MODULI PROBLEMS AND GEOMETRIC INVARIANT THEORY

VICTORIA HOSKINS

1. Introduction

In this course, we study moduli problems in algebraic geometry and constructions of moduli spaces using geometric invariant theory. A moduli problem is essentially a classification problem: we want to classify certain geometric objects up to some notion of equivalence (key examples are vector bundles on a fixed variety up to isomorphism or hypersurfaces in $\mathbb{P}^n$ up to projective transformations). We are also interested in understanding how these objects deform in families and this information is encoded in a moduli functor. An ideal solution to a moduli problem is a (fine) moduli space, which is a scheme that represents this functor. However, there are many simple moduli problems which do not admit such a solution. Often we must restrict our attention to well-behaved objects to construct a moduli space. Typically the construction of moduli spaces is given by taking a group quotient of a parameter space, where the orbits correspond to the equivalence classes of objects.

Geometric invariant theory (GIT) is a method for constructing group quotients in algebraic geometry and it is frequently used to construct moduli spaces. The core of this course is the construction of GIT quotients. Eventually we return to our original motivation of moduli problems and construct moduli spaces using GIT. We complete the course by constructing moduli spaces of projective hypersurfaces and moduli spaces of (semistable) vector bundles over a smooth complex projective curve.

Let us recall the quotient construction in topology: given a group G acting on a topological space X, we can give the orbit space $X/G := \{G \cdot x : x \in X\}$ the quotient topology, so that the quotient map $\pi : X \to X/G$ is continuous. In particular, π gives a quotient in the category of topological spaces. More generally, we can suppose G is a Lie group and X has the structure of a smooth manifold. In this case, the quotient X/G will not always have the structure of a smooth manifold (for example, the presence of non-closed orbits, usually gives a non-Hausdorff quotient). However, if G acts properly and freely, then X/G has a smooth manifold structure, such that π is a smooth submersion.

In this course, we are interested in actions of an affine algebraic group G (that is, an affine scheme with a group structure such that multiplication and inversion are algebraic morphisms). More precisely, we're interested in algebraic G-actions on an algebraic variety (or scheme of finite type) X over an algebraically closed field k. As most affine groups are non-compact, their actions typically have some non-closed orbits. Consequently, the topological quotient X/G will not be Hausdorff. However one could also ask whether we should relax the idea of having an orbit space, in order to get a quotient with better geometrical properties. More precisely, we ask for a categorical quotient in the category of finite type k-schemes; that is, a G-invariant morphism $\pi : X \to Y$ which is universal (i.e., every other G-invariant morphism $X \to Z$ factors uniquely through π). With this definition, it is not necessary for Y to be an orbit space and so we can allow π to identify some orbits in order to get an algebraic quotient.

Geometric invariant theory, as developed by Mumford in [25], shows that for a reductive group G acting on a quasi-projective scheme X (with respect to an ample linearisation) one can construct an open subvariety $U \subset X$ and a categorical quotient $U/\!/G$ of the G-action on U which is a quasi-projective scheme. In general, the quotient will not be an orbit space but it contains an open subscheme V/G which is the orbit space for an open subset $V \subset U$. If X is an affine scheme, we have that $U = X$ and the categorical quotient is also an affine scheme and if X is a projective scheme, the categorical quotient is also projective. We briefly summarise the main techniques involved in GIT.

Let $X = \operatorname{Spec} A$ be an affine scheme of finite type over an algebraically closed field k; then $A = \mathcal{O}(X) := \mathcal{O}_X(X)$ is a finitely generated k-algebra. An algebraic G-action on X induces G-action on the ring $\mathcal{O}(X)$ of regular functions on X. For any G-invariant morphism $f : X \to Z$ of schemes, the image of the associated homomorphism $f^* : \mathcal{O}(Z) \to \mathcal{O}(X)$ is contained in the subalgebra $\mathcal{O}(X)^G$ of G-invariant functions. In particular, if $\mathcal{O}(X)^G$ is finitely generated as a k-algebra, then the associated affine scheme $\operatorname{Spec} \mathcal{O}(X)^G$ is also of finite type over k and the inclusion $\mathcal{O}(X)^G \hookrightarrow \mathcal{O}(X)$ induces a morphism $X \to X//G := \operatorname{Spec} \mathcal{O}(X)^G$, which is categorical quotient of the G-action on X. The affine GIT quotient $X \to X//G$ identifies any orbits whose closures meet, but restricts to an orbit space on an open subscheme of so-called stable points.

An important problem in GIT is determining when the ring of invariants $\mathcal{O}(X)^G$ is finitely generated; this is known as Hilbert's 14th problem. For $G = \mathrm{GL}_n$ over the complex numbers, Hilbert showed that the invariant ring is always finitely generated. However, for a group G built using copies of the additive group $\mathbb{G}_a$, Nagata gave a counterexample where $\mathcal{O}(X)^G$ is non-finitely generated. Furthermore, Nagata proved for any reductive group G, the ring of invariants $\mathcal{O}(X)^G$ is finitely generated. Consequently, (classical) GIT is concerned with the action of reductive groups; for developments on the theory for non-reductive groups, see [6].

The affine GIT quotient serves as a guide for the general approach: as every scheme is constructed by gluing affine schemes, the general theory is obtained by gluing affine GIT quotients. Ideally, we would to cover X by G-invariant open affine sets and glue the corresponding affine GIT quotients. The open G-affine sets are given by non-vanishing loci of invariant sections of a line bundle L on X, to which we have lifted the G-action. However, usually we cannot cover the whole of X with such open subsets, but rather only an open subset X^{ss} of X of so-called semistable points. In this case, we have a categorical quotient of X^{ss} which restricts to an orbit space on the stable locus X^s.

The definitions of (semi)stability are given in terms of the existence of invariant sections of a line bundle with certain properties. However, as calculating rings of invariants is difficult, one often instead makes use of a numerical criterion for semistability known as the Hilbert–Mumford criterion. More precisely, the Hilbert–Mumford criterion reduces the semistability of points in a projective scheme to the study of the weights of all 1-dimensional subtori $\mathbb{G}_m \subset G$.

The techniques of GIT have been used to construct many moduli spaces in algebraic geometry and finally we return to the construction of some important moduli spaces. The main examples we cover in this course are the GIT construction of moduli spaces of hypersurfaces and the moduli spaces of (semistable) vector bundles on a smooth complex projective curve.

The main references for this course are the books of Newstead [31] and Mukai [24] on moduli problems and GIT, and the book of Mumford [25] on GIT.

Notation and conventions. Throughout we fix an algebraically closed field k; at certain points in the text we will assume that the characteristic of the field is zero in order to simplify the proofs. By a scheme, we always mean a finite type scheme over k. By a variety, we mean a reduced separated (finite type) scheme over k; in particular, we do not assume varieties are irreducible. We let $\mathcal{O}(X) := \mathcal{O}_X(X)$ denote the ring of regular functions on a scheme X. For a projective scheme X with ample line bundle L, we let $R(X, L)$ denote the homogeneous coordinate ring of X given by taking the direct sum of the spaces of sections of all non-negative powers of L.

Acknowledgements. First and foremost, I would like to thank Frances Kirwan for introducing me to geometric invariant theory. These lectures notes were written for a masters course at the Freie Universität Berlin in the winter semester of 2015 and I am grateful to have had some excellent students who actively participated in the course and engaged with the topic: thank you to Anna-Lena, Christoph, Claudius, Dominic, Emmi, Fei, Felix, Jennifer, Koki, Maciek, Maik, Maksymilian, Markus, and Vincent. Thanks also to Eva Martínez for all her hard work running the exercise sessions and to the students for persevering with the more challenging exercises. I am very grateful to Giangiacomo Sanna for many useful comments and suggestions on the script. Finally, I would especially like to thank Simon Pepin Lehalleur for numerous helpful discussions and contributions, and for his encyclopaedic knowledge of algebraic groups.

2. Moduli problems

2.1. **Functors of points.** In this section, we will make use of some of the language of *category theory*. We recall that a morphism of categories $\mathcal{C}$ and $\mathcal{D}$ is given by a *(covariant) functor* $F : \mathcal{C} \to \mathcal{D}$, which associates to every object $C \in \mathcal{C}$ an object $F(C) \in \mathcal{D}$ and to each morphism $f : C \to C'$ in $\mathcal{C}$ a morphism $F(f) : F(C) \to F(C')$ in D such that F preserves identity morphisms and composition. A *contravariant functor* $F : \mathcal{C} \to \mathcal{D}$ reverses arrows: so F sends $f : C \to C'$ to $F(f) : F(C') \to F(C)$.

The notion of a morphism of (covariant) functors $F, G : \mathcal{C} \to \mathcal{D}$ is given by a *natural transformation* $\eta : F \to G$ which associates to every object $C \in \mathcal{C}$ a morphism $\eta_C : F(C) \to G(C)$ in D which is compatible with morphisms $f : C \to C'$ in $\mathcal{C}$, i.e. we have a commutative square

$$\begin{tikzcd} F(C) \arrow[r, "\eta_C"] \arrow[d, "F(f)"'] & G(C) \arrow[d, "G(f)"] \\ F(C') \arrow[r, "\eta_{C'}"'] & G(C'). \end{tikzcd}$$

We note that if F and G were contravariant functors, the vertical arrows in this square would be reversed. If η_C is an isomorphism in $\mathcal{D}$ for all $C \in \mathcal{C}$, then we call η a *natural isomorphism* or simply an isomorphism of functors.

Remark 2.1. The focus of this course is moduli problems, rather than category theory and so we are doing naive category theory (in the sense that we allow the objects of a category to be a class). This is analogous to doing naive set theory without a consistent axiomatic approach. However, for those interested in category theory, this can all be handled in a consistent manner, where one pays more careful attention to the size of the set of objects. One approach to this more formal category theory can be found in the book of Kashiwara and Schapira [18]. Strictly speaking, in this case, one should work with the category of 'small' sets.

Let Set denote the category of sets and let Sch denote the category of schemes (of finite type over k).

Definition 2.2. The functor of points of a scheme X is a contravariant functor $h_X := \mathrm{Hom}(-, X) : \mathrm{Sch} \to \mathrm{Set}$ from the category of schemes to the category of sets defined by

$$\begin{array}{rcllcl} h_X(Y) & := & \mathrm{Hom}(Y, X) & & & \\ h_X(f : Y \to Z) & := & h_X(f) : & h_X(Z) & \to & h_X(Y) \\ & & & g & \mapsto & g \circ f. \end{array}$$

Furthermore, a morphism of schemes $f : X \to Y$ induces a natural transformation of functors $h_f : h_X \to h_Y$ given by

$$\begin{array}{rccc} h_{f,Z} : & h_X(Z) & \to & h_Y(Z) \\ & g & \mapsto & f \circ g. \end{array}$$

Contravariant functors from schemes to sets are called *presheaves on* Sch and form a category, with morphisms given by natural transformations; this category is denoted $\mathrm{Psh(Sch)} := \mathrm{Fun}(\mathrm{Sch}^{\mathrm{op}}, \mathrm{Set})$, the category of presheaves on Sch. The above constructions can be phrased as follows: there is a functor $h : \mathrm{Sch} \to \mathrm{Psh(Sch)}$ given by

$$X \mapsto h_X \qquad (f : X \to Y) \mapsto h_f : h_X \to h_Y.$$

In fact, there is nothing special about the category of schemes here. So for any category $\mathcal{C}$, there is a functor $h : \mathcal{C} \to \mathrm{Psh}(\mathcal{C})$.

Example 2.3. For a scheme X, we have that $h_X(\operatorname{Spec} k) := \mathrm{Hom}(\operatorname{Spec} k, X)$ is the set of k-points of X and, for another scheme Y, we have that $h_X(Y)$ is the set of Y-valued points of X. Let $X = \mathbb{A}^1$ be the affine line; then the functor of points $h_{\mathbb{A}^1}$ associates to a scheme Y the set of functions on Y (i.e. morphisms $Y \to \mathbb{A}^1$). Similarly, for the scheme $\mathbb{G}_m = \mathbb{A}^1 - \{0\}$, the functor $h_{\mathbb{A}^1}$ associates to a scheme Y the set of invertible functions on Y.

Lemma 2.4 (The Yoneda Lemma). *Let $\mathcal{C}$ be any category. Then for any $C \in \mathcal{C}$ and any presheaf $F \in \mathrm{Psh}(\mathcal{C})$, there is a bijection*

$$\{\text{natural transformsations } \eta : h_C \to F\} \longleftrightarrow F(C).$$

given by $\eta \mapsto \eta_C(Id_C)$.

Proof. Let us first check that this is surjective: for an object $s \in F(C)$, we define a natural transformation $\eta = \eta(s) : h_C \to F$ as follows. For $C' \in \mathcal{C}$, let $\eta_{C'} : h_C(C') \to F(C')$ be the morphism of sets which sends $f : C' \to C$ to $F(f)(s)$ (recall that $F(f) : F(C) \to F(C')$). This is compatible with morphisms and, by construction, $\eta_C(\mathrm{id}_C) = F(\mathrm{id}_C)(s) = s$.

For injectivity, suppose we have natural transformations $\eta, \eta' : h_C \to F$ such that $\eta_C(\mathrm{Id}_C) = \eta'_C(\mathrm{Id}_C)$. Then we claim $\eta = \eta'$; that is, for any C' in $\mathcal{C}$, we have $\eta_{C'} = \eta'_{C'} : h_C(C') \to F(C')$. Let $g : C' \to C$, then as η is a natural transformation, we have a commutative square

$$\begin{tikzcd} h_C(C) \arrow[r, "\eta_C"] \arrow[d, "h_C(g)"'] & F(C) \arrow[d, "F(g)"] \\ h_C(C') \arrow[r, "\eta_{C'}"'] & F(C'). \end{tikzcd}$$

It follows that

$$(F(g) \circ \eta_C)(\mathrm{id}_C) = (\eta_{C'} \circ h_C(g))(\mathrm{Id}_C) = \eta_{C'}(g)$$

and similarly, as η' is a natural transformation, that $(F(g) \circ \eta'_C)(\mathrm{id}_C) = \eta'_{C'}(g)$. Hence

$$\eta_{C'}(g) = F(g)(\eta_C(\mathrm{id}_C)) = F(g)(\eta'_C(\mathrm{id}_C)) = \eta'_{C'}(g)$$

as required. □

The functor $h : \mathcal{C} \to \mathrm{Psh}(\mathcal{C})$ is called the Yoneda embedding, due to the following corollary.

Corollary 2.5. *The functor $h : \mathcal{C} \to \mathrm{Psh}(\mathcal{C})$ is fully faithful.*

Proof. We recall that a functor is fully faithful if for every C, C' in $\mathcal{C}$, the morphism

$$\mathrm{Hom}_{\mathcal{C}}(C, C') \to \mathrm{Hom}_{\mathrm{Psh}(\mathcal{C})}(h_C, h'_C)$$

is bijective. This follows immediately from the Yoneda Lemma if we take $F = h'_C$. □

Exercise 2.6. Show that if there is a natural isomorphism $h_C \to h'_C$, then there is a canonical isomorphism $C \to C'$.

The presheaves in the image of the Yoneda embedding are known are representable functors.

Definition 2.7. A presheaf $F \in \mathrm{Psh}(\mathcal{C})$ is called representable if there exists an object $C \in \mathcal{C}$ and a natural isomorphism $F \cong h_C$.

Question: Is every presheaf $F \in \mathrm{Psh}(\mathrm{Sch})$ representable by a scheme X?

The question has a negative answer, as we will soon see below. However, we are most interested in answering this question for special functors F, known as moduli functors, which classify certain geometric families. Before we introduce these moduli functors, we start with the naive notion of a moduli problem.

2.2. **Moduli problem.** A moduli problem is essentially a classification problem: we have a collection of objects and we want to classify these objects up to equivalence. In fact, we want more than this, we want a moduli space which encodes how these objects vary continuously in families; this information is encoded in a moduli functor.

Definition 2.8. A (naive) moduli problem (in algebraic geometry) is a collection $\mathcal{A}$ of objects (in algebraic geometry) and an equivalence relation $\sim$ on $\mathcal{A}$.

Example 2.9.

(1) Let $\mathcal{A}$ be the set of k-dimensional linear subspaces of an n-dimensional vector space and $\sim$ be equality.

(2) Let $\mathcal{A}$ be the set of n ordered distinct points on $\mathbb{P}^1$ and $\sim$ be the equivalence relation given by the natural action of the automorphism group PGL_2 of $\mathbb{P}^1$.
(3) Let $\mathcal{A}$ to be the set of hypersurfaces of degree d in $\mathbb{P}^n$ and $\sim$ can be chosen to be either equality or the relation given by projective change of coordinates (i.e. corresponding to the natural PGL_{n+1}-action).
(4) Let $\mathcal{A}$ be the collection of vector bundles on a fixed scheme X and $\sim$ be the relation given by isomorphisms of vector bundles.

Our aim is to find a scheme M whose k-points are in bijection with the set of equivalence classes $\mathcal{A}/\sim$. Furthermore, we want M to also encode how these objects vary continuously in 'families'. More precisely, we refer to $(\mathcal{A}, \sim)$ as a naive moduli problem, because there is often a natural notion of families of objects over a scheme S and an extension of $\sim$ to families over S, such that we can pullback families by morphisms $T \to S$.

Definition 2.10. Let $(\mathcal{A}, \sim)$ be a naive moduli problem. Then an extended moduli problem is given by

(1) sets $\mathcal{A}_S$ of families over S and an equivalence relation $\sim_S$ on $\mathcal{A}_S$, for all schemes S,
(2) pullback maps $f^* : \mathcal{A}_S \to \mathcal{A}_T$, for every morphism of schemes $T \to S$,

satisfying the following properties:

(i) $(\mathcal{A}_{\operatorname{Spec} k}, \sim_{\operatorname{Spec} k}) = (\mathcal{A}, \sim)$;
(ii) for the identity $\mathrm{Id} : S \to S$ and any family $\mathcal{F}$ over S, we have $\mathrm{Id}^*\mathcal{F} = \mathcal{F}$;
(iii) for a morphism $f : T \to S$ and equivalent families $\mathcal{F} \sim_S \mathcal{G}$ over S, we have $f^*\mathcal{F} \sim_T f^*\mathcal{G}$;
(iv) for morphisms $f : T \to S$ and $g : S \to R$, and a family $\mathcal{F}$ over R, we have an equivalence $(g \circ f)^*\mathcal{F} \sim_T f^*g^*\mathcal{F}$.

For a family $\mathcal{F}$ over S and a point $s : \operatorname{Spec} k \to S$, we write $\mathcal{F}_s := s^*\mathcal{F}$ to denote the corresponding family over $\operatorname{Spec} k$.

Lemma 2.11. *A moduli problem defines a functor $\mathcal{M} \in \mathrm{Psh}(\mathrm{Sch})$ given by*

$$\mathcal{M}(S) := \{\text{families over } S\}/\sim_S \qquad \mathcal{M}(f : T \to S) = f^* : \mathcal{M}(S) \to \mathcal{M}(T).$$

We will often refer to a moduli problem simply by its moduli functor. There can be several different extensions of a naive moduli problem.

Example 2.12. Let us consider the naive moduli problem given by vector bundles (i.e. locally free sheaves) on a fixed scheme X up to isomorphism. Then this can be extended in two different ways. The natural notion for a family over S is a locally free sheaf $\mathcal{F}$ over $X \times S$ flat over S, but there are two possible equivalence relations:

$$\begin{aligned} \mathcal{F} \sim'_S \mathcal{G} &\iff \mathcal{F} \cong \mathcal{G} \\ \mathcal{F} \sim_S \mathcal{G} &\iff \mathcal{F} \cong \mathcal{G} \otimes \pi_S^*\mathcal{L} \text{ for a line bundle } \mathcal{L} \to S \end{aligned}$$

where $\pi_S : X \times S \to S$. For the second equivalence relation, since $\mathcal{L} \to S$ is locally trivial, there is a cover S_i of S such that $\mathcal{F}|_{X \times S_i} \cong \mathcal{G}|_{X \times S_i}$. It turns out that the second notion of equivalence offers the extra flexibility we will need in order to construct moduli spaces.

Example 2.13. Let $\mathcal{A}$ consist of 4 ordered distinct points (p_1, p_2, p_3, p_4) on $\mathbb{P}^1$. We want to classify these quartuples up to the automorphisms of $\mathbb{P}^1$. We recall that the automorphism group of $\mathbb{P}^1$ is the projective linear group PGL_2, which acts as Möbius transformations. We define our equivalence relation by $(p_1, p_2, p_3, p_4) \sim (q_1, q_2, q_3, q_4)$ if there exists an automorphisms $f : \mathbb{P}^1 \to \mathbb{P}^1$ such that $f(p_i) = q_i$ for $i = 1, \ldots, 4$. We recall that for any 3 distinct points (p_1, p_2, p_3) on $\mathbb{P}^1$, there exists a unique Möbius transformation $f \in \mathrm{PGL}_2$ which sends (p_1, p_2, p_3) to $(0, 1, \infty)$ and the cross-ratio of 4 distinct points (p_1, p_2, p_3, p_4) on $\mathbb{P}^1$ is given by $f(p_4) \in \mathbb{P}^1 - \{0, 1, \infty\}$, where f is the unique Möbius transformation that sends (p_1, p_2, p_3) to $(0, 1, \infty)$. Therefore, we see that the set $\mathcal{A}/\sim$ is in bijection with the set of k-points in the quasi-projective variety $\mathbb{P}^1 - \{0, 1, \infty\}$.

In fact, we can naturally speak about families of 4 distinct points on $\mathbb{P}^1$ over a scheme S: this is given by a proper flat morphism $\pi : \mathcal{X} \to S$ such that the fibres $\pi^{-1}(s) \cong \mathbb{P}^1$ are smooth rational curves and 4 disjoint sections $(\sigma_1, \dots, \sigma_4)$ of π. We say two families $(\pi : \mathcal{X} \to S, \sigma_1, \dots, \sigma_4)$ and $(\pi' : \mathcal{X}' \to S, \sigma'_1, \dots, \sigma'_4)$ are equivalent over S if there is an isomorphism $f : \mathcal{X} \to \mathcal{X}'$ over S (i.e. $\pi = \pi' \circ f$) such that $f \circ \sigma_i = \sigma'_i$.

There is a tautological family over the scheme $S = \mathbb{P}^1 - \{0, 1, \infty\}$: let $\pi : \mathbb{P}^1 - \{0, 1, \infty\} \times \mathbb{P}^1 \to \mathbb{P}^1$ be the projection map and choose sections $(\sigma_1(s) = 0, \sigma_2(s) = 1, \sigma_3(s) = \infty, \sigma_4(s) = s)$. It turns out that this family over $\mathbb{P}^1 - \{0, 1, \infty\}$ encodes all families parametrised by schemes S (in the language to come, $\mathcal{U}$ is a *universal family* and $\mathbb{P}^1 - \{0, 1, \infty\}$ is a *fine moduli space).*

Exercise 2.14. Define an analogous notion for families of n ordered distinct points on $\mathbb{P}^1$ and let the corresponding moduli functor be denoted $\mathcal{M}_{0,n}$ (this is the moduli functor of n ordered distinct points on the curve $\mathbb{P}^1$ of genus 0). For $n = 3$, show that $\mathcal{M}_{0,3}(\operatorname{Spec} k)$ is a single element set and so is in bijection with the set of k-points of $\operatorname{Spec} k$. Furthermore, show there is a tautological family over $\operatorname{Spec} k$.