Geometric Invariant Theory

1 Introduction

Geometric Invariant Theory is a useful technique to construct quotients by group actions in Algebraic Geometry. It is widely applied in many branches of Mathematics, in particular in the theory of moduli spaces.

For example, suppose you want to classify smooth curves of fixed genus $g \geq 2$ up to isomorphism; then you can embed any one of them in projective space (with the same numerical invariants!), so you can regard them as (some; and understanding which is an important part of the theory) points of a Hilbert scheme (this is the scheme representing closed subschemes of $\mathbb{P}^N$ with a fixed Hilbert polynomial; the easiest examples are $\mathbb{P}^m$ identifying hypersurfaces with their coefficients and $\mathrm{Gr}(k, n)$ parametrizing linear k-dimensional subspaces of $\mathbb{C}^n$). But you did not care about the polarization to start with, therefore you must now take the quotient by the action of PGL_N changing coordinates of $\mathbb{P}^N$.

For simplicity we will work over $\mathbb{C}$.

2 Reductive groups

Definition 2.1. A Lie group G is called *reductive* if it is the complexification of a compact Lie group K.

In terms of Lie algebras, this means that $\mathfrak{g} = \mathfrak{k} \otimes_{\mathbb{R}} \mathbb{C}$. Examples are $S^1 \subseteq \mathbb{C}^*$, $SU(n) \subseteq SL_n(\mathbb{C})$ and $U(n) \subseteq GL_n(\mathbb{C})$.

Reductive groups have the useful property that their linear representations split into direct sum of irreducible sub-representations. The key point is the following (and induction).

Lemma 2.2. *Let G be a complex reductive group and $\rho\colon G \to GL(V)$ a linear representation. Let $W \subseteq V$ be a G-invariant subspace. Then there exists a G-invariant complement W', s.t. $V = W \oplus W'$ as G-modules.*

Proof. Pick any Hermitian metric h on V. Then average on K (i.e. consider $< v, w >= \int_K h(k.v, k.w)$, using the Haar measure). $<,>$ is a K-invariant measure on V. Take $W' = W^{\perp}$ with respect to $<,>$; this is a K-invariant subspace.

Exercise: using that G is the complexification of K, show that W' is in fact G-invariant. □

Aside remark: in the context of linear algebraic groups it is usual to define a group G reductive when its unipotent radical (maximal connected, normal, unipotent subgroup) is trivial. In characteristic 0 this turns out to be equivalent to being linearly reductive (i.e. the property discussed above); in positive characteristic this is no more true.

The first example of non-reductive group that comes to mind is probably $\mathbb{G}_a$. The standard embedding $(\mathbb{C}, +) \to GL_2(\mathbb{C})$, $a \mapsto \left(\begin{smallmatrix} 1 & a \\ 0 & 1 \end{smallmatrix}\right)$, gives a representation which is not completely reducible ($<(1,0)>$ is an invariant subspace with no complement).

3 Introduction and first examples

Let G be a complex reductive group acting on a complex projective variety $X \subseteq \mathbb{P}^n$. In the following we shall suppose not only that the action extends to $\mathbb{P}^n$, but in fact that we have a "linearization", i.e. the action extends to the vector space $\mathbb{C}^{n+1}$ over $\mathbb{P}^n$ (i.e. we assume to have a group homomorphism $G \to SL_{n+1}(\mathbb{C})$ lifting $G \to PGL_{n+1}(\mathbb{C})$); or, equivalently, we may look at $\mathcal{O}_{\mathbb{P}^n}(-1) \subseteq \mathbb{P}^n \times \mathbb{C}^{n+1}$ with the diagonal action. More in general, we may ask that the action extends to a line bundle L on X (the previous case is essentially $L = \mathcal{O}_X(1)$); a natural choice could be $\mathcal{K}_X = \Lambda^{\dim(X)}\Omega^1_X$.

An *aside*: choosing a different linearization may change the notion of (semi)stability, therefore producing a different quotient $X//G$; it is possible (and interesting) to study how these different quotients are related, and it turns out that there usually are birational maps between two of them (flips). The so-called VGIT (variation of GIT) gives us an insight into the birational geometry of these varieties and suggests a path to carry out the MMP on them.

As a first attempt we can consider the topological quotient X/G; this will not be well-behaved in general, e.g. not separated. The picture we bear in mind is the following: there exist non-closed orbits having lower-dimensional orbits in their closure; it may well happen that one such smaller orbit lies in the closure of more than one.

Example 3.1. Let the torus $\mathbb{C}^*$ act on $\mathbb{C}^2$ by $\lambda.(x, y) = (\lambda x, \lambda^{-1} y)$. Orbits are C_α: hyperbola $xy - \alpha$ for $\alpha \in \mathbb{C} \setminus \{0\}$, E_x: the punctured x-axis, E_y: the punctured y-axis and O: the origin. Notice that it is not enough to take out the lowest dimensional orbit in order to get a separated quotient topological space: in this example, both E_x and E_y are limits of C_α for $\alpha \to 0$.

Since affine algebraic varieties are completely determined by the regular functions on them, and since the projection map is supposed to pull back functions on the quotient to G-invariant functions on X, then it might make sense to consider $\operatorname{Spec} \mathbb{C}[X]^G$ as a plausible quotient. In this example it is easy to check that $\mathbb{C}[x, y]^{\mathbb{C}^*} = \mathbb{C}[xy] \simeq \mathbb{C}[\alpha]$, i.e. the putative quotient is the affine line.

It is good to remark that this operation can be made inside the category of affine algebraic varieties. Reducedness is obvious; what is not trivial here is the well-known result that if G is a reductive group acting on a $\mathbb{C}$-algebra of finite type, then the subring of invariants is again finitely generated. This is basically a consequence of the Hilbert basis theorem and of the existence of "averaging" operators, i.e. a collection of natural and G-invariant projections $V \to V^G$, for any linear representation of G (so-called Reynolds operators). If one wants to pass to the quasi-projective setup, a crucial ingredient is the existence of affine G-invariants open coverings (e.g. they do exist if G is finite).

Different strategies have been considered to face the problem of constructing a meaningful quotient in the example above (and further):

1. Kapranov's Chow quotient parametrizes invariant conics in this case, therefore the three bad orbits are identified. The quotient is $\mathbb{A}^1_{\mathbb{C}}$, with coordinate α.

2. GIT quotient: first remove undesired orbits (which are called *unstable*), then consider the space of remaining (*semistable*) orbits up to some equivalence relation (which is basically identifying orbits the closures of which intersect).

3. Symplectic quotient: remove unstable orbits and then pick the closed orbits among the remaining ones.

It turns out that GIT and symplectic reduction are closely related, since equivalence classes of semi-stable orbits contain one unique closed orbit (the lowest-dimensional one).

Example 3.2. Consider $\mathbb{C}^*$ acting on $\mathbb{C}^n$ via $\lambda \mapsto \lambda \,\mathrm{Id}$. In this case the invariant functions are only constant ones, therefore the "affine" quotient would be the point $\mathrm{Spec}(\mathbb{C})$. This is unsatisfactory for our geometric intuition, since the orbit space should resemble $\mathbb{P}^{n-1}$ (this represents most orbits, in fact all punctured lines in $\mathbb{C}^n$; the origin is the only point that is missed). $\mathbb{C}$ is indeed the space of regular functions on $\mathbb{P}^{n-1}$. The reason is that projective varieties have too few (global) regular functions to describe their geometric structure. Another remark: all closures of the orbits in this example meet in the origin. We will soon see how to embed $\mathbb{C}^n \subseteq \mathbb{P}^n$ and extend the $\mathbb{C}^*$-action in a way that the projective theory of GIT gives us the desired answer.

4 Quotients

We had a slogan that the goal of GIT is to construct quotients in algebraic geometry; one would like to consider the quotient topological space (which is the natural way to endow the set of orbits with a topology) and give it the structure of an algebraic variety. This is not always possible. On the other hand, at least in the affine setting, one is tempted to consider the spectrum of the algebra of invariants as a quotient, but this sometimes happen to record little information about the topological quotient. What we are going to do is to get rid of bad orbits and obtain a reasonable quotient space of the remaining ones. Let's now go through some definitions of reasonable quotients and the properties we would expect from them.

Definition 4.1. Let $\mathcal{C}$ be a category of algebro-geometric objects (we have in mind 1. (quasi-)affine algebraic varieties, 2. (quasi-)projective varieties, 3. schemes over some base, say $\mathbb{C}$). Let G be a group (object in $\mathcal{C}$) acting on $X \in \mathrm{ob}(\mathcal{C})$. A morphism $\pi\colon X \to Y$ in $\mathcal{C}$ is called a *categorical quotient* if, for every G-invariant arrow $f\colon X \to Z$ (i.e. equivariant with respect to the trivial action of G on Z) in $\mathcal{C}$, there exists a map $Y \to Z$ that factorizes f through π.

Uniqueness follows from the universal property.

Exercise: let $\varphi\colon G \times X \to X$ be the action map. Write down the diagrams that this map is required to satisfy (i.e. identity, compatibility).

Definition 4.2. A map $\pi\colon X \to Y$ is called a *good quotient* if

1. π is G-invariant (i.e. constant on orbits) and surjective;

2. π is affine and, for every affine open $U \subseteq Y$, $\pi^\sharp$ identifies $\mathbb{C}[U]$ as the subring of G-invariant functions inside $\mathbb{C}[\pi^{-1}(U)]$;

3. if W is closed and G-invariant in X, then $\pi(X)$ is closed (i.e. π *submersive*; it is sometimes required to be universally submersive);

4. if W_1 and W_2 are closed, G-invariants subsets of X such that they do not intersect, then also $\pi(W_1) \cap \pi(W_2) = \varnothing$.

Definition 4.3. A map $\pi\colon X \to Y$ is called a *geometric quotient* if it is a good quotient and Y is an orbit space.

Basically, a good quotient may fail to be a geometric quotient by grouping together a bunch of orbits. The basic results of GIT for reductive groups are the following:

- in the affine setting, one obtains a good quotient by taking $X \to \operatorname{Spec} \mathbb{C}[X]^G$;
- in the projective case, one can throw away some orbits and obtain a good quotient (that is also a projective variety) of the other ones (semistable orbits); the quotient map restricts to a geometric quotient on the set of stable orbits (the quotient is only quasi-projective).

Furthermore, it is possible to give a numerical characterization of (semi)-stable points.

5 GIT in the projective setting

Let us suppose that a complex reductive group G acts on a projective variety X, with a linearization $G \curvearrowright \mathcal{O}_X(1)$ (the pullback of the hyperplane bundle). Since the quotient is supposed to be described by G-invariant functions on X, in the projective case it makes sense to set $X//G = \operatorname{Proj}(\bigoplus_{l\geq 0} H^0(X, \mathcal{O}_X(l))^G)$, where H^0 is standard notation for global sections and $\mathcal{O}_X(l) = \mathcal{O}_X(1)^{\otimes l}$.

Define a rational map $X \dashrightarrow X//G$ by $x \mapsto \mathrm{ev}_x$. Recall that

$$\bigoplus_{l\geq 0} H^0(X, \mathcal{O}_X(l))^G$$

is finitely generated, in particular we may assume it is generated in some degree d (the Proj construction is unchanged under taking the subring $\bigoplus_{l\geq 0} R_{dl}$ of $\bigoplus_{l\geq 0} R_l$); evaluate invariant sections of $O_X(d)$ at x to obtain a map

$$\mathrm{ev}_x \in \operatorname{Hom}(H^0(X, \mathcal{O}_X(d))^G, \mathbb{C}_x).$$

This gives a rational map from X to $\mathbb{P}(H^0(X, \mathcal{O}_X(d))^G)^*$, which is not defined on points for which $\mathrm{ev}_x = 0$. This motivates the following

Definition 5.1. A point $x \in X$ is called *unstable* if every $s \in H^0(X, \mathcal{O}_X(d))^G$ vanishes on x. It is called *semi-stable* if it is not unstable. The locus of semi-stable points in X is denoted by X^{ss}.

Remark that this is a sensible definition because $\bigoplus_{k\geq 0} H^0(X, \mathcal{O}_X(k))^G$ is generated in degree d; if there is one section (of any degree) not vanishing at x, then there will also be a non-vanishing section of degree d. Furthermore, X^{ss} is open in X and we have a morphism $X^{ss} \to X//G$.

The next definition is motivated by the following: we want to pick those orbits which are separated from nearby ones by invariant functions; morally, we want to see around which points the space of orbits embeds into the GIT quotient.

Definition 5.2. A point $x \in X$ is *stable* if it is semi-stable, G acts on X_f with closed orbits (for some G-invariant function f s.t. $f(x) \neq 0$) and the stabilizer of x in G is finite.

The second part of this definition should mean that:

1. if we restrict to an affine open inside X_f (and trivialize $\mathcal{O}_X(1)$), for any two orbits O_1 and O_2 we can find G-invariant functions $s_1,\ s_2$ s.t. $s_{i|O_i} = 1$ and $s_{i|O_{2-i}} = 0$ for $i = 1, 2$;

2. this condition holds infinitesimally as well, i.e. for any vector $0 \neq v \in T_xX/T_x(G.x)$ we can find a G-invariant function s that vanishes on $G.x$ but s.t. $D_v(s) \neq 0$.

Remark that also the locus of stable points is open. Unfortunately it might happen to be empty.

Theorem 5.3. *Let G be a reductive group acting on a projective variety X with a linearization of the action.*

1. *The map $\pi\colon X^{ss} \to Y = X//G$ is a good quotient.*

2. *There exists an open subset $Y^s \subseteq Y$ s.t. $\pi^{-1}(Y^s) = X^s$ and the restriction of π to X^s is a geometric quotient.*

3. *For any $x_1,\ x_2 \in X^{ss}$, we have $\pi(x_1) = \pi(x_2)$ if and only if $\overline{G.x_1} \cap \overline{G.x_2} \cap X^{ss} \neq \varnothing$. A point $x \in X^{ss}$ is stable if and only if $G.x$ is closed in X^{ss} and x has finite stabilizer.*

The following proposition asserts that we can distinguish (semi-)stable points by looking at the orbits of their liftings in $\mathbb{C}^{n+1}$.

Proposition 5.4. *Let G act on $X \subseteq \mathbb{P}^n$ as above. Let $x \in X$ and $\widetilde{x}$ be any lifting of x to $\mathbb{C}^{n+1}$. Then*

- *x is semi-stable if and only if 0 lies in the closure of the G-orbit of $\widetilde{x}$.*

- *x is stable if and only if it is semi-stable, $G.\widetilde{x}$ is closed in $\mathbb{C}^{n+1}$ and $\widetilde{x}$ has finite stabilizer.*

Observe that the action being linearized induces an action on $\mathcal{O}_X(-1) \subseteq \mathbb{P}^n \times \mathbb{C}^{n+1}$. A *one-parameter subgroup* (1PS) in G is a homomorphism of algebraic groups $\mathbb{G}_m = \mathbb{C}^* \to G$, i.e. a cocharacter. Any 1PS in G induces an action of $\mathbb{C}^*$ on X and it is clear from the above that (semi-)stable points for G are (semi-)stable for $\mathbb{C}^*$ as well. A non-trivial and useful result is that the behaviour under all the 1PS's determines stability under the action of the whole group.

Theorem 5.5 (Hilbert-Mumford criterion). *Let G and X be as above. A point $x \in X$ is (semi-)stable under the action of G if and only if it is (semi-)stable under the action of $\mathbb{C}^*$ induced by every 1PS of G.*

The action of $\mathbb{C}^*$ on $\mathbb{C}^{n+1}$ can be diagonalized in some basis $\{e_0, \ldots, e_n\}$; suppose that $\mathbb{C}^*$ acts with weight p_i on each e_i. We may take a lifting of x and express it in this basis, $\widetilde{x} = \sum x_i e_i$; then $\lambda.\widetilde{x} = \sum \lambda^{p_i} x_i e_i$. Define $\mu^+(x) = \max\{-p_i : x_i \neq 0\}$ and $\mu^-(x) = \max\{p_i : x_i \neq 0\}$. Then $\mu^+(x) > 0 \Leftrightarrow \lim_{\lambda \to 0} \lambda.x$ does not exist and $\mu^+(x) = 0 \Leftrightarrow \lim_{\lambda \to 0} \lambda.x$ exists and is not 0. Analogous statements hold with μ^- and $\lim_{\lambda \to \infty}$. Therefore

Lemma 5.6. *A point $x \in X$ is semi-stable for the action of $\mathbb{C}^*$ if and only if $\mu^+(x) \geq 0$ and $\mu^-(x) \geq 0$. It is stable iff strict inequalities hold.*

It may be useful to notice that, for a 1PS $\varphi \colon \mathbb{C}^* \to G$, $x \in X$ and $g \in G$, one has $\mu^\pm(g.x, \varphi) = \mu^\pm(x, g^{-1}\varphi g)$. Therefore one can change the 1PS by conjugation in G whenever this makes the computations easier; in particular any 1PS can be conjugated into a chosen maximal subtorus of G (e.g. diagonal matrices will work for the special linear group).

From the above it follows that the Hilbert-Mumford criterion may be expressed in the following numerical form.

Theorem 5.7. *Let $x \in X$. For any 1PS $\varphi \colon \mathbb{C}^* \to G$, set $x_0 = \lim_{\lambda \to 0} \varphi(\lambda).x$. Then x_0 is a fixed point for the 1PS, therefore $\mathbb{C}^*$ acts on $\mathcal{O}(-1)_{x_0}$ with a certain weight, say p. Then x is stable iff $p < 0$ for any 1PS and semi-stable iff $p \leq 0$.*

Example 5.8. Let us go back to the example of $\mathbb{C}^*$ acting on $\mathbb{C}^n$ via $\lambda \mapsto \lambda Id$. One possibility is to consider $\mathbb{C}^*$ as acting on the trivial line bundle over $\mathbb{C}^n$ but with a non-trivial character, say with weight $-p$. Then it is an **exercise** to check that $H^0(\mathbb{C}^n, L^k)^{\mathbb{C}^*}$ are homogeneous polynomials of degree pk; one concludes that $\mathbb{C}^n//\mathbb{C}^* = \mathbb{P}^{n-1}$ with the line bundle $\mathcal{O}(p)$ - this is obvious for $p = 1$, for $p > 1$ prove that homogeneous primes of a graded ring $R = \bigoplus_{k \geq 0} R_k$ correspond one-to-one to homogeneous primes of $S = \bigoplus_{k \geq 0} R_{kd}$ for any fixed integer d.

This can also be seen as follows: we can embed $\mathbb{C}^n$ as an affine chart of $\mathbb{P}^n$, mapping $\underline{z}$ to $[\underline{z} : 1]$. The action of $\mathbb{C}^*$ can be extended to $\lambda.[\underline{z} : w] = [\lambda\underline{z} : \lambda^{-n}w]$. Show that unstable points are $[0 : \ldots : 0 : 1]$ and all those of the form $[\underline{z} : 0]$ for $\underline{z} \in \mathbb{C}^n \setminus \{0\}$; all other points are stable and the quotient is $\mathbb{P}^{n-1}$.

Example 5.9 (Configurations of n points in $\mathbb{P}^1$). A length n subscheme of $\mathbb{P}^1$ is (semi-)stable for the natural action of SL_2 iff each point has multiplicity $< n/2$ ($\leq$).

Consider a torus $\mathbb{C}^* \to SL_2$ and fix a basis in which it is diagonal $\lambda \mapsto \left(\begin{smallmatrix} \lambda^p & 0 \\ 0 & \lambda^{-p} \end{smallmatrix}\right)$. Length n subschemes are represented by degree n homogeneous polynomials in two variables, i.e. $\mathbb{P}H^0(\mathbb{P}^1, \mathcal{O}(n))$. Write $f(x, y) = a_0 x^n + \ldots + a_n y^n$ in the chosen basis. Observe that $\lambda.(x^i y^{n-1}) = \lambda^{p(2i-n)} x^i y^{n-i}$, so $\lambda.f \to \infty$ for $\lambda \to 0$ unless $a_0 = \ldots = a_{\lceil n/2 \rceil} = 0$, i.e. unless $[1 : 0]$ has multiplicity $> n/2$.

Alternatively, we could have said that f tends to $a_i x^{n-i} y^i$ for the smallest i s.t. $a_i \neq 0$; and $\mathbb{C}^*$ acts on the fiber over $x^{n-i} y^i$ with weight $2i - n$.

Exercise Consider the case of $S^k\mathbb{P}^n$, configurations of k points in $\mathbb{P}^n$. We can see them as union of k hyperplanes in the dual projective space.

Exercise Consider the action of SL_r on $\mathrm{Hom}(\mathbb{C}^r, \mathbb{C}^n)$ for $r < n$. Show that semi-stable points correspond to injective homomorphisms. There are no strictly semi-stable points. The GIT quotient is $\mathrm{Gr}(r, n)$.