

Invariant Theory with Applications

Jan Draisma and Dion Gijswijt

October 8 2009

Contents

Chapter 1

Lecture 1. Introducing invariant theory

The first lecture gives some flavor of the theory of invariants. Basic notions such as (linear) group representation, the ring of regular functions on a vector space and the ring of invariant functions are defined, and some instructive examples are given.

1.1 Polynomial functions

Let V be a complex vector space. We denote by $V^* := \{f : V \rightarrow \mathbb{C} \text{ linear map}\}$ the dual vector space. Viewing the elements of V^* as functions on V , and taking the usual pointwise product of functions, we can consider the algebra of all $\mathbb{C}$ -linear combinations of products of elements from V^* .

Definition 1.1.1. The *coordinate ring* $\mathcal{O}(V)$ of the vectorspace V is the algebra of functions $F : V \rightarrow \mathbb{C}$ generated by the elements of V^* . The elements of $\mathcal{O}(V)$ are called *polynomial* or *regular* functions on V .

If we fix a basis $e_1, \dots, e_n$ of V , then a dual basis of V^* is given by the coordinate functions $x_1, \dots, x_n$ defined by $x_i(c_1e_1 + \dots + c_ne_n) := c_i$. For the coordinate ring we obtain $\mathcal{O}(V) = \mathbb{C}[x_1, \dots, x_n]$. This is a polynomial ring in the x_i , since our base field $\mathbb{C}$ is infinite.

Exercise 1.1.2. Show that indeed $\mathbb{C}[x_1, \dots, x_n]$ is a polynomial ring. In other words, show that the x_i are algebraically independent over $\mathbb{C}$: there is no nonzero polynomial $p \in \mathbb{C}[X_1, \dots, X_n]$ in n variables $X_1, \dots, X_n$, such that $p(x_1, \dots, x_n) = 0$. Hint: this is easy for the case $n = 1$. Now use induction on n .

We call a regular function $f \in \mathcal{O}(V)$ *homogeneous of degree d* if $f(tv) = t^d f(v)$ for all $v \in V$ and $t \in \mathbb{C}$. Clearly, the elements of V^* are regular of degree

1, and the product of polynomials f, g homogeneous of degrees d, d' yields a homogeneous polynomial of degree $d + d'$. It follows that every regular function f can be written as a sum $f = c_0 + c_1 f_1 + \cdots + c_k f_k$ of regular functions f_i homogeneous of degree i . This decomposition is unique (disregarding the terms with zero coefficient). Hence we have a direct sum decomposition $\mathcal{O}(V) = \bigoplus_{d \in \mathbb{N}} \mathcal{O}(V)_d$, where $\mathcal{O}(V)_d := \{f \in \mathcal{O}(V) \mid f \text{ homogeneous of degree } d\}$, making $\mathcal{O}(V)$ into a *graded algebra*.

Exercise 1.1.3. Show that indeed the decomposition of a regular function f into its homogeneous parts is unique.

In terms of the basis $x_1, \dots, x_n$, we have $\mathcal{O}(V)_d = \mathbb{C}[x_1, \dots, x_n]_d$, where $\mathbb{C}[x_1, \dots, x_n]_d$ consists of all polynomials of total degree d and has as basis the monomials $x_1^{d_1} x_2^{d_2} \cdots x_n^{d_n}$ for $d_1 + d_2 + \cdots + d_n = d$.

1.2 Representations

Central objects in this course are linear representations of groups. For any vector space V we write $\mathrm{GL}(V)$ for the group of all invertible linear maps from V to itself. When we have a fixed basis of V , we may identify V with $\mathbb{C}^n$ and $\mathrm{GL}(V)$ with the set of invertible matrices $n \times n$ matrices $\mathrm{GL}(\mathbb{C}^n) \subset \mathrm{Mat}_n(\mathbb{C})$.

Definition 1.2.1. Let G be a group and let X be a set. An *action of G on X* is a map $\alpha : G \times X \rightarrow X$ such that $\alpha(1, x) = x$ and $\alpha(g, \alpha(h, x)) = \alpha(gh, x)$ for all $g, h \in G$ and $x \in X$.

If α is clear from the context, we will usually write gx instead of $\alpha(g, x)$. What we have just defined is sometimes called a *left action* of G on X ; *right actions* are defined similarly.

Definition 1.2.2. If G acts on two sets X and Y , then a map $\phi : X \rightarrow Y$ is called *G -equivariant* if $\phi(gx) = g\phi(x)$ for all $x \in X$ and $g \in G$. As a particular case of this, if X is a subset of Y satisfying $gx \in X$ for all $x \in X$ and $g \in G$, then X is called *G -stable*, and the inclusion map is G -equivariant.

Example 1.2.3. The symmetric group S_4 acts on the set $\binom{[4]}{2}$ of unordered pairs of distinct numbers in $[4] := \{1, 2, 3, 4\}$ by $g\{i, j\} = \{g(i), g(j)\}$. Think of the edges in a tetrahedron to visualise this action. The group S_4 also acts on the set $X := \{(i, j) \mid i, j \in [4] \text{ distinct}\}$ of all ordered pairs by $g(i, j) = (g(i), g(j))$ —think of directed edges—and the map $X \rightarrow \binom{[4]}{2}$ sending (i, j) to $\{i, j\}$ is S_4 -equivariant.

Definition 1.2.4. Let G be a group and let V be a vector space. A *(linear) representation of G on V* is a group homomorphism $\rho : G \rightarrow \mathrm{GL}(V)$.

If ρ is a representation of G , then the map $(g, v) \mapsto \rho(g)v$ is an action of G on V . Conversely, if we have an action α of G on V such that $\alpha(g, \cdot) : V \rightarrow V$ is a linear map for all $g \in G$, then the map $g \mapsto \alpha(g, \cdot)$ is a linear representation.

As with actions, instead of $\rho(g)v$ we will often write gv . A vector space with an action of G by linear maps is also called a G -module.

Given a linear representation $\rho : G \rightarrow \mathrm{GL}(V)$, we obtain a linear representation $\rho^* : G \rightarrow \mathrm{GL}(V^*)$ on the dual space V^* , called the *dual representation* or *contragredient representation* and defined by

$$(\rho^*(g)x)(v) := x(\rho(g)^{-1}v) \text{ for all } g \in G, x \in V^* \text{ and } v \in V. \quad (1.1)$$

Exercise 1.2.5. Let $\rho : G \rightarrow \mathrm{GL}_n(\mathbb{C})$ be a representation of G on $\mathbb{C}^n$. Show that with respect to the dual basis, ρ^* is given by $\rho^*(g) = (\rho(g)^{-1})^\top$, where $A^\top$ denotes the transpose of the matrix A .

1.3 Invariant functions

Definition 1.3.1. Given a representation of a group G on a vector space V , a regular function $f \in \mathcal{O}(V)$ is called G -invariant or simply *invariant* if $f(v) = f(gv)$ for all $g \in G, v \in V$. We denote by $\mathcal{O}(V)^G \subseteq \mathcal{O}(V)$ the subalgebra of invariant functions. The actual representation of G is assumed to be clear from the context.

Observe that $f \in \mathcal{O}(V)$ is invariant, precisely when it is constant on the orbits of V under the action of G . In particular, the constant functions are invariant.

The representation of G on V induces an action on the (regular) functions on V by defining $(gf)(v) := f(g^{-1}v)$ for all $g \in G, v \in V$. This way the invariant ring can be described as the set of regular functions fixed by the action of G : $\mathcal{O}(V)^G = \{f \in \mathcal{O}(V) \mid gf = f \text{ for all } g \in G\}$. Observe that when restricted to $V^* \subset \mathcal{O}(V)$, this action coincides with the action corresponding to the dual representation. In terms of a basis $x_1, \dots, x_n$ of V^* , the regular functions are polynomials in the x_i and the action of G is given by $gp(x_1, \dots, x_n) = p(gx_1, \dots, gx_n)$ for any polynomial p . Since for every d , G maps the set of polynomials homogeneous of degree d to itself, it follows that the homogeneous parts of an invariant are invariant as well. This shows that $\mathcal{O}(V)^G = \bigoplus_d \mathcal{O}(V)_d^G$, where $\mathcal{O}(V)_d^G := \mathcal{O}(V)_d \cap \mathcal{O}(V)^G$.

Example 1.3.2. Consider the representation $\rho : \mathbb{Z}/3\mathbb{Z} \rightarrow \mathrm{GL}_2(\mathbb{C})$ defined by mapping 1 to the matrix $\begin{pmatrix} 0 & -1 \\ 1 & -1 \end{pmatrix}$ (and mapping 2 to $\begin{pmatrix} -1 & 1 \\ -1 & 0 \end{pmatrix}$ and 0 to the identity matrix). With respect to the dual basis x_1, x_2 , the dual representation is given by:

$$\rho^*(0) = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad \rho^*(1) = \begin{pmatrix} -1 & -1 \\ 1 & 0 \end{pmatrix}, \quad \rho^*(2) = \begin{pmatrix} 0 & 1 \\ -1 & -1 \end{pmatrix}. \quad (1.2)$$

The polynomial $f = x_1^2 - x_1x_2 + x_2^2$ is an invariant:

$$\rho^*(1)f = (-x_1 + x_2)^2 - (-x_1 + x_2)(-x_1) + (-x_1)^2 = x_1^2 - x_1x_2 + x_2^2 = f, \quad (1.3)$$

and since 1 is a generator of the group, f is invariant under all elements of the group. Other invariants are $x_1^2x_2 - x_1x_2^2$ and $x_1^3 - 3x_1x_2^2 + x_2^3$. These three invariants generate the ring of invariants, although it requires some work to show that.

A simpler example in which the complete ring of invariants can be computed is the following.

Example 1.3.3. Let D_4 be the symmetry group of the square, generated by a rotation r , a reflection s and the relations $r^4 = e, s^2 = e$ and $srs = r^3$, where e is the identity. The representation ρ of D_4 on $\mathbb{C}^2$ is given by

$$\rho(r) = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}, \quad \rho(s) = \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix}, \quad (1.4)$$

the dual representation is given by the same matrices:

$$\rho^*(r) = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}, \quad \rho^*(s) = \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix}. \quad (1.5)$$

It is easy to check that $x_1^2 + x_2^2$ and $x_1^2x_2^2$ are invariants, and so are all polynomial expressions in these two invariants. We will show that in fact $\mathcal{O}(\mathbb{C}^2)^{D_4} = \mathbb{C}[x_1^2 + x_2^2, x_1^2x_2^2] =: R$. It suffices to show that all homogeneous invariants belong to R .

Let $p \in \mathbb{C}[x_1, x_2]$ be a homogeneous invariant. Since $sp = p$, only monomials having even exponents for x_1 can occur in p . Since r^2s exchanges x_1 and x_2 , for every monomial $x_1^a x_2^b$ in p , the monomial $x_1^b x_2^a$ must occur with the same exponent. This proves the claim since every polynomial of the form $x_1^{2n} x_2^{2m} + x_1^{2m} x_2^{2n}$ is an element of R . Indeed, we may assume that $n \leq m$ and proceed by induction on $n + m$, the case $n + m = 0$ being trivial. If $n > 0$ we have $q = (x_1^2 x_2^2)^n (x_2^{2m-2n} + x_1^{2m-2n})$ and we are done. If $n = 0$ we have $2q = 2(x_1^{2m} + x_2^{2m}) = 2(x_1^2 + x_2^2)^m - \sum_{i=1}^{m-1} \binom{m}{i} (x_1^{2i} x_2^{2m-2i})$ and we are done by induction again.

1.4 Conjugacy classes of matrices

In this section we discuss the polynomial functions on the square matrices, invariant under conjugation of the matrix variable by elements of $\mathrm{GL}_n(\mathbb{C})$. This example shows some tricks that are useful when proving that certain invariants are equal. Denote by $M_n(\mathbb{C})$ the vectorspace of complex $n \times n$ matrices. We consider the action of $G = \mathrm{GL}_n(\mathbb{C})$ on $M_n(\mathbb{C})$ by conjugation: $(g, A) \mapsto gAg^{-1}$ for $g \in \mathrm{GL}_n(\mathbb{C})$ and $A \in M_n(\mathbb{C})$. We are interested in finding all polynomials in the entries of $n \times n$ matrices that are invariant under G . Two invariants are given by the functions $A \mapsto \det A$ and $A \mapsto \mathrm{tr} A$.

Let

$$\chi_A(t) := \det(tI - A) = t^n - s_1(A)t^{n-1} + s_2(A)t^{n-2} - \cdots + (-1)^n s_n(A) \quad (1.6)$$

be the characteristic polynomial of A . Here the s_i are polynomials in the entries of A . Clearly,

$$\chi_{gAg^{-1}}(t) = \det(g(tI - A)g^{-1}) = \det(tI - A) = \chi_A(t) \quad (1.7)$$

holds for all $t \in \mathbb{C}$. It follows that the functions $s_1, \dots, s_n$ are G -invariant. Observe that $s_1(A) = \operatorname{tr} A$ and $s_n(A) = \det A$.

Proposition 1.4.1. *The functions $s_1, \dots, s_n$ generate $\mathcal{O}(\operatorname{Mat}_n(\mathbb{C}))^{\operatorname{GL}_n(\mathbb{C})}$ and are algebraically independent.*

Proof. To each $c = (c_1, \dots, c_n) \in \mathbb{C}^n$ we associate the so-called *companion matrix*

$$A_c := \begin{pmatrix} 0 & \cdots & \cdots & 0 & -c_n \\ 1 & \ddots & & \vdots & -c_{n-1} \\ 0 & \ddots & \ddots & \vdots & \vdots \\ \vdots & \ddots & \ddots & 0 & c_2 \\ 0 & \cdots & 0 & 1 & c_1 \end{pmatrix} \in M_n(\mathbb{C}). \quad (1.8)$$

A simple calculation shows that $\chi_{A_c}(t) = t^n + c_{n-1}t^{n-1} + \cdots + c_1t + c_0$.

Exercise 1.4.2. Verify that $\chi_{A_c}(t) = t^n + c_{n-1}t^{n-1} + \cdots + c_1t + c_0$.

This implies that $s_i(A_c) = (-1)^i c_i$ and therefore

$$\{(s_1(A_c), s_2(A_c), \dots, s_n(A_c)) \mid A \in M_n(\mathbb{C})\} = \mathbb{C}^n. \quad (1.9)$$

It follows that the s_i are algebraically independent over $\mathbb{C}$. Indeed, suppose that $p(s_1, \dots, s_n) = 0$ for some polynomial p in n variables. Then

$$0 = p(s_1, \dots, s_n)(A) = p(s_1(A), \dots, s_n(A)) \quad (1.10)$$

for all A and hence $p(c_1, \dots, c_n) = 0$ for all $c \in \mathbb{C}^n$. But this implies that p itself is the zero polynomial.

Now let $f \in \mathcal{O}(\operatorname{Mat}_n(\mathbb{C}))^G$ be an invariant function. Define the polynomial p in n variables by $p(c_1, \dots, c_n) := f(A_c)$, and $P \in \mathcal{O}(\operatorname{Mat}_n(\mathbb{C}))^G$ by $P(A) := p(-s_1(A), s_2(A), \dots, (-1)^n s_n(A))$. By definition, P and f agree on all companion matrices, and since they are both G -invariant they agree on $W := \{gA_cg^{-1} \mid g \in G, c \in \mathbb{C}^n\}$. To finish the proof, it suffices to show that W is dense in $\operatorname{Mat}_n(\mathbb{C})$ since $f - P$ is continuous and zero on W . To show that W is dense in $\operatorname{Mat}_n(\mathbb{C})$, it suffices to show that the set of matrices with n distinct nonzero eigenvalues is a subset of W and is itself dense in $\operatorname{Mat}_n(\mathbb{C})$. This we leave as an exercise.

Exercise 1.4.3. Let $A \in \operatorname{Mat}_n(\mathbb{C})$ have n distinct nonzero eigenvalues. Show that A is conjugate to A_c for some $c \in \mathbb{C}^n$. Hint: find $v \in \mathbb{C}^n$ such that

$v, Av, A^2v, \dots, A^{n-1}v$ is a basis for $\mathbb{C}^n$. You might want to use the fact that the Vandermonde determinant

$$\det \begin{pmatrix} 1 & \dots & 1 \\ c_1 & \dots & c_n \\ c_1^2 & \dots & c_n^2 \\ \vdots & \ddots & \vdots \\ c_1^{n-1} & \dots & c_n^{n-1} \end{pmatrix} \quad (1.11)$$

is nonzero if $c_1, \dots, c_n$ are distinct and nonzero.

Exercise 1.4.4. Show that the set of matrices with n distinct nonzero eigenvalues is dense in the set of all complex $n \times n$ matrices. Hint: every matrix is conjugate to an upper triangular matrix.

□

1.5 Exercises

Exercise 1.5.1. Let G be a finite group acting on $V = \mathbb{C}^n$, $n \geq 1$. Show that $\mathcal{O}(V)^G$ contains a nontrivial invariant. That is, $\mathcal{O}(V)^G \neq \mathbb{C}$. Give an example of an action of an infinite group G on V with the property that only the constant functions are invariant.

Exercise 1.5.2. Let $\rho : \mathbb{Z}/2\mathbb{Z} \rightarrow \mathrm{GL}_2(\mathbb{C})$ be the representation given by $\rho(1) := \begin{pmatrix} -1 & 0 \\ 0 & -1 \end{pmatrix}$. Compute the invariant ring. That is, give a minimal set of generators for $\mathcal{O}(\mathbb{C}^2)^{\mathbb{Z}/2\mathbb{Z}}$.

Exercise 1.5.3. Let $U := \{ \begin{pmatrix} 1 & a \\ 0 & 1 \end{pmatrix} \mid a \in \mathbb{C} \}$ act on $\mathbb{C}^2$ in the obvious way. Denote the coordinate functions by x_1, x_2 . Show that $\mathcal{O}(\mathbb{C}^2)^U = \mathbb{C}[x_2]$.

Exercise 1.5.4. Let $\rho : \mathbb{C}^* \rightarrow \mathrm{GL}_3(\mathbb{C})$ be the representation given by $\rho(t) = \begin{pmatrix} t^{-2} & 0 & 0 \\ 0 & t^{-3} & 0 \\ 0 & 0 & t^4 \end{pmatrix}$. Find a minimal system of generators for the invariant ring.

Exercise 1.5.5. Let $\pi : \mathrm{Mat}_n(\mathbb{C}) \rightarrow \mathbb{C}^n$ be given by $\pi(A) := (s_1(A), \dots, s_n(A))$. Show that for every $c \in \mathbb{C}^n$ the fiber $\{A \mid \pi(A) = c\}$ contains a unique conjugacy class $\{gAg^{-1} \mid g \in \mathrm{GL}_n(\mathbb{C})\}$ of a diagonalizable (semisimple) matrix A .

Chapter 2

Lecture 2. Symmetric polynomials

In this chapter, we consider the natural action of the symmetric group S_n on the ring of polynomials in the variables $x_1, \dots, x_n$. The fundamental theorem of symmetric polynomials states that the elementary symmetric polynomials generate the ring of invariants. As an application we prove a theorem of Sylvester that characterizes when a univariate polynomial with real coefficients has only real roots.

2.1 Symmetric polynomials

Let the group S_n act on the polynomial ring $\mathbb{C}[x_1, \dots, x_n]$ by permuting the variables:

$$\sigma p(x_1, \dots, x_n) := p(x_{\sigma(1)}, \dots, x_{\sigma(n)}) \text{ for all } \sigma \in S_n. \quad (2.1)$$

The polynomials invariant under this action of S_n are called *symmetric polynomials*. As an example, for $n = 3$ the polynomial $x_1^2x_2 + x_1^2x_3 + x_1x_2^2 + x_1x_3^2 + x_2^2x_3 + x_2x_3^2 + 7x_1 + 7x_2 + 7x_3$ is symmetric, but $x_1^2x_2 + x_1x_3^2 + x_2^2x_3$ is not symmetric (although it is invariant under the alternating group).

In terms of linear representations of a group, we have a linear representation $\rho : S_n \rightarrow \text{GL}_n(\mathbb{C})$ given by $\rho(\sigma)e_i := e_{\sigma(i)}$, where $e_1, \dots, e_n$ is the standard basis of $\mathbb{C}^n$. On the dual basis $x_1, \dots, x_n$ the dual representation is given by $\rho^*(\sigma)x_i = x_{\sigma(i)}$, as can be easily checked. The invariant polynomial functions on $\mathbb{C}^n$ are precisely the symmetric polynomials.

Some obvious examples of symmetric polynomials are

$$s_1 := x_1 + x_2 + \dots + x_n \text{ and} \quad (2.2)$$

$$s_2 := x_1x_2 + x_1x_3 + \dots + x_1x_n + \dots + x_{n-1}x_n \quad (2.3)$$

More generally, for every $k = 1, \dots, n$, the k -th elementary symmetric poly-

mial

$$s_k := \sum_{i_1 < \dots < i_k} x_{i_1} \cdots x_{i_k} \quad (2.4)$$

is invariant. Recall that these polynomials express the coefficients of a univariate polynomial in terms of its roots:

$$\prod_{i=1}^n (t - x_i) = t^n + \sum_{k=1}^n (-1)^k s_k t^{n-k}. \quad (2.5)$$

Moreover, if g is any polynomial in n variables $y_1, \dots, y_n$, then $g(s_1, \dots, s_n)$ is again a polynomial in the x_i which is invariant under all coordinate permutations. A natural question is: which symmetric polynomials are expressible as a polynomial in the elementary symmetric polynomials. For example $x_1^2 + \dots + x_n^2$ is clearly symmetric and it can be expressed in terms of the s_i :

$$x_1^2 + \dots + x_n^2 = s_1^2 - 2s_2. \quad (2.6)$$

It is a beautiful fact that the elementary symmetric polynomials generate *all* symmetric polynomials.

Theorem 2.1.1 (Fundamental theorem of symmetric polynomials). *Every S_n -invariant polynomial $f(x_1, \dots, x_n)$ in the x_i can be written as $g(s_1, \dots, s_n)$, where $g = g(y_1, \dots, y_n)$ is a polynomial in n variables. Moreover, given f , the polynomial g is unique.*

The proof of this result uses the *lexicographic order* on monomials in the variables $\underline{x} = (x_1, \dots, x_n)$. We say that $\underline{x}^\alpha := x_1^{\alpha_1} \cdots x_n^{\alpha_n}$ is (lexicographically) larger than $\underline{x}^\beta$ if there is a k such that $\alpha_k > \beta_k$ and $\alpha_i = \beta_i$ for all $i < k$. So for instance $x_1^2 > x_1 x_2^4 > x_1 x_2^3 > x_1 x_2 x_3^5$, etc. The *leading monomial* $\text{lm}(f)$ of a non-zero polynomial f in the x_i is the largest monomial (with respect to this ordering) that has non-zero coefficient in f .

Exercise 2.1.2. Check that $\text{lm}(fg) = \text{lm}(f)\text{lm}(g)$ and that $\text{lm}(s_k) = x_1 \cdots x_k$.

Exercise 2.1.3. Show that there are no infinite lexicographically strictly decreasing chains of monomials.

Since every decreasing chain of monomials is finite, we can use this order to do induction on monomials, as we do in the following proof.

Proof of Theorem 2.1.1. Let f be any S_n -invariant polynomial in the x_i . Let $\underline{x}^\alpha$ be the leading monomial of f . Then $\alpha_1 \geq \dots \geq \alpha_n$ because otherwise a suitable permutation applied to $\underline{x}^\alpha$ would yield a lexicographically larger monomial, which has the same non-zero coefficient in f as $\underline{x}^\alpha$ by invariance of f . Now consider the expression

$$s_n^{\alpha_n} s_{n-1}^{\alpha_{n-1} - \alpha_n} \cdots s_1^{\alpha_1 - \alpha_2}. \quad (2.7)$$

The leading monomial of this polynomial equals

$$(x_1 \cdots x_n)^{\alpha_n} (x_1 \cdots x_{n-1})^{\alpha_{n-1} - \alpha_n} \cdots x_1^{\alpha_1 - \alpha_2}, \quad (2.8)$$

which is just $\underline{x}^\alpha$. Subtracting a scalar multiple of the expression from f therefore cancels the term with monomial $\underline{x}^\alpha$, and leaves an S_n -invariant polynomial with a strictly smaller leading monomial. After repeating this step finitely many times, we have expressed f as a polynomial in the s_k .

This shows existence of g in the theorem. For uniqueness, let $g \in \mathbb{C}[y_1, \dots, y_n]$ be a nonzero polynomial in n variables. It suffices to show that $g(s_1, \dots, s_n) \in \mathbb{C}[x_1, \dots, x_n]$ is not the zero polynomial. Observe that

$$\text{lm}(s_1^{\alpha_1} \cdots s_n^{\alpha_n}) = x_1^{\alpha_1 + \cdots + \alpha_n} x_2^{\alpha_2 + \cdots + \alpha_n} \cdots x_n^{\alpha_n}. \quad (2.9)$$

It follows that the leading monomials of the terms $s_1^{\alpha_1} \cdots s_n^{\alpha_n}$, corresponding to the monomials occurring with nonzero coefficient in $g = \sum_{\alpha} \underline{y}^\alpha$, are pairwise distinct. In particular, the largest such leading monomial will not be cancelled in the sum and is the leading monomial of $g(s_1, \dots, s_n)$. $\square$

Remark 2.1.4. The proof shows that in fact the coefficients of the polynomial g lie in the ring generated by the coefficients of f . In particular, when f has real coefficients, also g has real coefficients.

Exercise 2.1.5. Let $\pi : \mathbb{C}^n \rightarrow \mathbb{C}^n$ be given by

$$\pi(x_1, \dots, x_n) = (s_1(x_1, \dots, x_n), \dots, s_n(x_1, \dots, x_n)). \quad (2.10)$$

Use the fact that every univariate polynomial over the complex numbers can be factorised into linear factors to show that π is surjective. Use this to show that $s_1, \dots, s_n$ are algebraically independent over $\mathbb{C}$. Describe for $b \in \mathbb{C}^n$ the fiber $\pi^{-1}(b)$.

Remark 2.1.6. The above proof of the fundamental theorem of symmetric polynomials gives an algorithm to write a given symmetric polynomial as a polynomial in the elementary symmetric polynomials. In each step the initial monomial of the residual symmetric polynomial is decreased, ending with the zero polynomial after a finite number of steps. Instead of using the described lexicographic order on the monomials, other linear orders can be used. An example would be the *degree lexicographic order*, where we set $\underline{x}^\alpha > \underline{x}^\beta$ if either $\alpha_1 + \cdots + \alpha_n > \beta_1 + \cdots + \beta_n$ or equality holds and there is a k such that $\alpha_k > \beta_k$ and $\alpha_i = \beta_i$ for all $i < k$.

Example 2.1.7. We write $x_1^3 + x_2^3 + x_3^3$ as a polynomial in the s_i . Since the leading monomial is $x_1^3 x_2^0 x_3^0$ we subtract $s_3^0 s_2^0 s_1^3$ and are left with $-3(x_1^2 x_2 + x_1^2 x_3 + x_1 x_2^2 + x_1 x_3^2 + x_2^2 x_3 + x_2 x_3^2) - 6(x_1 x_2 x_3)$. The leading monomial is now $x_1^2 x_2$, so we add $3s_3^0 s_2^1 s_1^{2-1}$. This leaves $3x_1 x_2 x_3 = 3s_3^1 s_2^{1-1} s_1^{1-1}$, which is reduced to zero in the next step.

This way we obtain $x_1^3 + x_2^3 + x_3^3 = s_1^3 - 3s_1 s_2 + 3s_3$.

Exercise 2.1.8. Give an upper bound on the number of steps of the algorithm in terms of the number of variables n and the (total) degree of the input polynomial f .

2.2 Counting real roots

Given a (monic) polynomial $f(t) = t^n - a_1 t^{n-1} + \cdots + (-1)^n a_n$, the coefficients are elementary symmetric functions in the roots of f . Therefore, any property that can be expressed as a symmetric polynomial in the roots of f , can also be expressed as a polynomial in the coefficients of f . This way we can determine properties of the roots by just looking at the coefficients of f . For example: when are all roots of f distinct?

Definition 2.2.1. For a (monic) polynomial $f = (t - x_1) \cdots (t - x_n)$, define the *discriminant* $\Delta(f)$ of f by $\Delta(f) := \prod_{1 \leq i < j \leq n} (x_i - x_j)^2$.

Clearly, $\Delta(f) = 0$ if and only if all roots of f are distinct. It is not hard to see that $\Delta(f)$ is a symmetric polynomial in the roots of f . We will see later how f can be expressed in terms of the coefficients of f .

Exercise 2.2.2. Let $f(t) = t^2 - at + b$. Write $\Delta(f)$ as a polynomial in a and b .

Definition 2.2.3. Given n complex numbers $x_1, \dots, x_n$, the *Vandermonde matrix* A for these numbers is given by

$$A := \begin{pmatrix} 1 & x_1 & \cdots & x_1^{n-1} \\ 1 & x_2 & \cdots & x_2^{n-1} \\ \vdots & \vdots & \cdots & \vdots \\ 1 & x_n & \cdots & x_n^{n-1} \end{pmatrix}. \quad (2.11)$$

Lemma 2.2.4. Given numbers $x_1, \dots, x_n$, the Vandermonde matrix A has nonzero determinant if and only if the $x_1, \dots, x_n$ are distinct.

Proof. View the determinant of the Vandermonde matrix (called the *Vandermonde determinant*) as a polynomial p in the variables $x_1, \dots, x_n$. For any $i < j$, $p(x_1, \dots, x_n) = 0$ when $x_i = x_j$ and hence p is divisible by $(x_j - x_i)$. Expanding the determinant, we see that p is homogeneous of degree $\binom{n}{2}$, with lowest monomial $x_1^0 x_2^1 \cdots x_n^{n-1}$ having coefficient 1. It follows that

$$p = \prod_{1 \leq i < j \leq n} (x_j - x_i), \quad (2.12)$$

since the right-hand side divides p , and the two polynomials have the same degree and the same nonzero coefficient for $x_1^0 x_2^1 \cdots x_n^{n-1}$. $\square$

Exercise 2.2.5. Show that the Vandermonde matrix A of numbers $x_1, \dots, x_n$ satisfies $\det A = \prod_{1 \leq i < j \leq n} (x_j - x_i)$ by doing row- and column-operations on A and applying induction on n .

Definition 2.2.6. Let $f = (t - \alpha_1)(t - \alpha_2) \cdots (t - \alpha_n) \in \mathbb{C}[t]$ be a monic polynomial of degree n in the variable t . We define the *Bezoutiant matrix* $\text{Bez}(f)$ of f by

$$\text{Bez}(f) = (p_{i+j-2}(\alpha_1, \dots, \alpha_n))_{i,j=1}^n, \quad (2.13)$$

where $p_k(x_1, \dots, x_n) := x_1^k + \cdots + x_n^k$ for $k = 0, 1, \dots$ is the k -th *Newton polynomial*.

Since the entries of $\text{Bez}(f)$ are symmetric polynomials in the roots of f , it follows by the fundamental theorem of symmetric polynomials that the entries are polynomials (with integer coefficients) in the elementary symmetric functions and hence in the coefficients of f . In particular, when f has real coefficients, $\text{Bez}(f)$ is a real matrix. Another useful fact is that $\text{Bez}(f) = A^\top A$, where A is the Vandermonde matrix for the roots $\alpha_1, \dots, \alpha_n$ of f .

Exercise 2.2.7. Show that the discriminant of f satisfies: $\Delta(f) = \det \text{Bez}(f)$.

Example 2.2.8. Let $f = t^2 - at + b$ have roots α and β . So $a = \alpha + \beta$ and $b = \alpha\beta$. We compute $\text{Bez}(f)$. We have $p_0 = 2$, $p_1 = a$, $p_2 = a^2 - 2b$ so $\text{Bez}(f) = \begin{pmatrix} 2 & a \\ a & a^2 - 2b \end{pmatrix}$. The determinant equals $a^2 - 4b$ and the trace equals $a^2 - 2b + 2$. There are three cases for the eigenvalues $\lambda_1 \geq \lambda_2$ of $\text{Bez}(f)$:

- If $a^2 - 4b > 0$, we have $\lambda_1, \lambda_2 > 0$ and α, β are distinct real roots.
- If $a^2 - 4b = 0$, we have $\lambda_1 > 0, \lambda_2 = 0$ and $\alpha = \beta$.
- If $a^2 - 4b < 0$, we have $\lambda_1 > 0, \lambda_2 < 0$ and α and β are complex conjugate (nonreal) roots.

The determinant of $\text{Bez}(f)$ determines whether f has double roots. The matrix $\text{Bez}(f)$ can give us much more information about the roots of f . In particular, it describes when a polynomial with real coefficients has only real roots!

Theorem 2.2.9 (Sylvester). *Let $f \in \mathbb{R}[t]$ be a polynomial in the variable t with real coefficients. Let r be the number of distinct roots in $\mathbb{R}$ and $2k$ the number of distinct roots in $\mathbb{C} \setminus \mathbb{R}$. Then the Bezoutiant matrix $\text{Bez}(f)$ has rank $r + 2k$, with $r + k$ positive eigenvalues and k negative eigenvalues.*

proof of Theorem 2.2.9. Number the roots $\alpha_1, \dots, \alpha_n$ of f in such a way that $\alpha_1, \dots, \alpha_{2k+r}$ are distinct. We write m_i for the multiplicity of the root α_i , $i = 1, \dots, 2k+r$. Let A be the Vandermonde matrix for the numbers $\alpha_1, \dots, \alpha_n$, so that $\text{Bez}(f) = A^\top A$. We start by computing the rank of $\text{Bez}(f)$.

Denote by $\tilde{A}$ the $(2k+r) \times n$ submatrix of A consisting of the first $2k+r$ rows of A . An easy computation shows that

$$\text{Bez}(f) = A^\top A = \tilde{A}^\top \text{diag}(m_1, \dots, m_{2k+r}) \tilde{A}, \quad (2.14)$$

where $\text{diag}(m_1, \dots, m_{2k+r})$ is the diagonal matrix with the multiplicities of the roots on the diagonal. Since, $\tilde{A}$ contains a submatrix equal to the Vandermonde matrix for the distinct roots $\alpha_1, \dots, \alpha_{2k+r}$, it follows by Lemma 2.2.4 that the rows of $\tilde{A}$ are linearly independent. Since the diagonal matrix has full rank, it follows that $\text{Bez}(f)$ has rank $2k+r$.

To finish the proof, we write $A = B + iC$, where B and C are real matrices and i denotes a square root of -1 . Since f has real coefficients, $\text{Bez}(f)$ is a real matrix and hence

$$\text{Bez}(f) = B^\top B - C^\top C + i(C^\top B + B^\top C) = B^\top B - C^\top C. \quad (2.15)$$

We have

$$\text{rank}(B) \leq r + k, \quad \text{rank}(C) \leq k. \quad (2.16)$$

Indeed, for any pair $\alpha, \bar{\alpha}$ of complex conjugate numbers, the real parts of α^j and $\bar{\alpha}^j$ are equal and the imaginary parts are opposite. Hence B has at most $r + k$ different rows and C has (up to a factor -1) at most k different nonzero rows.

Denote the kernel of $\text{Bez}(f)$, B and C by N , N_B and N_C respectively. Clearly $N_B \cap N_C \subseteq N$. This gives

$$\begin{aligned} \dim N \geq \dim(N_B \cap N_C) &\geq \dim N_B + \dim N_C - n \\ &\geq (n - r - k) + (n - k) - n \\ &= n - r - 2k = \dim N. \end{aligned} \quad (2.17)$$

Hence we have equality throughout, showing that $\dim N_B = n - r - k$, $\dim N_C = n - k$ and $N_B \cap N_C = N$.

Write $N_B = N \oplus N'_B$ and $N_C = N \oplus N'_C$ as a direct sum of vector spaces. For all nonzero $u \in N'_C$, we have $u^\top C^\top C u = 0$ and $u^\top B^\top B u > 0$ and so $u^\top \text{Bez}(f) u > 0$. This shows that $\text{Bez}(f)$ has at least $\dim N'_C = r + k$ positive eigenvalues (see exercises). Similarly, $u^\top \text{Bez}(f) u < 0$ for all nonzero $u \in N'_B$ so that $\text{Bez}(f)$ has at least $\dim N'_B = k$ negative eigenvalues. Since $\text{Bez}(f)$ has $n - r - 2k$ zero eigenvalues, it has exactly $r + k$ positive eigenvalues and exactly k negative eigenvalues. $\square$

Exercise 2.2.10. Let B be a real $n \times n$ matrix and $x \in \mathbb{R}^n$. Show that $x^\top B^\top B x \geq 0$ and that equality holds if and only if $Bx = 0$.

Exercise 2.2.11. Let A be a real symmetric $n \times n$ matrix. Show that the following are equivalent:

- there exists a linear subspace $V \subseteq \mathbb{R}^n$ of dimension k such that $x^\top A x > 0$ for all nonzero $x \in V$,
- A has at least k positive eigenvalues.

Exercise 2.2.12. Use the previous exercise to show Sylvesters law of inertia: Given a real symmetric $n \times n$ matrix A and an invertible real matrix S , the two matrices A and $S^\top A S$ have the same number of positive, negative and zero eigenvalues. This implies that the signature of A can be easily determined by bringing it into diagonal form using simultaneous row and column operations.

2.3 Exercises

Exercise 2.3.1. Let $f(t) := t^3 + at + b$, where a, b are real numbers.

- Compute $\text{Bez}(f)$.
- Show that $\Delta(f) = -4a^3 - 27b^2$.

- Determine, in terms of a and b , when f has only real roots.

Exercise 2.3.2. Prove the following formulas due to Newton:

$$p_k - s_1 p_{k-1} + \cdots + (-1)^{k-1} s_{k-1} p_1 + (-1)^k k s_k = 0 \quad (2.18)$$

for all $k = 1, \dots, n$.

Show that for $k > n$ the following similar relation holds:

$$p_k - s_1 p_{k-1} + \cdots + (-1)^n s_n p_{k-n} = 0. \quad (2.19)$$

Hint: Let $f(t) = (1 - tx_1) \cdots (1 - tx_n)$ and compute $f'(t)/f(t)$ in two ways.

Chapter 3

Lecture 3. Multilinear algebra

We review some constructions from linear algebra, in particular the tensor product of vector spaces. Unless explicitly stated otherwise, all our vector spaces are over the field $\mathbb{C}$ of complex numbers.

Definition 3.0.3. Let $V_1, \dots, V_k, W$ be vector spaces. A map $\phi : V_1 \times \dots \times V_k \rightarrow W$ is called *multilinear* (or *k-linear* or bilinear if $k = 2$ or trilinear if $k = 3$) if for each i and all $v_1, \dots, v_{i-1}, v_{i+1}, \dots, v_k$ the map $V_i \rightarrow W, v_i \mapsto \phi(v_1, \dots, v_k)$ is linear.

Let U, V and T be vector spaces and let $\otimes : U \times V \rightarrow T$ be a bilinear map. The map $\otimes$ is said to have the *universal property* if for every bilinear map $\phi : U \times V \rightarrow W$ there exists a *unique* linear map $f : T \rightarrow W$ such that $\phi = f \circ \otimes$.

$$\begin{array}{ccc} U \times V & \xrightarrow{\phi} & W \\ \otimes \downarrow & \nearrow f & \\ T & & \end{array}$$

We will usually write $u \otimes v := \otimes(u, v)$ for $(u, v) \in U \times V$. Although $\otimes$ will in general not be surjective, the image linearly spans T .

Exercise 3.0.4. Show that if $\otimes : U \times V \rightarrow T$ has the universal property, the vectors $u \otimes v, u \in U, v \in V$ span T .

Given U and V , the pair $(T, \otimes)$ is *unique up to a unique isomorphism*. That is, given two bilinear maps $\otimes : U \times V \rightarrow T$ and $\otimes' : U \times V \rightarrow T'$ that both have the universal property, there is a unique linear isomorphism $f : T \rightarrow T'$ such that $f(u \otimes v) = u \otimes' v$ for all $u \in U, v \in V$. This can be seen as follows. Since $\otimes'$ is bilinear, there exists by the universal property of $\otimes$, a unique linear map $f : T \rightarrow T'$ such that $\otimes' = f \circ \otimes$. It suffices to show that f is a bijection. By the

universal property of $\otimes'$ there is a linear map $f' : T' \rightarrow T$ such that $\otimes' = f' \circ \otimes$. Now $\otimes \circ f' \circ f = \otimes$, which implies that $f' \circ f : T \rightarrow T$ is the identity since the image of $\otimes$ spans T (or alternatively, by using the universal property of $\otimes$, and the bilinear map $\otimes$ itself). Hence f is injective. Similarly, $f \circ f'$ is the identity on T' and hence f is surjective.

Definition 3.0.5. Let U, V be vector spaces. The *tensor product* of U and V is a vector space T together with a bilinear map $\otimes : U \times V \rightarrow T$ having the universal property. The space T , which is uniquely determined by U and V up to an isomorphism, is denoted by $U \otimes V$.

Often we will refer to $U \otimes V$ as the tensor product of U and V , implicitly assuming the map $\otimes : U \times V \rightarrow U \otimes V$.

So far, we have not shown that the tensor product $U \otimes V$ exists at all, nor did we gain insight into the dimension of this space in terms of the dimensions of U and V . One possible construction of $U \otimes V$ is as follows.

Start with the vector space F (for *free* or *formal*) formally spanned by pairs (u, v) as u, v run through U, V , respectively. Now take the subspace R (for *relations*) of F spanned by all elements of the form

$$(c_1 u + u', c_2 v + v') - c_1 c_2 (u, v) - c_1 (u, v') - c_2 (u', v) - (u', v') \quad (3.1)$$

with $c_1, c_2 \in \mathbb{C}$, $v, v' \in V, u, u' \in U$. Now any map $\phi : U \times V \rightarrow W$ factors through the injection $i : U \times V \rightarrow F$ and a unique linear map $g : F \rightarrow W$. The kernel of g contains R if and only if ϕ is bilinear, and in that case the map g factors through the quotient map $\pi : F \rightarrow F/R$ and a unique linear map $f : F/R \rightarrow W$. Taking for $\otimes$ the bilinear map $\pi \circ i : (u, v) \mapsto u \otimes v$, the space F/R together with the map $\otimes$ is the tensor product of U and V .

As for the dimension of $U \otimes V$, let $(u_i)_{i \in I}$ be a basis of U . Then by using bilinearity of the tensor product, every element $T \in U \otimes V$ can be written as a $t = \sum_{i \in I} u_i \otimes w_i$ with w_i non-zero for only finitely many i . We claim that the w_i in such an expression are unique. Indeed, for $k \in I$ let ξ_k be the linear function on U determined by $u_i \mapsto \delta_{ik}$, $i \in I$. The bilinear map $U \times V \rightarrow V$, $(u, v) \mapsto \xi_k(u)v$ factors, by the universal property, through a unique linear map $f : U \otimes V \rightarrow V$. This map sends all terms in the expression $\sum_{i \in I} u_i \otimes w_i$ for T to zero except the term with $i = k$, which is mapped to w_k . Hence $w_k = f_k(t)$ and this shows the uniqueness of the w_k .

Exercise 3.0.6. Use a similar argument to show that if $(v_j)_{j \in J}$ is a basis for V , then the set of all elements of the form $u_i \otimes v_j$, $i \in I$, $j \in J$ form a basis of $U \otimes V$.

This exercise may remind you of matrices. Indeed, there is a natural map ϕ from $U \otimes V^*$, where V^* is the dual of V , into the space $\text{Hom}(V, U)$ of linear maps $V \rightarrow U$, defined as follows. Given a pair $u \in U$ and $f \in V^*$, $\phi(u \otimes f)$ is the linear map sending v to $f(v)u$. Here we are implicitly using the universal property: the linear map $v \mapsto f(v)u$ depends bilinearly on f and u , hence there is a unique linear map $U \otimes V^* \rightarrow \text{Hom}(V, U)$ that sends $u \otimes f$ to $v \mapsto f(v)u$.

Note that if f and u are both non-zero, then the image of $u \otimes f$ is a linear map of rank one.

- Exercise 3.0.7.** 1. Show that ϕ is injective. Hint: after choosing a basis $(u_i)_i$ use that a general element of $U \otimes V^*$ can be written in a unique way as $\sum_i u_i \otimes f_i$.
2. Show that ϕ is surjective onto the subspace of $\text{Hom}(V, U)$ of linear maps of finite rank, that is, having finite-dimensional image.

Making things more concrete, if $U = \mathbb{C}^m$ and $V = \mathbb{C}^n$ and $u_1, \dots, u_m$ is the standard basis of U and $v_1, \dots, v_n$ is the standard basis of V with dual basis $x_1, \dots, x_n$, then the tensor $u_i \otimes x_j$ corresponds to the linear map with matrix E_{ij} , the matrix having zeroes everywhere except for a 1 in position (i, j) .

Remark 3.0.8. A common mistake is to assume that all elements of $U \otimes V$ are of the form $u \otimes v$. The above shows that the latter elements correspond to rank-one linear maps from V^* to U , or to rank-one matrices, while $U \otimes V$ consists of all finite-rank linear maps from V^* to U —a much larger set.

Next we discuss tensor products of linear maps. If $A : U \rightarrow U'$ and $B : V \rightarrow V'$ are linear maps, then the map $U \times V \rightarrow U' \otimes V'$, $(u, v) \mapsto Au \otimes Bv$ is bilinear. Hence, by the universal property of $U \otimes V$ there exists a unique linear map $U \otimes V \rightarrow U' \otimes V'$ that sends $u \otimes v$ to $Au \otimes Bv$. This map is denoted $A \otimes B$.

Example 3.0.9. If $\dim U = m, \dim U' = m', \dim V = n, \dim V' = n'$ and if A, B are represented by an $m' \times m$ -matrix $(a_{ij})_{ij}$ and an $n' \times n$ -matrix $(b_{kl})_{kl}$, respectively, then $A \otimes B$ can be represented by an $m'n' \times mn$ -matrix, with rows labelled by pairs (i, k) with $i \in [m'], k \in [n']$ and columns labelled by pairs (j, l) with $j \in [m], l \in [n]$, whose entry at position $((i, k), (j, l))$ is $a_{ij}b_{kl}$. This matrix is called the *Kronecker product* of the matrices $(a_{ij})_{ij}$ and $(b_{kl})_{kl}$.

Exercise 3.0.10. Assume, in the setting above, that $U = U', m' = m$ and $V = V', n' = n$ and A, B are diagonalisable with eigenvalues $\lambda_1, \dots, \lambda_m$ and $\mu_1, \dots, \mu_n$, respectively. Determine the eigenvalues of $A \otimes B$.

Most of what we said about the tensor product of two vector spaces carries over *verbatim* to the tensor product $V_1 \otimes \dots \otimes V_k$ of k . This tensor product can again be defined by a universal property involving k -linear maps, and its dimension is $\prod_i \dim V_i$. Its elements are called k -tensors. We skip the boring details, but do point out that for larger k there is no longer a close relationship with k -tensors with linear maps—in particular, the *rank* of a k -tensor T , usually defined as the minimal number of terms in any expression of T as a sum of *pure tensors* $v_1 \otimes \dots \otimes v_k$, is only poorly understood. For instance, *computing* the rank, which for $k = 2$ can be done using Gaussian elimination, is very hard in general. If all V_i are the same, say V , then we also write $V^{\otimes k}$ for $V \otimes \dots \otimes V$ (k factors).

Given three vector spaces U, V, W , we now have several ways to take their tensor product: $(U \otimes V) \otimes W$, $U \otimes (V \otimes W)$ and $U \otimes V \otimes W$. Fortunately,

these tensor products can be identified. For example, there is a unique linear isomorphism $f : U \otimes V \otimes W \rightarrow (U \otimes V) \otimes W$ such that $f(u \otimes v \otimes w) = (u \otimes v) \otimes w$ for all $u \in U, v \in V, w \in W$.

Indeed, consider the trilinear map $U \times V \times W \rightarrow (U \otimes V) \otimes W$ defined by $(u, v, w) \mapsto (u \otimes v) \otimes w$. By the universal property, there is a unique linear map $f : U \otimes V \otimes W \rightarrow (U \otimes V) \otimes W$ such that $f(u \otimes v \otimes w) = (u \otimes v) \otimes w$ for all u, v, w .

Now for fixed $w \in W$, the bilinear map $\phi_w : U \times V \rightarrow U \otimes V \otimes W$ defined by $\phi_w(u, v) := u \otimes v \otimes w$ induces a linear map $g_w : U \otimes V \rightarrow U \otimes V \otimes W$ such that $u \otimes v$ is mapped to $u \otimes v \otimes w$. Hence the bilinear map $\phi : (U \otimes V) \times W \rightarrow U \otimes V \otimes W$ given by $\phi(x, w) := g_w(x)$ induces a linear map $g : (U \otimes V) \otimes W \rightarrow U \otimes V \otimes W$ sending $(u \otimes v) \otimes w$ to $u \otimes v \otimes w$. It follows that $f \circ g$ and $g \circ f$ are the identity on $(U \otimes V) \otimes W$ and $U \otimes V \otimes W$ respectively. This shows that f is an isomorphism.

Exercise 3.0.11. Let V be a vector space. Show that for all p, q there is a unique linear isomorphism $V^{\otimes p} \otimes V^{\otimes q} \rightarrow V^{\otimes(p+q)}$ sending $(v_1 \otimes \cdots \otimes v_p) \otimes (v_{p+1} \otimes \cdots \otimes v_{p+q})$ to $v_1 \otimes \cdots \otimes v_{p+q}$.

The direct sum $TV := \bigoplus_{k=0}^{\infty} V^{\otimes k}$ is called the *tensor algebra* of V , where the natural linear map $V^{\otimes k} \times V^{\otimes l} \rightarrow V^{\otimes k} \otimes V^{\otimes l} = V^{\otimes(k+l)}$ plays the role of (non-commutative but associative) multiplication. We move on to other types of tensors.

Definition 3.0.12. Let V be a vector space. A k -linear map $\omega : V^k \rightarrow W$ is called *symmetric* if $\omega(v_1, \dots, v_k) = \omega(v_{\pi(1)}, \dots, v_{\pi(k)})$ for all permutations $\pi \in \text{Sym}(k)$.

The k -th *symmetric power* of V is a vector space $S^k V$ together with a symmetric k -linear map $V^k \rightarrow S^k V$, $(v_1, \dots, v_k) \mapsto v_1 \cdots v_k$ such that for all vector spaces W and symmetric k -linear maps $\psi : V^k \rightarrow W$ there is a unique linear map $\phi : S^k V \rightarrow W$ such that $\psi(u_1, \dots, u_k) = \phi(u_1 \cdots u_k)$.

Uniqueness of the k -th symmetric power of V can be proved in exactly the same manner as uniqueness of the tensor product. For existence, let R be the subspace of $V^{\otimes k} := V \otimes \cdots \otimes V$ spanned by all elements of the form

$$v_1 \otimes \cdots \otimes v_k - v_{\pi(1)} \otimes \cdots \otimes v_{\pi(k)}, \quad \pi \in \text{Sym}(k).$$

Then the composition of the maps $V^k \rightarrow V^{\otimes k} \rightarrow V^{\otimes k}/R$ is a symmetric k -linear map and if $\psi : V^k \rightarrow W$ is any such map, then ψ factors through a linear map $V^{\otimes k} \rightarrow W$ since it is k -linear, which in turn factors through a unique linear map $V^{\otimes k}/R \rightarrow W$ since ψ is symmetric. This shows existence of symmetric powers, and, perhaps more importantly, the fact that they are quotients of tensor powers of V . This observation will be very important in proving the first fundamental theorem for $\text{GL}(V)$.

There is also an analogue of the tensor product of maps: if A is a linear map $U \rightarrow V$, then the map $U^k \rightarrow S^k V$, $(u_1, \dots, u_k) \mapsto Au_1 \cdots Au_k$ is multilinear and symmetric. Hence, by the universal property of symmetric powers, it factors

through the map $U^k \rightarrow S^k U$ and a linear map $S^k U \rightarrow S^k V$. This map, which sends $u_1 \cdots u_k$ to $Au_1 \cdots Au_k$, is the k -th symmetric power $S^k A$ of A .

If $(v_i)_{i \in I}$ is a basis of V , then using multilinearity and symmetry every element t of $S^k V$ can be written as a linear combination $\sum_{\alpha} c_{\alpha} v^{\alpha}$ of the elements $v^{\alpha} := \prod_{i \in I} v_i^{\alpha_i}$ —the product order is immaterial—where $\alpha \in \mathbb{N}^I$ satisfies $|\alpha| := \sum_{i \in I} \alpha_i = k$ and only finitely many coefficients c_{α} are non-zero. We claim that the c_{α} are unique, so that the v^{α} , $|\alpha| = k$ a basis of $S^k V$. Indeed, let $\alpha \in \mathbb{N}^I$ with $|\alpha| = k$. Then there is a unique k -linear map $\psi_{\alpha} : V^k \rightarrow \mathbb{C}$ which on a tuple $(v_{i_1}, \dots, v_{i_k})$ takes the value 1 if $|\{j \mid i_j = i\}| = \alpha_i$ for all $i \in I$ and zero otherwise. Moreover, ψ_{α} is symmetric and therefore induces a linear map $\phi_{\alpha} : S^k V \rightarrow \mathbb{C}$. We find that $c_{\alpha} = \phi_{\alpha}(t)$, which proves the claim.

This may remind you of polynomials. Indeed, if $V = \mathbb{C}^n$ and $x_1, \dots, x_n$ is the basis of V^* dual to the standard basis of V , then $S^k V^*$ is just the space of *homogeneous polynomials* in the x_i of degree k . The algebra of all polynomial functions on V therefore is canonically isomorphic to $SV^* := \bigoplus_{k=0}^{\infty} S^k V^*$. The product of a homogeneous polynomials of degree k and homogeneous polynomials of degree l corresponds to the unique bilinear map $S^k V^* \times S^l V^* \rightarrow S^{k+l} V^*$ making the diagram

$$\begin{array}{ccc} (V^*)^{\otimes k} \times (V^*)^{\otimes l} & \longrightarrow & (V^*)^{\otimes k+l} \\ \downarrow & & \downarrow \\ S^k V^* \times S^l V^* & \longrightarrow & S^{k+l} V^* \end{array}$$

commute, and this corresponds to multiplying polynomials of degrees k and l . Thus SV^* is a quotient of the tensor algebra TV (in fact, the maximal commutative quotient).

Above we have introduced $S^k V$ as a *quotient* of $V^{\otimes k}$. This should not be confused with the *subspace* of $V^{\otimes k}$ spanned by all *symmetric tensors*, defined as follows. For every permutation $\pi \in S_k$ there is a natural map $V^k \rightarrow V^k$ sending $(v_1, \dots, v_k)$ to $(v_{\pi^{-1}(1)}, \dots, v_{\pi^{-1}(k)})$. Composing this map with the natural k -linear map $V^k \rightarrow V^{\otimes k}$ yields another k -linear map $V^k \rightarrow V^{\otimes k}$, and hence a linear map $V^{\otimes k} \rightarrow V^{\otimes k}$, also denoted π . A tensor ω in $V^{\otimes k}$ is called *symmetric* if $\pi\omega = \omega$ for all $\pi \in S_k$. The restriction of the map $V^{\otimes k} \rightarrow S^k V$ to the subspace of symmetric tensors is an isomorphism with inverse determined by $v_1 \cdots v_k \mapsto \frac{1}{k!} \sum_{\pi \in S_k} \pi(v_1 \otimes \cdots \otimes v_k)$. (Note that this inverse would not be defined in characteristic less than k .)

Exercise 3.0.13. Show that the subspace of symmetric tensors in $V^{\otimes k}$ is spanned by the tensors $v \otimes v \cdots \otimes v$, where $v \in V$.

3.1 Exercises

Exercise 3.1.1. Let $U \otimes V$ be the tensor product of the vector spaces U and V . Let $u_1, \dots, u_s$ and $u'_1, \dots, u'_t$ be two systems of linearly independent vectors

in U and let $v_1, \dots, v_s$ and $v'_1, \dots, v'_t$ be two systems of linearly independent vectors in V . Suppose that

$$u_1 \otimes v_1 + \dots + u_s \otimes v_s = u'_1 \otimes v'_1 + \dots + u'_t \otimes v'_t. \quad (3.2)$$

Show that $s = t$.

Exercise 3.1.2. a) Let $T \in V_1 \otimes V_2 \otimes V_3$ be an element of the tensor product of V_1 , V_2 and V_3 . Suppose that there exist $v_1 \in V_1$, $v_3 \in V_3$, $T_{23} \in V_2 \otimes V_3$ and $T_{12} \in V_1 \otimes V_2$ such that

$$T = v_1 \otimes T_{23} = T_{12} \otimes v_3. \quad (3.3)$$

Show that there exist a $v_2 \in V_2$ such that $T = v_1 \otimes v_2 \otimes v_3$.

b) Suppose that $T \in V_1 \otimes V_2 \otimes V_3$ can be written as a sum of at most d_1 tensors of the form $v_1 \otimes T_{23}$, where $v_1 \in V_1$, $T_{23} \in V_2 \otimes V_3$, and also as a sum of at most d_3 tensors of the form $T_{12} \otimes v_3$, where $v_3 \in V_3$, $T_{12} \in V_1 \otimes V_2$. Show that T can be written as the sum of at most $d_1 d_3$ tensors of the form $v_1 \otimes v_2 \otimes v_3$, where $v_i \in V_i$.

Exercise 3.1.3. Let U, V, W be vector spaces. Denote by $B(U \times V, W)$ the linear space of bilinear maps from $U \times V$ to W . Show that the map $f \mapsto f \circ \otimes$ is a linear isomorphism between $\text{Hom}(U \otimes V, W)$ and $B(U \times V, W)$.

Exercise 3.1.4. Let U, V be vector spaces. Show that the linear map $\phi : U^* \otimes V^* \rightarrow (U \otimes V)^*$ given by $\phi(f \otimes g)(u \otimes v) := f(u)g(v)$ is an isomorphism.

Chapter 4

Lecture 4. Representations

Central objects in this course are linear representations of groups. We will only consider representations on complex vector spaces. Recall the following definition.

Definition 4.0.5. Let G be a group and let V be a vector space. A *(linear) representation of G on V* is a group homomorphism $\rho : G \rightarrow \mathrm{GL}(V)$.

If ρ is a representation of G , then the map $(g, v) \mapsto \rho(g)v$ is an action of G on V . A vector space with an action of G by linear maps is also called a G -module. Instead of $\rho(g)v$ we will often write gv .

Definition 4.0.6. Let U and V be G -modules. A linear map $\phi : U \rightarrow V$ is called a G -module morphism or a G -linear map if $\phi(gu) = g\phi(u)$ for all $u \in U$ and $g \in G$. If ϕ is invertible, then it is called an isomorphism of G -modules. The linear space of all G -linear maps from U to V is denoted $\mathrm{Hom}(U, V)^G$.

The multilinear algebra constructions from Section 3 carry over to representations. For instance, if $\rho : G \rightarrow \mathrm{GL}(U)$ and $\sigma : G \rightarrow \mathrm{GL}(V)$ are representations, then the map $\rho \otimes \sigma : G \rightarrow \mathrm{GL}(U \otimes V)$, $g \mapsto \rho(g) \otimes \sigma(g)$ is also a representation. Similarly, for any natural number k the map $g \mapsto S^k \rho(g)$ is a representation of G on $S^k V$. Also, the dual space V^* of all linear functions on V carries a natural G -module structure: for $f \in V^*$ and $g \in G$ we let gf be the linear function defined by $gf(v) = f(g^{-1}v)$. The inverse ensures that the action is a left action rather than a right action: for $g, h \in G$ and $v \in V$ we have

$$(g(hf))(v) = (hf)(g^{-1}v) = f(h^{-1}g^{-1}v) = f((gh)^{-1}v) = ((gh)f)(v),$$

so that $g(hf) = (hg)f$.

Exercise 4.0.7. Show that the set of fixed points in $\mathrm{Hom}(U, V)$ under the action of G is precisely $\mathrm{Hom}(U, V)^G$.

Example 4.0.8. Let V, U be G -modules. Then the space $\mathrm{Hom}(V, U)$ of linear maps $V \rightarrow U$ is a G module with the action $(g\phi)(v) := g\phi(g^{-1}v)$. The space

$U \otimes V^*$ is also a G -module with action determined by $g(u \otimes f) = (gu) \otimes (gf)$. The natural map $\Psi : U \otimes V^* \rightarrow \text{Hom}(V, U)$ determined by $\Psi(u \otimes f)v = f(v)u$ is a morphism of G -modules. To check this it suffices to observe that

$$\Psi(g(u \otimes f))v = \Psi((gu) \otimes (gf))v = (gf)(v) \cdot gu = f(g^{-1}v) \cdot gu$$

and

$$(g\Psi(u \otimes f))v = g\Psi(u \otimes f)(g^{-1}v) = g(f(g^{-1}v)u) = f(g^{-1}v) \cdot gu.$$

The map Ψ is an G -module isomorphism of $U \otimes V^*$ with the space of finite-rank linear maps from V to U . In particular, if U or V is finite-dimensional, then Ψ is an isomorphism.

Example 4.0.9. Let G be a group acting on a set X . Consider the vectorspace

$$\mathbb{C}X := \left\{ \sum_{x \in X} \lambda_x x \mid \lambda_x \in \mathbb{C} \text{ for all } x \in X \text{ and } \lambda_x = 0 \text{ for almost all } x \right\} \quad (4.1)$$

consisting of all formal finite linear combinations of elements from X . The natural action of G given by $g(\sum_x \lambda_x x) := \sum_x \lambda_x gx$ makes $\mathbb{C}X$ into a G module. In the special case where $X = G$ and G acts on itself by multiplication on the left, the module $\mathbb{C}G$ is called the *regular representation of G* .

Definition 4.0.10. A G -submodule of a G -module V is a G -stable subspace, that is, a subspace U such that $gU \subseteq U$ for all $g \in G$. The *quotient* V/U then carries a natural structure of G -module, as well, given by $g(v + U) := (gv) + U$.

Definition 4.0.11. A G -module V is called *irreducible* if it has precisely two G -submodules (namely, 0 and V).

Exercise 4.0.12. Show that for finite groups G , every irreducible G -module has finite dimension.

Note that, just like 1 is not a prime number and the empty graph is not connected, the zero module is not irreducible. In this course we will be concerned only with G -modules that are either finite-dimensional or *locally finite*.

Definition 4.0.13. A G -module V is called *locally finite* if every $v \in V$ is contained in a finite-dimensional G -submodule of V .

Proposition 4.0.14. For a locally finite G -module V the following statements are equivalent.

1. for every G -submodule U of V there is a G -submodule W of V such that $U \oplus W = V$;
2. V is a (potentially infinite) direct sum of finite-dimensional irreducible G -submodules.

In this case we call V *completely reducible*; note that we include that condition that V be locally finite in this notion.

Proof. First assume (1). Let $\mathcal{M}$ be the collection of all finite-dimensional irreducible G -submodules of V . The collection of subsets S of $\mathcal{M}$ for which the sum $\sum_{U \in S} U$ is direct satisfies the condition of Zorn's Lemma: the union of any chain of such subsets S is again a subset of $\mathcal{M}$ whose sum is direct. Hence by Zorn's Lemma there exists a maximal subset S of $\mathcal{M}$ whose sum is direct. Let U be its (direct) sum, which is a G -submodule of V . By (1) U has a direct complement W , which is also a G -submodule. If W is non-zero, then it contains a non-zero finite-dimensional submodule (since it is locally finite), and for dimension reasons the latter contains an irreducible G -submodule W' . But then $S \cup \{W'\}$ is a subset of $\mathcal{M}$ whose sum is direct, contradicting maximality of S . Hence $W = 0$ and $V = U = \bigoplus_{M \in S} M$, which proves (2).

For the converse, assume (2) and write V as the direct sum $\bigoplus_{M \in S} M$ of irreducible finite-dimensional G -modules. Let U be any submodule of V . Then the collections of subsets of S whose sum intersects U only in 0 satisfies the condition of Zorn's Lemma. Hence there is a maximal such subset S' . Let W be its sum. We claim that $U + W = V$ (and the sum is direct by construction). Indeed, if not, then some element M of S is not contained in $U + W$. But then $M \cap (U + W) = \{0\}$ by irreducibility of M and therefore the sum of $S' \cup \{M\}$ still intersects U trivially, contradicting the maximality of S' . This proves (1). $\square$

Remark 4.0.15. It is not hard to prove that direct sums, submodules, and quotients of locally finite G -modules are again locally finite, and that they are also completely reducible if the original modules were.

Example 4.0.16. A typical example of a module which is not completely reducible is the following. Let G be the group of invertible upper triangular 2×2 -matrices, and let $V = \mathbb{C}^2$. Then the subspace spanned by the first standard basis vector e_1 is a G -submodule, but it does not have a direct complement that is G -stable.

Note that the group in this example is infinite. This is not a coincidence, as the following fundamental results show.

Proposition 4.0.17. *Let G be a finite group and let V be a finite-dimensional G -module. Then there exists a Hermitian inner product $(\cdot|\cdot)$ on V such that $(gu|gv) = (u|v)$ for all $g \in G$ and $u, v \in V$.*

Proof. Let $(\cdot|\cdot)'$ be any Hermitian inner product on V and take

$$(u|v) := \sum_{g \in G} (gu|gv)'.$$

Straightforward computations shows that $(\cdot|\cdot)$ is G -invariant, linear in its first argument, and semilinear in its second argument. For positive definiteness, note that for $v \neq 0$ the inner product $(v|v) = \sum_{g \in G} (gv|gv)$ is positive since every entry is positive. $\square$

Theorem 4.0.18. *For a finite group G any G -module is completely reducible.*

Proof. Let V be a G -module. Then every $v \in V$ lies in the finite-dimensional subspace spanned by its orbit $Gv = \{gv \mid g \in G\}$, which moreover is G -stable. Hence V is locally finite. By Zorn's lemma there exists a submodule U of V which is maximal among all direct sums of finite-dimensional irreducible submodules of V . If U is not all of V , then let W be a finite-dimensional submodule of V not contained in U , and let $(\cdot|\cdot)$ be a G -invariant Hermitian form on W . Then $U \cap W$ is a G -submodule of W , and therefore so is the orthogonal complement $(U \cap W)^\perp$ of $U \cap W$ in W —indeed, one has $(gw|U \cap W) = (w|g^{-1}(U \cap W)) \subseteq (w|U \cap W) = \{0\}$ for $g \in G$ and $w \in (U \cap W)^\perp$, so that $gw \in (U \cap W)^\perp$. Let W' be an irreducible submodule of $(U \cap W)^\perp$. Then $U \oplus W'$ is a larger submodule of V which is the direct sum of irreducible submodules of V , a contradiction. Hence $V = U$ is completely reducible. $\square$

4.1 Schur's lemma and isotypic decomposition

The following easy observation due to the German mathematician Issai Schur (1875-1941) is fundamental to representation and invariant theory.

Lemma 4.1.1 (Schur's Lemma). *Let V and U be irreducible finite-dimensional G modules for some group G . Then either V and U are isomorphic and $\text{Hom}(V, U)^G$ is one-dimensional, or they are not isomorphic and $\text{Hom}(V, U)^G = \{0\}$.*

Proof. Suppose that $\text{Hom}(V, U)^G$ contains a non-zero element ϕ . Then $\ker \phi$ is a G -submodule of V unequal to all of V and hence equal to $\{0\}$. Also, $\text{im } \phi$ is a G -submodule of U unequal to $\{0\}$, hence equal to U . It follows that ϕ is an isomorphism of G -modules. Now suppose that ϕ' is a second element of $\text{Hom}(V, U)^G$. Then $\psi := \phi' \circ \phi^{-1}$ is a G -morphism from U to itself; let $\lambda \in \mathbb{C}$ be an eigenvalue of it. Then $\psi - \lambda I$ is a G -morphism from U to itself, as well, and its kernel is a non-zero submodule, hence all of U . This shows that $\psi = \lambda I$ and therefore $\phi' = \lambda \phi$. Hence $\text{Hom}(V, U)^G$ is one-dimensional, as claimed. $\square$

If G is a group and V is a completely reducible G -module, then the decomposition of V as a direct sum of irreducible G -modules need not be unique. For instance, if V is the direct sum $U_1 \oplus U_2 \oplus U_3$ where the first two are isomorphic irreducible modules and the third is an irreducible module not isomorphic to the other two, then V can also be written as $U_1 \oplus \Delta \oplus U_3$, where $\Delta = \{u_1 + \phi(u_1) \mid u_1 \in U_1\}$ is the diagonal subspace of $U_1 \oplus U_2$ corresponding to an isomorphism ϕ from U_1 to U_2 .

However, there is always a coarser decomposition of V into G -modules which is unique. For this, let $\{U_i\}_{i \in I}$ be a set of representatives of the isomorphism classes of G -modules, so that every irreducible finite-dimensional G -module is isomorphic to U_i for some unique $i \in I$. For every $i \in I$ let V_i be the (non-direct) sum of all G -submodules of V that are isomorphic to U_i . Clearly each V_i is a G -submodule of V and, since V is a direct sum of irreducible G -modules, $\sum_{i \in I} V_i = V$. Using Zorn's lemma one sees that V_i can also be written as $\bigoplus_{j \in J_i} V_{ij}$ for irreducible submodules V_{ij} , $j \in J_i$ that are all isomorphic to U_i .

We claim that $V = \bigoplus_{i \in I} V_i$. To see this, suppose that $V_{i_0} \cap \sum_{i \neq i_0} V_i \neq \{0\}$ and let U be an irreducible submodule of this module. Then the projection of U onto some $V_{i_0, j}$ along the direct sum of the remaining direct summands of V_{i_0} is non-zero, and similarly the projection of U onto $V_{i_1, j}$ for some i_1, j along the remaining summands of V_{i_1} is non-zero. By Schur's lemma U is then both isomorphic to $V_{i_0, j}$ and to $V_{i_1, j}$, a contradiction. Hence $V_{i_0} \cap \sum_{i \neq i_0} V_i$ is zero, as claimed.

The space V_i is called the *isotypic component of V of type U_i* , and it has the following pretty description. The map $\text{Hom}(U_i, V)^G \times U_i \rightarrow V$, $(\phi, u) \mapsto \phi(u)$ is bilinear, and therefore gives rise to a linear map $\Psi : \text{Hom}(U_i, V)^G \otimes U_i \rightarrow V$. This linear map is a linear isomorphism onto V_i .

Exercise 4.1.2. Let U, V, W be G -modules. Show that $\text{Hom}(U \oplus V, W)^G \cong \text{Hom}(U, W) \oplus \text{Hom}(V, W)$ and $\text{Hom}(W, U \oplus V) \cong \text{Hom}(W, U) \oplus \text{Hom}(W, V)$.

4.2 Exercises

Exercise 4.2.1. • Let V be a G -module and $\langle, \rangle$ a G -invariant inner product on V . Show that for any two non-isomorphic, irreducible submodules $V_1, V_2 \subset V$ we have $V_1 \perp V_2$, that is, $\langle v_1, v_2 \rangle = 0$ for all $v_1 \in V_1, v_2 \in V_2$.

- Give an example where $V_1 \not\perp V_2$ for (isomorphic) irreducible G -modules V_1 and V_2 .

Exercise 4.2.2. Let the symmetric group on 3 letters S_3 act on $\mathbb{C}[x_1, x_2, x_3]_2$ by permuting the variables. This action makes $\mathbb{C}[x_1, x_2, x_3]_2$ into a S_3 -module. Give a decomposition of this module into irreducible submodules.

Exercise 4.2.3. Let G be an abelian group. Show that every irreducible G -module has dimension 1. Show that G has a *faithful* irreducible representation if and only if G is cyclic. A representation ρ is called faithful if it is injective.

Exercise 4.2.4. Let G be a finite group and V an irreducible G -module. Show that there is a *unique* G -invariant inner product on V , unique up to multiplication by scalars.

Exercise 4.2.5. Let G be a finite group, and let $\mathbb{C}G$ be the regular representation of G and let $\mathbb{C}G = W_1^{m_1} \oplus \cdots \oplus W_k^{m_k}$ be the isotypic decomposition of $\mathbb{C}G$. Show that for every irreducible G -module W , there is an i such that W is isomorphic to W_i and show that $m_i = \dim W_i$. Hint: for all $w \in W$ the linear map $\mathbb{C}G \rightarrow W$ given by $\sum_g \lambda_g g \mapsto \sum_g \lambda_g gw$ is a G -linear map.

Chapter 5

Finite generation of the invariant ring

In all examples we have met so far, the invariant ring was generated by a finite number of invariants. In this section we prove Hilbert's theorem that under reasonable conditions, this is always the case. For the proof we will need another theorem by Hilbert.

Recall that for a ring R and a subset $S \subseteq R$, the *ideal generated by S* is defined as

$$(S) := \{r_1 s_1 + \cdots + r_k s_k \mid k \in \mathbb{N}, r_1, \dots, r_k \in R, s_1, \dots, s_k \in S\}. \quad (5.1)$$

You may want to check that this indeed defines an ideal in R . An ideal $I \subseteq R$ is called *finitely generated* if there is a finite set S such that $I = (S)$.

Definition 5.0.6. A ring R is called *Noetherian* if every ideal I in R is finitely generated.

Exercise 5.0.7. Show that a ring R is Noetherian if and only if there is no infinite ascending chain of ideals $I_1 \subsetneq I_2 \subsetneq I_3 \subsetneq \cdots$.

We will be mostly interested in polynomial rings over $\mathbb{C}$ in finitely many indeterminates, for which the following theorem is essential.

Theorem 5.0.8 (Hilbert's Basis Theorem). *The polynomial ring $\mathbb{C}[x_1, \dots, x_n]$ is Noetherian.*

We will deduce this statement from the following result.

Lemma 5.0.9 (Dixon's Lemma). *If $m_1, m_2, m_3, \dots$ is an infinite sequence of monomials in the variables $x_1, \dots, x_n$, then there exist indices $i < j$ such that $m_i \mid m_j$.*

Proof. We proceed by induction on n . For $n = 0$ all monomials are 1, so we can take any $i < j$. Suppose that the statement is true for $n - 1 \geq 0$. Define

the infinite sequences $e_1 \leq e_2 \leq \dots$ and $i_1 < i_2 < \dots$ as follows: e_1 is the smallest exponent of x_n in any of the monomials m_i , and i_1 is the smallest index i for which the exponent of x_n in m_i equals e_1 . For $k > 1$ the exponent e_k is the smallest exponent of x_n in any of the m_i with $i > i_{k-1}$ and i_k is the smallest index $i > i_{k-1}$ for which the exponent of x_n in m_i equals e_k . Now the monomials in the sequence $m_{i_1}/x_n^{e_1}, m_{i_2}/x_n^{e_2}, \dots$ do not contain x_n . Hence by induction there exist $j < l$ such that $m_{i_j}/x_n^{e_j} \mid m_{i_l}/x_n^{e_l}$. As $e_j \leq e_l$ we then also have $m_{i_j} \mid m_{i_l}$, and of course $i_j < i_l$, as claimed. $\square$

Proof of Hilbert's Basis Theorem. Let $I \subseteq \mathbb{C}[x_1, \dots, x_n]$ be an ideal. For any polynomial f in $\mathbb{C}[x_1, \dots, x_n]$ we denote by $\text{lm}(f)$ the leading monomial of f : the lexicographically largest monomial having non-zero coefficient in f . By Dixon's lemma, the set of $\mid$ -minimal monomials in $\{\text{lm}(f) \mid f \in I\}$ is finite. Hence there exist finitely many polynomials $f_1, \dots, f_k \in I$ such that for all $f \in I$ there exists an i with $\text{lm}(f_i) \mid \text{lm}(f)$. We claim that the ideal $J := (f_1, \dots, f_k)$ generated by the f_i equals I . If not, then take an $f \in I \setminus J$ with the lexicographically smallest leading monomial among all counter examples. Take i such that $\text{lm}(f_i) \mid \text{lm}(f)$, say $\text{lm}(f) = m \text{lm}(f_i)$. Subtracting a suitable scalar multiple of $m f_i$, which lies in J , from f gives a polynomial with a lexicographically smaller leading monomial, and which is still in $I \setminus J$. But this contradicts the minimality of $\text{lm}(f)$. $\square$

Remark 5.0.10. More generally, Hilbert showed that for R Noetherian, also $R[x]$ is Noetherian (which you may want to prove yourself!). Since clearly any field is a Noetherian ring, this implies the previous theorem by induction on the number of indeterminates.

With this tool in hand, we can now return to our main theorem of this section.

Theorem 5.0.11 (Hilbert's Finiteness Theorem). *Let G be a group and let W be a finite dimensional G -module with the property that $\mathbb{C}[W]$ is completely reducible. Then $\mathbb{C}[W]^G := \{f \in \mathbb{C}[W] \mid gf = f\}$ is a finitely generated subalgebra of $\mathbb{C}[W]$. That is, there exist $f_1, \dots, f_k \in \mathbb{C}[W]^G$ such that every G -invariant polynomial on W , is a polynomial in the f_i .*

The proof uses the so-called *Reynolds operator* ρ , which is defined as follows. We assume that the vector space $\mathbb{C}[W]$ is completely reducible. Consider its isotypic decomposition $\mathbb{C}[W] = \bigoplus_{i \in I} V_i$ and let $1 \in I$ correspond to the trivial 1-dimensional G -module, so that $\mathbb{C}[W]^G = V_1$. Now let ρ be the projection from $\mathbb{C}[W]$ onto V_1 along the direct sum of all V_i with $i \neq 1$. This is a G -equivariant linear map. Moreover, we claim that

$$\rho(f \cdot h) = f \cdot \rho(h) \text{ for all } f \in V_1, \quad (5.2)$$

where the multiplication is multiplication in $\mathbb{C}[W]$. Indeed, consider the map $\mathbb{C}[W] \rightarrow \mathbb{C}[W]$, $h \mapsto fh$. This is a G -module morphism, since $g(fh) = (gf) \cdot (gh) = f \cdot (gh)$, where the first equality reflects that G acts by automorphisms

on $\mathbb{C}[W]$ and the second equality follows from the invariance of f . Hence if we write h as $\sum_i h_i$ with $h_i \in V_i$, then $fh_i \in V_i$ by Schur's lemma, and therefore the component of $fh = \sum_i (fh_i)$ in V_1 is just fh_1 . In other words $\rho(fh) = f\rho(h)$, as claimed.

Exercise 5.0.12. Show that for a finite group G , the Reynolds operator is just $f \mapsto \frac{1}{|G|} \sum_{g \in G} gf$.

Proof of Hilbert's finiteness theorem. Let $I' := \bigoplus_{d>0} \mathbb{C}[W]_d^G$ be the ideal in $\mathbb{C}[W]^G$ consisting of all invariants with zero constant term. Denote by $I := \mathbb{C}[W]I'$ the ideal in $\mathbb{C}[W]$ generated by I' . Since W is finite dimensional, it follows from Hilbert's basis theorem that there exist $f_1, \dots, f_k \in I$ that generate the ideal I . We may assume that the f_i belong to I' . Indeed, if $f_i \notin I'$, we can write $f_i = \sum_j f_{ij}g_{ij}$ for certain $f_{ij} \in I'$ and $g_{ij} \in \mathbb{C}[W]$ and replace f_i with the f_{ij} to obtain a finite generating set of I with fewer elements in $I \setminus I'$.

We observe that the ideal I' is generated by the f_i . Indeed, let $h \in I' \subseteq I$ and write $h = \sum_i g_i f_i$ for some $g_i \in \mathbb{C}[W]$. Using the Reynolds operator ρ we find: $h = \rho(h) = \sum_i \rho(g_i f_i) = \sum_i f_i \rho(g_i)$.

The proof is now completed by exercise 5.0.13. $\square$

Exercise 5.0.13. Let $A \subseteq \mathbb{C}[W]$ be a subalgebra, and let $A^+ := \bigoplus_{d \geq 1} A \cap \mathbb{C}[W]_d$ be the ideal of polynomials in A with zero coefficient. Suppose that the ideal A^+ is finitely generated. Show that A is finitely generated as an algebra over $\mathbb{C}$.

5.1 Noethers degree bound

For a *finite* group G , any G -module V is completely reducible as we have seen in the previous lecture. This implies by Hilbert's theorem that for finite groups, the invariant ring is always finitely generated. In this section, we prove a result of Noether stating that for finite groups G , the invariant ring is already generated by the invariants of degree at most $|G|$, which implies a bound on the number of generators needed.

Theorem 5.1.1 (Noether's degree bound). *Let G be a finite group, and let W be a (finite dimensional) G -module. Then the invariant ring $\mathbb{C}[W]^G$ is generated by the homogeneous invariants of degree at most $|G|$.*

Proof. We choose a basis $x_1, \dots, x_n$ of W^* so that $\mathbb{C}[W] = \mathbb{C}[x_1, \dots, x_n]$. For any n -tuple $\alpha = (\alpha_1, \dots, \alpha_n)$ of nonnegative integers we have an invariant

$$j_\alpha := \sum_{g \in G} g(x_1^{\alpha_1} \cdots x_n^{\alpha_n}) \quad (5.3)$$

homogeneous of degree $|\alpha| := \alpha_1 + \cdots + \alpha_n$. Clearly, the invariants j_α span the vector space $\mathbb{C}[W]^G$, since for any invariant $f = \sum_\alpha c_\alpha x_1^{\alpha_1} \cdots x_n^{\alpha_n}$, we have

$$f = \frac{1}{|G|} \sum_{g \in G} gf = \frac{1}{|G|} \sum_\alpha c_\alpha j_\alpha. \quad (5.4)$$

It will therefore suffice to prove that every j_α is a polynomial in the j_β with $|\beta| \leq |G|$.

Let $z_1, \dots, z_n$ be n new variables and define for $j \in \mathbb{N}$ the polynomials

$$p_j(x_1, \dots, x_n, z_1, \dots, z_n) := \sum_{g \in G} (gx_1 \cdot z_1 + \dots + gx_n \cdot z_n)^j. \quad (5.5)$$

So these are the Newton polynomials (see Lecture 2), where we have substituted the expressions $(gx_1 \cdot z_1 + \dots + gx_n \cdot z_n)$ for the $|G|$ variables. Expanding p_j and sorting terms with respect to the variables z_i , we see that $p_j = \sum_{|\alpha|=j} f_\alpha z_1^{\alpha_1} \dots z_n^{\alpha_n}$, where

$$f_\alpha = \binom{j}{\alpha_1, \dots, \alpha_n} j_\alpha. \quad (5.6)$$

Now let $j > |G|$. Recall that p_j is a polynomial in $p_1, \dots, p_{|G|}$. This implies that also the coefficients $f_\alpha, |\alpha| = j$ of p_j are polynomials in the coefficients $f_\beta, |\beta| \leq |G|$ of $p_1, \dots, p_{|G|}$. This finishes the proof, since $\binom{j}{\alpha_1, \dots, \alpha_n} \neq 0$ when $\alpha_1 + \dots + \alpha_n = j$. $\square$

Exercise 5.1.2. Show that for all cyclic groups, the bound in the theorem is met in some representation.

5.2 Exercises

For a finite group G , define $\beta(G)$ to be the minimal number m such that for every (finite dimensional) G -module W , the invariantring $\mathbb{C}[W]^G$ is generated by the invariants of degree at most m . By Noether's theorem, we always have $\beta(G) \leq |G|$.

Exercise 5.2.1. Let G be a finite abelian group. We use additive notation. Define the Davenport constant $\delta(G)$ to be the maximum length m of a non-shortable expression $0 = g_1 + \dots + g_m, g_1, \dots, g_m \in G$. Non-shortable means that there is no strict non-empty subset I of $\{1, \dots, m\}$ such that $\sum_{i \in I} g_i = 0$. Show that $\delta(G) = \beta(G)$. Compute $\delta((\mathbb{Z}/2\mathbb{Z})^n)$.

Chapter 6

Affine varieties and the Quotient Map

6.1 Affine varieties

Definition 6.1.1. An *affine variety* is a subset of some $\mathbb{C}^n$ which is the common zero set of a collection of polynomials in the coordinates $x_1, \dots, x_n$ on $\mathbb{C}^n$.

Suppose that S is a subset of $\mathbb{C}[\underline{x}] := \mathbb{C}[x_1, \dots, x_n]$ and let $p \in \mathbb{C}^n$ be a common zero of the elements of S . Then any finite combination $\sum_i a_i f_i$ where the f_i are in S and the a_i are in $\mathbb{C}[x]$ also vanishes on p . The collection of all such polynomials is the ideal generated by S . So the study of affine varieties leads naturally to the study of ideals in the polynomial ring $\mathbb{C}[x_1, \dots, x_n]$. We have seen in Week 5 that such ideals are always finitely generated.

Exercise 6.1.2. Show that the collection of affine varieties in $\mathbb{C}^n$ satisfy the following three properties:

1. $\mathbb{C}^n$ and $\emptyset$ are affine varieties;
2. the union of two affine varieties is an affine variety; and
3. the intersection of arbitrarily many affine varieties is an affine variety.

These conditions say that the affine varieties in $\mathbb{C}^n$ form the closed subsets in a *topology* on $\mathbb{C}^n$. This topology is called the *Zariski topology*, after the Polish-American mathematician Otto Zariski (1899-1986). We will interchangeably use the terms affine (sub)variety in $\mathbb{C}^n$ and Zariski-closed subset of $\mathbb{C}^n$. Moreover, in the last case we will often just say closed subset; when we mean closed subset in the Euclidean sense rather than in the Zariski-sense, we will explicitly mention that.

Exercise 6.1.3. The Zariski-topology on $\mathbb{C}^n$ is very different from the Euclidean topology on $\mathbb{C}^n$, as the answers to the following problems show:

1. determine the Zariski-closed subsets of $\mathbb{C}$;
2. prove that $\mathbb{R}^n$ is Zariski-dense in $\mathbb{C}^n$ (that is, the smallest Zariski-closed subset of $\mathbb{C}^n$ containing $\mathbb{R}^n$ is $\mathbb{C}^n$ itself); and
3. show that every non-empty Zariski-open subset of $\mathbb{C}^n$ (that is, the complement of a Zariski-closed set) is dense in $\mathbb{C}^n$.

On the other hand, in some other aspects the Zariski topology resembles the Euclidean topology:

1. show that Zariski-open subsets of $\mathbb{C}^n$ are also open in the Euclidean topology;
2. determine the image of the map $\phi : \mathbb{C}^2 \rightarrow \mathbb{C}^3$, $(x_1, x_2) \rightarrow (x_1, x_1x_2, x_1(1+x_2))$, and show that its Zariski closure coincides with its Euclidean closure.

If you solved the last exercise correctly, then you found that the image is some Zariski-closed subset minus some Zariski-closed subset plus some other Zariski closed subset. In general, the subsets of $\mathbb{C}^n$ that are generated by the Zariski-closed sets under (finitely many of) the operations $\cup, \cap$, and complement, are called *constructible sets*. An important result due to the French mathematician Claude Chevalley (1909-1984) says that the image of a constructible set under a polynomial map $\mathbb{C}^n \rightarrow \mathbb{C}^m$ is again a constructible set. Another important fact is that the Euclidean closure of a constructible set equals its Zariski closure.

From undergraduate courses we know that $\mathbb{C}$ is an *algebraically closed field*, that is, that every non-constant univariate polynomial $f \in \mathbb{C}[x]$ has a root. The following multivariate analogue of this statement is the second major theorem of Hilbert's that we will need.

Theorem 6.1.4 (Hilbert's weak Nullstellensatz). *Let I be an ideal in $\mathbb{C}[\underline{x}]$ that is not equal to all of $\mathbb{C}[\underline{x}]$. Then there exists a point $\xi = (\xi_1, \dots, \xi_n)$ such that $f(\xi) = 0$ for all $f \in I$.*

The theorem is also true with $\mathbb{C}$ replaced by any other algebraically closed field. But we will give a self-contained proof that uses the fact that $\mathbb{C}$ is not countable.

Lemma 6.1.5. *Let U, V be vector spaces over $\mathbb{C}$ of countably infinite dimension, let $A(x) : U \otimes \mathbb{C}[x] \rightarrow V \otimes \mathbb{C}[x]$ be a $\mathbb{C}[x]$ -linear map, and let $v(x) \in V \otimes \mathbb{C}[x]$ be a target vector. Suppose that for all $\xi \in \mathbb{C}$ there is a $u \in U$ such that $A(\xi)u = v(\xi)$. Then there exists a $u(x) \in U \otimes \mathbb{C}(x)$ such that $Au(x) = v(x)$.*

Proof. Suppose, on the contrary, that no such $u(x)$ exists. This means that the image under $A(x)$ of the $\mathbb{C}(x)$ -vector space $U \otimes \mathbb{C}(x)$ does not contain $v(x)$. Let $F(x)$ be a $\mathbb{C}(x)$ -linear function on $V \otimes \mathbb{C}(x)$ taking the value 0 on $A(U \otimes \mathbb{C}(x))$ and 1 on $v(x)$; such a function exists and is determined by its values $f_1(x), f_2(x), f_3(x), \dots \in \mathbb{C}(x)$ on a $\mathbb{C}$ -basis $v_1, v_2, v_3, \dots$ of V . Since $\mathbb{C}$ is uncountable there is a value $\xi \in \mathbb{C}$ where all f_i are defined, so that $F(\xi)$

is a well-defined linear function on V . Now we have $F(\xi)A(\xi)u = 0$ for all $u \in U$ but $F(\xi)v(\xi) = 1$, contradicting the assumption that $A(\xi)u = v(\xi)$ has a solution. $\square$

Proof of the weak Nullstellensatz. We proceed by induction on n . For $n = 0$ the statement is just that any proper ideal of $\mathbb{C}$ is 0. Now suppose that $n > 0$ and that the statement is true for $n - 1$. By Hilbert's basis theorem, the ideal I is generated by finitely many polynomials $f_1, \dots, f_k$. If there exists a value $\xi \in \mathbb{C}$ for x_n such that the ideal in $\mathbb{C}[x_1, \dots, x_{n-1}]$ generated by $f_{1,\xi} := f_1(x_1, \dots, x_{n-1}, \xi), \dots, f_{k,\xi} := f_k(x_1, \dots, x_{n-1}, \xi)$ does not contain 1, then we can use the induction hypothesis and we are done. Suppose therefore that no such ξ exists, that is, that 1 can be written as a $\mathbb{C}[x_1, \dots, x_{n-1}]$ -linear combination

$$1 = \sum_j c_{j,\xi} f_{j,\xi}$$

for every $\xi \in \mathbb{C}$. We will use this fact in two ways. First, note that this means that

$$\sum_j c_{j,\xi} f_j = 1 + (x_n - \xi)g_\xi$$

for some polynomial $g_\xi \in \mathbb{C}[x_1, \dots, x_n]$. Put differently, $(x_n - \xi)$ has a *multiplicative inverse modulo the ideal I* for each $\xi \in \mathbb{C}$. But then every univariate polynomial in x_n , being a product of linear ones since $\mathbb{C}$ is algebraically closed, has such a multiplicative inverse. Since 1 does not lie in I , this implies that $I \cap \mathbb{C}[x_n] = \{0\}$.

Second, by Lemma 6.1.5 applied to $U = \mathbb{C}[x_1, \dots, x_{n-1}]^k$, $V = \mathbb{C}[x_1, \dots, x_{n-1}]$, $x = x_n$, and $A(c_1, \dots, c_k) = \sum_{i=1}^k c_i f_i$, we can write

$$1 = \sum_{i=1}^k c_i(x_n) f_i,$$

where each $c_i(x_n)$ lies in $\mathbb{C}[x_1, \dots, x_{n-1}](x_n)$. Letting $D(x_n) \in \mathbb{C}[x_n] \setminus 0$ be a common denominator of the c_i and setting $c'_i := D c_i \in \mathbb{C}[x_1, \dots, x_n]$, we find that

$$D(x_n) = \sum_{i=1}^k c'_i f_i \in I.$$

But this contradicts our earlier conclusion that I does not contain non-zero polynomials in x_n only. $\square$

The Nullstellensatz has many applications to combinatorial problems.

Exercise 6.1.6. Let $G = (V, E)$ be a finite, undirected graph with vertex set V and edge set $E \subseteq \binom{V}{2}$. A *proper k -colouring* of G is a map $c : V \rightarrow [k]$ with the property that $c(i) \neq c(j)$ whenever $\{i, j\} \in E$. To G we associate

the polynomial ring $\mathbb{C}[x_i \mid i \in V]$ and its ideal I generated by the following polynomials:

$$x_i^k - 1 \text{ for all } i \in V; \text{ and } x_i^{k-1} + x_i^{k-2}x_j + \dots + x_j^{k-1} \text{ for all } \{i, j\} \in E.$$

Prove that G has a proper k -colouring if and only if $1 \notin I$.

Two important maps set up a beautiful duality between geometry and algebra. First, we have the map $\mathcal{V}$ that sends a subset $S \subseteq \mathbb{C}[\underline{x}]$ to the variety $\mathcal{V}(S)$ that it defines; and second, the map $\mathcal{I}$ that sends a subset $X \subseteq \mathbb{C}^n$ to the ideal $\mathcal{I}(X) \subseteq \mathbb{C}[\underline{x}]$ of all polynomials that vanish on all points in X . The following properties are straightforward:

1. if $S \subseteq S'$ then $\mathcal{V}(S) \supseteq \mathcal{V}(S')$;
2. if $X \subseteq X'$ then $\mathcal{I}(X) \supseteq \mathcal{I}(X')$;
3. $X \subseteq \mathcal{V}(\mathcal{I}(X))$;
4. $S \subseteq \mathcal{I}(\mathcal{V}(S))$;
5. $\mathcal{V}(\mathcal{I}(\mathcal{V}(S))) = \mathcal{V}(S)$; and
6. $\mathcal{I}(\mathcal{V}(\mathcal{I}(X))) = \mathcal{I}(X)$.

For instance, in (5) the containment $\supseteq$ follows from (3) applied to $X = \mathcal{V}(S)$ and the containment $\subseteq$ follows from (4) applied to S and then (1) applied to $S \subseteq S' := \mathcal{I}(\mathcal{V}(S))$.

This shows that $\mathcal{V}$ and $\mathcal{I}$ set up an inclusion-reversing bijection between sets of the form $\mathcal{V}(S) \subseteq \mathbb{C}^n$ —that is, affine varieties in $\mathbb{C}^n$ —and sets of the form $\mathcal{I}(X) \subseteq \mathbb{C}[\underline{x}]$. Sets of the latter form are always ideals, but not all ideals are of this form, as the following example shows.

Example 6.1.7. Suppose that $n = 1$, fix a natural number k , and let I_k be the ideal in $\mathbb{C}[x_1]$ generated by x_1^k . Then $\mathcal{V}(I) = \{0\}$ and $\mathcal{I}(\mathcal{V}(I))$ is the ideal generated by x_1 . So for $k > 1$ the ideal I_k is not of the form $\mathcal{I}(X)$ for any subset of $\mathbb{C}^n$.

This example exhibits a necessary condition for an ideal to be of the form $\mathcal{I}(X)$ for some set X —it must be *radical*.

Definition 6.1.8. The *radical* of an ideal $I \subseteq \mathbb{C}[\underline{x}]$ is the set of all polynomials f of which some positive power lies in I ; it is denoted $\sqrt{I}$. The ideal I is called *radical* if $I = \sqrt{I}$.

Indeed, suppose that $I = \mathcal{I}(X)$ and suppose that $f \in \mathbb{C}[\underline{x}]$ has $f^k \in I$ for some $k > 0$. Then f^k vanishes on X and hence so does f , and hence $f \in \mathcal{I}(X) = I$. This shows that I is radical.

Exercise 6.1.9. Show that, for general ideals I , $\sqrt{I}$ is an ideal containing I .

The second important result of Hilbert's that we will need is that the condition that I be radical is also *sufficient* for I to be the vanishing ideal of some set X .

Theorem 6.1.10 (Hilbert's Nullstellensatz). *Suppose that $I \subseteq \mathbb{C}[\underline{x}]$ is a radical ideal. Then $\mathcal{I}(\mathcal{V}(I)) = I$.*

Proof. This follows from the weak Nullstellensatz using *Rabinowitsch's trick* from 1929. Let g be a polynomial vanishing on all common roots of the polynomials in I . Introducing an auxiliary variable t , we have that the ideal in $\mathbb{C}[\underline{x}, t]$ generated by I and $tg - 1$ does not have any common zeroes. Hence by the weak Nullstellensatz 1 can be written as

$$1 = a(tg - 1) + \sum_{i=1}^k c_i(\underline{x}, t)f_i, \quad a, c_i \in \mathbb{C}[\underline{x}, t], \quad f_i \in I.$$

Replacing t on both sides by $1/g$ we have

$$1 = \sum_j c_j(\underline{x}, 1/g)f_j.$$

Multiplying both sides with a suitable power g^d eliminates g from the denominators and hence expresses g^d as a $\mathbb{C}[\underline{x}]$ -linear combination of the f_j . Hence $g^d \in I$ and therefore $f \in I$ since I is radical. $\square$

We have thus set up an inclusion-reversing bijection between closed subsets of $\mathbb{C}^n$ and radical ideals in $\mathbb{C}[\underline{x}]$. It is instructive to see what this bijection does with the smallest closed subsets consisting of a single point $p = (p_1, \dots, p_n) \in \mathbb{C}^n$. The ideal $\mathcal{I}(p) := \mathcal{I}(\{p\})$ of polynomials vanishing on p is generated by $x_1 - p_1, \dots, x_n - p_n$ (check this). This is a *maximal ideal* (that is, an ideal which is maximal among the *proper* ideals of $\mathbb{C}[x_1, \dots, x_n]$), since the quotient by it is the field $\mathbb{C}$. This follows from the fact that, by definition, $\mathcal{I}$ is the kernel of the homomorphism of $\mathbb{C}$ -algebras $\mathbb{C}[x_1, \dots, x_n] \rightarrow \mathbb{C}$, $f \mapsto f(p)$ and that this homomorphism is surjective. Conversely, suppose that I is a maximal ideal. Then it is radical—indeed, if the radical were strictly larger than I , it would contain 1 by maximality, but then some power of 1 would be in I , a contradiction. Hence by the Nullstellensatz there exists a non-empty subset X of $\mathbb{C}^n$ such that $I = \mathcal{I}(X)$. But then for any point p in X we have that $\mathcal{I}(p)$ is a radical ideal containing I , hence equal to I by maximality. We have thus proved the following corollary of the Nullstellensatz.

Corollary 6.1.11. *The map sending p to $\mathcal{I}(p)$ is a bijection between points in $\mathbb{C}^n$ and maximal ideals in $\mathbb{C}[\underline{x}]$.*

6.2 Regular functions and maps

Definition 6.2.1. Let X be a Zariski-closed subset of $\mathbb{C}^n$. Then a *regular function* on X is by definition a $\mathbb{C}$ -valued function of the form $f|_X$ where $f \in \mathbb{C}[\underline{x}]$. A

regular map from X to $\mathbb{C}^m$ is map $X \rightarrow \mathbb{C}^m$ of the form $x \mapsto (f_1(x), \dots, f_m(x))$ with $f_1, \dots, f_m$ regular functions on X .

Regular functions form a commutative $\mathbb{C}$ -algebra with 1, denoted $\mathbb{C}[X]$ (or sometimes $\mathcal{O}(X)$) and sometimes called the *coordinate ring* of X . By definition, $\mathbb{C}[X]$ is the image of the restriction map $\mathbb{C}[\underline{x}] \rightarrow \{\mathbb{C}\text{-valued functions on } X\}$, $f \mapsto f|_X$. Hence it is isomorphic to the quotient algebra $\mathbb{C}[\underline{x}]/\mathcal{I}(X)$.

Example 6.2.2. 1. If X is a d -dimensional subspace of $\mathbb{C}^n$, then $\mathcal{I}(X)$ is generated by the space $X^0 \subseteq (\mathbb{C}^n)^*$ of linear functions vanishing on X . If $y_1, \dots, y_d \in (\mathbb{C}^n)^*$ span a vector space complement of X^0 , then modulo $\mathcal{I}(X)$ every polynomial in the x_i is equal to a unique polynomial in the y_j . This shows that $\mathbb{C}[X] = \mathbb{C}[y_1, \dots, y_d]$ is a polynomial ring in d variables.

2. Consider the variety X of $(m+1) \times (m+1)$ -matrices of the shape

$$\begin{bmatrix} x & 0 \\ 0 & y \end{bmatrix}$$

with x an $m \times m$ -matrix and y a complex number satisfying $\det(x)y = 1$. Then $\mathbb{C}[X] = \mathbb{C}[(x_{ij})_{ij}, y]/(\det(x)y - 1)$. The map $y \mapsto 1/\det(x)$ sets up an isomorphism of this algebra with the algebra of rational functions in the variables x_{ij} generated by the x_{ij} and $1/\det(x)$. We therefore also write $\mathbb{C}[X] = \mathbb{C}[(x_{ij})_{ij}, 1/\det(x)]$. Note that X is a group with respect to matrix multiplication, isomorphic to GL_n . This is the fundamental example of an *algebraic group*; here *algebraic* refers to the variety structure of X .

3. Consider the variety $X = M_{k,m}^{\leq l}$ of all $k \times m$ -matrices all of whose $(l+1) \times (l+1)$ -minors (that is, determinants of $(l+1) \times (l+1)$ -submatrices) vanish. Elementary linear algebra shows that X consists of all matrices of rank at most l , and that such matrices can always be written as AB with $A \in M_{k,l}$, $B \in M_{l,m}$.

6.3 The quotient map

Let G be a group and let V be a finite-dimensional G -module such that $\mathbb{C}[W] = \bigoplus S^k W^*$ is a completely reducible G -module. By Hilbert's finiteness theorem, we know that the algebra $\mathbb{C}[W]^G$ of G -invariant polynomials is a finitely generated algebra. Let $f_1, \dots, f_k$ be a generating set of this algebra. Then we have a polynomial map

$$\pi : W \rightarrow \mathbb{C}^k, \quad w \mapsto (f_1(w), \dots, f_k(w)).$$

This map is called the *quotient map*, because in some sense, which will become clear below, the image of this map parameterises G -orbits in W .

Example 6.3.1. Take $G := \{-1, 1\}$ with its action on $W := \mathbb{C}^2$ where -1 sends (x, y) to $(-x, -y)$. The invariant ring $\mathbb{C}[W]^G$ is generated by the polynomials

$f_1 := x^2$, $f_2 := xy$, $f_3 := y^2$ (check this). Thus the quotient map is $\pi : \mathbb{C}^2 \rightarrow \mathbb{C}^3$, $(x, y) \mapsto (x^2, xy, y^2)$. Let u, v, w be the standard coordinates on $\mathbb{C}^3$, and note that the image of π is contained in the affine variety $Z \subseteq \mathbb{C}^3$ with equation $v^2 - uw$. Indeed, π is surjective onto Z : let u, v, w be complex numbers such that $v^2 = uw$. Let $x \in \mathbb{C}$ be a square root of u . If $x \neq 0$, then set $y := v/x$ so that $v = xy$ and $w = v^2/u = x^2 y^2 / x^2 = y^2$. If $x = 0$, then let y be a square root of w . In both cases $\pi(x, y) = (u, v, w)$. Note that the fibres of π over every point (u, v, w) are orbits of G .

Example 6.3.2. Take $G := S_n$ with its standard action on $W := \mathbb{C}^n$. Recall that the invariants are generated by the elementary symmetric polynomials $\sigma_1(x) = \sum_i x_i, \dots, \sigma_n(x) = \prod_i x_i$. This gives the quotient map

$$\pi : \mathbb{C}^n \rightarrow \mathbb{C}^n, (x_1, \dots, x_n) \mapsto (\sigma_1(x), \dots, \sigma_n(x)).$$

We claim that the image of this map is all of $\mathbb{C}^n$. Indeed, for any n -tuple $(c_1, \dots, c_n) \in \mathbb{C}^n$ consider the polynomial

$$f(t) := t^n - c_1 t^{n-1} + \dots + (-1)^n c_n.$$

As $\mathbb{C}$ is algebraically closed, this polynomial has n roots (counted with multiplicities), so that we can also write

$$f(t) = (t - x_1) \cdots (t - x_n);$$

expanding this expression for f and comparing with the above gives $\pi(x_1, \dots, x_n) = (c_1, \dots, c_n)$. Note furthermore that any other n -tuple $(x'_1, \dots, x'_n)$ with this property is necessarily some permutation of the x_i , since the roots of f are determined by $(c_1, \dots, c_n)$.

This means that we can think of $\pi : \mathbb{C}^n \rightarrow \mathbb{C}^n$ as follows: every S_n -orbit on the first copy of $\mathbb{C}^n$ is “collapsed” by π to a single point on the second copy of $\mathbb{C}^n$, and conversely, the fibre of π above any point in the second copy of $\mathbb{C}^n$ consists of a single S_n -orbit. Thus *the second copy of $\mathbb{C}^n$ parameterises S_n -orbits on $\mathbb{C}^n$* .

Example 6.3.3. Let the group $G := \mathrm{SL}_2$ act on the $W := M_2(\mathbb{C})$ of 2×2 -matrices by left multiplication. There are three types of orbits: First, $0 \in M_2(\mathbb{C})$ is an orbit by itself. Second, if $A \in M_2(\mathbb{C})$ has rank 1, then left multiplying by a suitable matrix $g \in \mathrm{SL}_2$ gives

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

Third, if A has rank 2, then left-multiplying by a suitable $g \in \mathrm{SL}_2$ gives

$$gA = \begin{bmatrix} 1 & 0 \\ 0 & \det(A) \end{bmatrix}.$$

Note that the previous equality is the special case of this where $\det(A) = 0$.

Now $\det$ is a polynomial on $M_2(\mathbb{C})$ which is SL_2 -invariant. We claim that it generates the invariant ring. Indeed, suppose $f \in \mathbb{C}[a_{11}, a_{12}, a_{21}, a_{22}]$ is any invariant. Then for A in any of the last two orbits we find that

$$f(A) = f(gA) = f(1, 0, 0, \det(A)) =: g(\det A)$$

where g is a polynomial in 1 variable. Since both f and g are continuous, we find that this equality also holds for the zero matrix. Hence $f(A) = g(\det A)$ for all A , so f is in the algebra generated by $\det A$.

In this case the quotient map $\pi : M_2(\mathbb{C}) \rightarrow \mathbb{C}$ is just the map $A \mapsto \det(A)$. The fibre above a point $d \in \mathbb{C}^*$ is just the set of 2×2 -matrices of determinant d , which is a Zariski-closed set and a single SL_2 -orbit. The fibre above $0 \in \mathbb{C}$ is the set of all matrices of rank ≤ 1 . This is, of course, also a Zariski closed set, but *not* a single orbit—indeed, it consists of the closed orbit consisting of the zero matrix and the non-closed orbit consisting of all matrices of rank 1. Note that the latter orbit has 0 in its closure.

These two examples illustrate the general situation: for *finite* G the fibres of the quotient map are precisely the orbits of G , while for *infinite* G they are certain G -stable closed sets.

Theorem 6.3.4. *Let Z denote the Zariski closure of $\pi(W)$, that is, the set of all points in $\mathbb{C}^m$ that satisfy all polynomial relations that are satisfied by the invariants $f_1, \dots, f_m$. The quotient map π has the following properties:*

1. $\pi(gw) = \pi(w)$ for all $g \in G$, $w \in W$;
2. the fibres of π are G -stable, Zariski-closed subsets of W ;
3. for any regular (polynomial) map $\psi : W \rightarrow \mathbb{C}^m$ that satisfies $\psi(gw) = \psi(w)$ for all $g \in G$ there exists a unique regular map $\phi : Z \rightarrow \mathbb{C}^m$ such that $\phi \circ \pi = \psi$.
4. π is surjective onto Z ;

Proof. 1. $\pi(gw) = (f_1(gw), \dots, f_k(gw)) = (f_1(w), \dots, f_k(w))$ because the f_i are invariant.

2. If $w \in \pi^{-1}(z)$, then $\pi(gw) = \pi(w) = z$, so $gw \in \pi^{-1}(z)$.

3. Let $y_1, \dots, y_m$ be the standard coordinates on $\mathbb{C}^m$. Then $y_i \circ \psi$ is a G -invariant polynomial on W for all i . As the f_j generate these polynomials, we may write $y_i \circ \psi$ as $g_i(f_1, \dots, f_k)$ for some k -variate polynomial g_i . Now the regular map $\phi : Z \rightarrow \mathbb{C}^m$, $z \mapsto (g_1(z), \dots, g_m(z))$ has the required property. Notice that the g_i need not be unique. However, the map $Z \rightarrow \mathbb{C}^m$ with the required property is unique: if ϕ_1, ϕ_2 both have the property, then necessarily $\phi_1(\pi(w)) = \phi_2(\pi(w)) = \psi(w)$ for all $w \in W$, so that ϕ_1 and ϕ_2 agree on the subset $\mathrm{im} \pi$ of Z . Since Z is the Zariski closure of this set, ϕ_1 and ϕ_2 need to agree everywhere. (In fact $Z = \mathrm{im} \pi$ as we will see shortly.)

4. Let $z \in Z$. This means that the coordinates of z satisfy all polynomial relations satisfied by $f_1, \dots, f_k$, hence there exists a homomorphism $\phi : \mathbb{C}[W]^G = \mathbb{C}[f_1, \dots, f_k] \rightarrow \mathbb{C}$ of $\mathbb{C}$ -algebras sending f_i to z_i . The kernel of this homomorphism is a maximal ideal M_z in $\mathbb{C}[W]^G$. We claim that there exists a maximal ideal M' in $\mathbb{C}[W]$ whose intersection with $\mathbb{C}[W]^G$ is M_z . Indeed, let I be the ideal in $\mathbb{C}[W]$ generated by M_z . We only need to show that $I \neq \mathbb{C}[W]$; then the axiom of choice implies the existence of a maximal ideal containing I . Suppose, on the contrary, that $I \ni 1$, and write

$$1 = \sum_{i=1}^l a_i h_i \text{ with all } a_i \in \mathbb{C}[W], h_i \in M_z.$$

Since $\mathbb{C}[W]$ is completely reducible as a G -module, there exists a Reynolds operator ρ . Applying ρ to both sides of the equality yields

$$1 = \sum_{i=1}^l \rho(a_i) h_i,$$

where the $\rho(a_i)$ are in $\mathbb{C}[W]^G$. But this means that 1 lies in M_z , a contradiction to the maximality of the latter ideal. This proves the claim that such an M' exists. The maximal ideal M' is the kernel of evaluation at some point $w \in W$ by the discussion of maximal ideals after the Nullstellensatz. Thus we have found a point $w \in W$ with the property that evaluating f_i at w gives z_i . Thus $\pi(w) = z$ and we are done. $\square$

Remark 6.3.5. By (3) Z is “independent” of the choice of generators of $\mathbb{C}[W]^G$.

Exercise 6.3.6. Let $G = \mathbb{C}^*$ act on $W = \mathbb{C}^4$ by $t(x_1, x_2, y_1, y_2) = (tx_1, tx_2, t^{-1}y_1, t^{-1}y_2)$.

1. Find generators $f_1, \dots, f_k$ of the invariant ring $\mathbb{C}[W]^G$.
2. Determine the image Z of the quotient map $\pi : W \rightarrow \mathbb{C}^k$, $w \mapsto (f_1(w), \dots, f_k(w))$.
3. For every $z \in Z$ determine the fibre $\pi^{-1}(z)$.