda}$ for some weight λ. To find this weight, we need only inspect the B action (or just the T action) on the fiber of Q_d at the standard flag, corresponding to the coset $1B$. This fiber is $\bigwedge^d(V/E_{n-d})$. It is spanned by the single vector $e_{n-d+1} \wedge \cdots \wedge e_n$, on which T acts with weight $w_0\lambda = (0, \dots, 0, 1, \dots, 1)$, with d ones and $n - d$ zeroes. Therefore we have the following result.

Proposition: The G-equivariant line bundle $\bigwedge^d Q_d$ on G/B is isomorphic to L_{λ_d} where $\lambda_d = \varepsilon_1 + \cdots + \varepsilon_d = (1, \dots, 1, 0, \dots, 0)$ corresponds to the partition (1^d).

We remark that in a homogeneous polynomial representation of degree d of G, the scalar matrices tI must act as scalars $t^d I$. Equivalently, every weight space V_λ has $|\lambda| = d$. Similar considerations apply to B and T. The defining representation V of G is clearly homogeneous of degree 1. The equivariant bundle Q_d corresponds to the B-module V/E_{n-d}—note that E_{n-d} is a B-submodule of V, so this makes sense. This is again homogeneous of degree 1, and its d-th exterior power is homogeneous of degree d. Thus we should expect, as we have just seen, that λ for the corresponding line bundle L_λ should be a partition of d.

Corollary: For every partition λ, the space of global sections $H^0(G/B, L_\lambda)$ is non-zero.
Proof: Every such L_λ is a tensor product $L_{\lambda_1}^{\mu_1} \otimes \cdots \otimes L_{\lambda_n}^{\mu_n}$, where in fact μ is the conjugate partition λ'. Since $L_{\lambda_d} = M_d$, it has a non-zero global section.

We have seen that $L_{(1,\ldots,1)}$ is the determinant representation of G on the trivial bundle. By tensoring with $L^k_{(1,\ldots,1)}$, for any integer k, we see that the Corollary also applies for λ of the form partition plus $(k, \ldots, k)$, that is, for λ any non-increasing integer sequence. Shifting λ by $(k, \ldots, k)$ has the effect of tensoring $H^0(G/B, L_\lambda)$ by the k-th power of the determinant representation.

2.8. Irreducibile representations.

Theorem: The representation of G on $H^0(G/B, L_\lambda)$ is irreducible whenever it is non-zero.
Proof: By Corollary 2.4, there is an N-invariant section in $H^0(G/B, L_\lambda)$. Equally well, by symmetry, there is an N^--invariant section σ. Then σ is determined on all of $N^- E$ by its value in the fiber over the standard flag E. Since $N^- E$ is open, it is dense, and thus σ is determined completely by its value in the fiber over E. In particular, since this fiber is 1-dimensional, any two N^--invariant sections in $H^0(G/B, L_\lambda)$ are scalar multiples of each other. Again by Corollary 2.4 (with N^- in place of N) this implies that $H^0(G/B, L_\lambda)$ is irreducible.

Now let V be an arbitrary irreducible rational representation of G and let λ be a maximal weight, that is, λ is a partition plus $(k, \ldots, k)$ for some integer k, and the partition in question is maximal in dominance order. Then as we have seen, V contains a non-zero N-invariant weight vector $v \in V_\lambda$. Let $W = \mathbb{C}v$ be the span of v. Then W is a 1-dimensional B-submodle of V, since N fixes v, and v is a weight vector for T. Hence we can construct the line bundle $L_{w_0\lambda}$ on G/B as $G \times_B W$.
There is a morphism ϕ of algebraic varieties from $G \times_B W$ to V given by

$$B(g, tv) \mapsto tgv.$$

This is well-defined since a general element of $B(g, v)$ has the form (gb^{-1}, bv). Its image is the union of $\{0\}$ and the G orbit of v in V.
The map ϕ is linear in the scalar t. Hence if $\alpha \in V^*$ is a linear functional on V, then the composite $\alpha \circ \phi$ is linear on each fiber of the line bundle $G \times_B W$, that is, it represents a global section of the dual bundle $L_{-w_0\lambda}$.
In particular, by taking a basis of weight vectors including v in V, and its dual basis (also of weight vectors) in V^*, we can choose α of weight $-\lambda$, with $\alpha(v) \neq 0$. Then obviously the section $\sigma = \alpha \circ \phi \in H^0(G/B, L_{-w_0\lambda})$ is non-zero. This shows that the map from V^* to $H^0(G/B, L_{-w_0\lambda})$ sending α to $\alpha \circ \phi$ is non-zero, and it is clearly linear and G-equivariant. Since V^* and $H^0(G/B, L_{-w_0\lambda})$ are both irreducible, this map is an isomorphism, by Schur's

Lemma. Furthermore, $-w_0\lambda$ is a maximal weight (indeed, the unique maximal weight, by irreducibility and Corollary 2.4) of V^*. Of course, since V was arbitrary, so is V^*, so we have proved the following fundamental theorem of Borel and Weil.

Theorem: Every irreducible rational representation V of GL_n is isomorphic to $H^0(G/B, L_\lambda)$, where λ is the unique maximal weight of V.

Corollary: The irreducible representations V^λ with highest weight λ are non-isomorphic for distinct λ.

Proof: Immediate from the uniqueness of λ.

Corollary: A rational G module V is irreducible if and only if $\dim V^N = 1$.

Proof: We have already seen the "if". For the only if, use the Borel-Weil theorem and the fact shown in the proof of the previous theorem, that $H^0(G/B, L_\lambda)$ has a unique N-invariant vector, up to a scalar multiple.