London Mathematical Society
Lecture Note Series 345

Algebraic and Analytic Geometry

Amnon Neeman

London Mathematical Society

CAMBRIDGE
UNIVERSITY PRESS

Algebraic and Analytic Geometry

Amnon Neeman
Australian National University

CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE UNIVERSITY PRESS
Cambridge, New York, Melbourne, Madrid, Cape Town, Singapore, São Paulo

Cambridge University Press
The Edinburgh Building, Cambridge CB2 8RU, UK

Published in the United States of America by Cambridge University Press, New York

www.cambridge.org
Information on this title: www.cambridge.org/9780521709835

First published 2007

Printed in the United Kingdom at the University Press, Cambridge

A catalogue record for this book is available from the British Library

ISBN 978-0-521-70983-5 paperback

Preface

This book came out of a course I taught, twice, at the Australian National University. I taught it first in the Fall of 2004, and then again, because of interest from some students and colleagues, in the Fall of 2005. The course was a one-semester affair, and the students were fourth-year undergraduates.

Given that these were undergraduate students in their final year, this could be one of the last few mathematics courses they would ever see. They might well decide to pursue interests having nothing to do with mathematics; they could, for all I know, choose to become doctors, or lawyers, or bankers, or politicians. My task was to present to them an overview of algebraic geometry. It would be premature to give them a thorough grounding in the field; a broad, panoramic picture seemed far more appropriate, and if possible the panorama should include glimpses into a wide assortment of pretty vistas, into more specialized areas, each of which is beautiful in its own right. I tried to cover interesting topics, without delving into too much detail on any one of them.

The first order of business was to choose the subject matter for the course. I had the option of teaching classical algebraic geometry; there are several excellent textbooks to choose from, written specifically for students at this level. But I wanted to teach modern algebraic geometry, and there really are no undergraduate treatments of the field. The consensus seems to be that this topic is beyond undergraduates, suitable only for courses at the graduate level.

So here I was, soon to face a class of math majors in their final year, and I had decided to teach them some modern algebraic geometry, even though there was no available textbook. I had to assemble the material myself. In so doing, I had to take into account the mathematics the students are likely to have seen in their first three years at university. Usually this would include some background on point-set topology, maybe a course on analytic functions in one complex variable, possibly a course on functional analysis, which would probably cover the Hahn–Banach Theorem and the Open Mapping Theorem, possibly a little about manifolds, maybe a rudimentary course on algebraic topology, and perhaps some basic algebra—groups, rings, fields, modules, if I were lucky maybe even the Hilbert Basis Theorem and the Nullstellensatz.

Algebraic geometry is a meeting place in which all the previous strands of knowledge magically converge; why not present it this way? This was my guiding philosophy in planning the course, and later the book. Bearing in mind that only an unusual math major will have seen every one of the topics listed, I tried not to lean too heavily on any one of them. But strong math majors should have met a large portion of these subjects, and my hope was that I was building on familiar ground.

Whenever we teach a course, especially a wide-ranging survey course, there will be the keen students, the enthusiastic ones who want to go a little beyond the discussions presented in class, who might even wish to pursue the subject further some day. This book was written for them. It covers the same topics treated in the course, and it tries to remain at the same level, demanding from the reader no prerequisites beyond what was assumed in my course. But the book has far more detail than the course, with many of the proofs included. I mention this because, if you decide to use the book to teach a course, and the course you have in mind is to be modeled approximately on the one I gave, then it would be a mistake to follow the book slavishly. In your lectures you will be presenting only selected portions of the book; you will have to pick and choose, deciding which parts of the material to present, and what to leave out.

I do not presume to make these decisions for you; the author has no authority to direct the reader how to use a book. As a rough guide let me, nevertheless, tell you what I did. Perhaps even more helpful: let me tell you what I would do if I come to teach the course again. The Introduction is worth presenting, just to give the students an overview; I spent the first class doing that. Chapter 2 rated two classes; one for an overview, and one for the proof of Theorem 2.3.2. I spent a few weeks on Chapter 3; it contains the definition of schemes, and a thorough understanding is worthwhile. In the case of Chapters 4, 5 and 6, I stated the results and mostly skipped the proofs. The results are about the complex topology and the sheaf of holomorphic functions on schemes of finite type over $\mathbb{C}$. Everything is plausible enough and, in my opinion, sketchy arguments sufficed. Now that the book is available, the need to cover this material thoroughly is even less than it was when

I was teaching the course; the students would be able fill in the details as carefully as they want, of any of the parts which you choose to omit in class.

I spent most of the semester on Chapters 7 and 8. Chapter 7 introduces coherent sheaves, both algebraic and analytic, while Chapter 8 gives a little window on geometric invariant theory, and uses it to study the properties of projective space. My feeling was that these were topics well worth treating a little more completely.

By the time you have finished with Chapter 8 you should have a reasonable idea how much time you have left. Based on that, I would decide how much of Chapter 9 to cover; most of the chapter can be omitted. Time permitting, I would present a little of the material on representations of linear algebraic groups, treat the special case of the multiplicative group $\mathbb{G}_m$, and present the proof of Hilbert's theorem on the finite generation of rings of invariants. Just how much is presented, and how many of the proofs, would very much depend on how pressed I were for time. The chapter is very skippable.

If the ground is thoroughly prepared, then Chapter 10 should not demand much time. It is the punchline of the course and is worth doing well. I do have to admit, however, that in both of my attempts to teach the material I skipped the computational parts; I stated, without proof, the results of Lemma 10.6.4. This meant, among other things, that I basically skipped all of Section 10.6. This is a much more significant omission than it may seem; if you choose to dispense with the proofs of Section 10.6 then you can also leave out all the preparatory material, which takes up several earlier sections in the book. For example there would be no compelling need to say much about the Fréchet topology, of the vector space of sections of a coherent analytic sheaf. And the various computations, of what I tend to refer to as the "elementary" or "concrete" examples, all become optional. My rough estimate is that it renders about 15% of the book dispensable.

Let me explain that I did not pull this figure out of thin air; the book used to be approximately this much shorter. At the suggestion of the reviewer I expanded the book to include Section 10.6, together with all the preparatory material. I am grateful to the reviewer for proposing this improvement; it certainly makes the treatment far more self-contained.

The first draft of the book was completed in 2005. Then, in 2006, I had an unusual fourth-year, undegraduate student. Michael Carmody wanted to do his senior thesis with me, but he told me, in advance, that this would be his last mathematical year. He had decided that his passion was for philosophy. This meant that, after his fourth year finished, his intention was to start a PhD in philosophy.

This made him another student ideal for this type of book. Giving him a solid grounding in the field, the sort that would prepare him for research, was not a priority; it seemed far more appropriate to present him with a panoramic view. I therefore gave him the manuscript of this book to read. He took about four months to get through it. We met every couple of weeks. At these meetings he would present me with long lists of misprints, as well as with some points in the mathematics which he found unclear. I took his comments extremely seriously; whenever he found anything confusing, I would rewrite the text to elucidate the point. I owe him a tremendous debt for his help. Anyway, I was pleased that he managed to plough through the book, almost unaided, in about four months. It meant that the book is accessible to its intended audience.

Before I end the Preface I should thank the many people whose help has been invaluable. Let me begin with the students who took the courses; I have already thanked Michael Carmody, but special mention goes also to Joanne Hall, Jason Lo and Kester Tong, whose questions helped inform what I wrote. I would like to thank my colleagues Eugene Lerman and Shahar Mendelson for making me teach the course a second time, only one year after the first, and for many comments on the manuscript as it was being written. I am grateful also to my (graduate) student Daniel Murfet, for pointing out several notational inconsistencies. Thanks go also to Boris Chorny, Jonathan Manton and Greg Stevenson, who kept sending me corrections, as well as ideas for more substantial improvements, right up until the very last minute. I wish to thank the anonymous reviewer for some wonderful suggestions, and my editor, Diana Gillooly, for her patient help and good humor. Speaking of patience and good humor: I would like offer my warmest thanks to my family, for putting up with me during the months when this manuscript was being written. Thank you Terry (my wife), and Ted, Joe and Jeremy (our sons). Jeremy was always ready with a good joke. For example: when Cambridge University Press and I were in the process of choosing a title for the book, it was Jeremy who finally said that, if we really wanted the book to sell, then we should name it *Harry Potter and Algebraic Geometry.* Diana Gillooly improved this to *Harry Potter and the Proof of GAGA,* which sounds more mysterious. And Jeremy also proposed that, if I wanted to lighten the mood by having a joke somewhere in the book, it could start out with: "Three algebraic varieties walk into a bar..."

Enough of the Preface; let the Quidditch match begin.

Contents

1
Introduction

Algebraic geometry is an old subject. There are many introductory books about it, at various levels. There are even some introductory texts which, like the present one, are addressed primarily to advanced undergraduates or beginning graduate students. The idea of these elementary introductions is to sell the subject. We do not yet attempt to train people in the field, only to convince them that it is fascinating and well worth the effort required to learn it.

There is no doubt that learning algebraic geometry entails substantial effort. The modern way of approaching the subject makes use of several technical machines, and a well-trained algebraic geometer needs to master at least one of these machines, preferably more than one. The very elementary introductions to the field try to avoid the machinery. They are generally very classical, using mathematics from the nineteenth century and the first half of the twentieth, before algebraic geometry underwent the Industrial Revolution and became so mechanized. The classical introductory books talk a great deal about curves, using the Riemann-Roch theorem to study them. They also might deal a little with simple singularities and their resolutions.

There are also many excellent books which do a thorough job teaching the foundations. These are for the serious graduate student, who already knows that this is the subject in which she wants to write her PhD. The current book is addressed to the uncertain graduate student, who is trying to decide if she really wants to spend the next four years of her life learning how to use sheaf cohomology to solve problems in algebraic geometry. The idea of the book is not to avoid the machinery, but rather to give an impressive illustration of its power.

The current book goes right for the mechanical apparatus, and tries to persuade the beginner of its value. We chose a theorem whose proof needs the machine. We chose an interesting, powerful theorem, and present a proof of it. Since we want to impress the reader with the value of the machine we explain the proof fairly completely, developing along the way the parts of the machine needed in the proof. But, bearing in mind that we want the book to be accessible to beginners, we keep the mechanical parts minimal. We only develop those parts of the machine which are unavoidable in the proof.

Let us therefore immediately make our disclaimer: this is not the right book for a serious graduate student, preparing herself for a research career in algebraic geometry. The treatment we give here, of the foundations of the subject, is much too patchy and has far too many glaring holes. For a thorough introduction to the foundations the student is referred to the many excellent books pitched at a higher level.

The theorem we chose to prove in this book is Serre's GAGA theorem. The GAGA stands for *Géométrie algébrique et géométrie analytique,* the title of [7], the 1956 paper by Jean-Pierre Serre containing the proof. In the remainder of the introduction we will try to explain what the theorem is about, and why it is surprising and important. There is, however, a problem: we do not yet have the language to state the theorem in the generality in which Serre proved it. Instead we will state three theorems, all of which are essentially immediate consequences of Serre's GAGA. We can state the consequences already, and will try to explain their importance. As the book progresses we will develop the language necessary for the more general theorem, and most of the machinery needed in the proof. We should make one more disclaimer: there will be some small parts of the proof we will not fully explain. We will give references and say something about the ideas.

Before we proceed any further we should say something about the prerequisites for reading the book. Let us start with the algebraic prerequisites: the reader is assumed to know what are commutative rings and ideals, what are ring homomorphisms, and the relation between ideals and kernels of ring homomorphisms; in particular given a ring R and an ideal $I \subset R$ the reader should know how to form the ring homomorphism $R \longrightarrow R/I$. The reader should also know what it means for an ideal to be either prime or maximal. Given a field k, our most important example of a ring will be the polynomial ring $k[x_1, x_2, \ldots, x_n]$; the reader is assumed to know this ring. The reader should also know about modules, homomorphisms between modules and exact sequences of modules. When we need anything more sophisticated we will state the results we need and give references. We made sure to appeal only to facts in commutative algebra which are contained in Atiyah and Macdonald's book [1], and the vast majority of what we need may be found in the first few chapters.

We also assume a basic familiarity with homological algebra. Atiyah and Macdonald's book [1] does not cover this, but there are many excellent accounts of the subject; see for example Weibel [9]. In most of the book we use very little; we assume a familiarity with exact sequences, an acquaintance with the 5–lemma, and an ability to do simple diagram chasing to prove sequences exact. Towards the end of the book the level of sophistication goes up a notch; starting in Section 10.4 we will look at the homology of chain complexes, we will appeal to the result that homotopic maps of chain complexes induce equal maps in homology, and we will rely on the fact that a short exact sequence of chain complexes gives a long exact sequence in homology.

The reader is assumed to know some basic point-set topology; we will freely refer to topological spaces, open and closed sets and continuous maps. We expect the reader to know what a homeomorphism is. Given a topological space X and a subset U, the reader should know what is the subspace topology on U. Given a topological space X, a set Y and surjective map $X \longrightarrow Y$, the reader should know how to form the quotient topology on Y. We will also feel free to talk about connectedness and compactness of topological spaces, and about Hausdorff spaces, and we will assume the reader knows what all these concepts mean.

And finally, while it is not absolutely indispensable, it would help to have seen at least one course on functions of one or more complex variables. We will freely refer to "holomorphic functions", and occasionally also to "meromorphic functions". For a reader who has never met them before let us briefly introduce them. Let $U \subset \mathbb{C}^n$ be an open set. A holomorphic function on U is a function $f : U \longrightarrow \mathbb{C}$ so that, for every point $p \in U$, the Taylor series of the function f at the point p converges near p to the function f. More formally this means that for every point $p \in U$ there exists a real number $\delta > 0$ so that, on the ball of radius δ centered at $p = (p_1, p_2, \ldots, p_n)$, the function f is given by a convergent power series

$$f(z_1, z_2, \ldots, z_n) = \sum_{(i_1, i_2, \ldots, i_n) \in \mathbb{N}^n} a_{i_1, i_2, \ldots, i_n} (z_1 - p_1)^{i_1} (z_2 - p_2)^{i_2} \cdots (z_n - p_n)^{i_n} .$$

This defines holomorphic functions. The ratio f/g of two holomorphic functions f and g, where g does not vanish on any open subset of U, is called a meromorphic function. Taken very literally it is not a function; if f and g are holomorphic on U then f/g is only defined on the subset of U where g does not vanish; in an abuse of language we say that f/g is meromorphic on U. In general, a meromorphic function h on U is only assumed to be locally of the form f/g. That is, every point $p \in U$ is contained in a ball $V \subset U$ so that, on the ball V, there exist holomorphic functions f_V and $g_V \neq 0$, and f_V/g_V agrees with the restriction to V of h.

1.1 Algebraic and analytic subspaces

Consider the closed subsets of $\mathbb{C}^n$, the n–dimensional complex space. There are many closed subsets. We are particularly interested in two classes of closed subsets: the closed algebraic and the closed analytic subspaces. Closed algebraic subspaces of $\mathbb{C}^n$ are the closed sets on which a finite number of polynomials vanish. Closed analytic subspaces are, at least locally, the closed subsets on which finitely many holomorphic functions vanish. Here are the definitions, given more precisely.

Definition 1.1.1. *A closed subset $X \subset \mathbb{C}^n$ is called a* closed algebraic subspace *if there are finitely many polynomial functions on $\mathbb{C}^n$, let us say $f_1, f_2, \ldots, f_r$, so that*

$$X = \{x \in \mathbb{C}^n \mid f_i(x) = 0 \quad \forall\, 1 \leq i \leq r\}.$$

The definition of analytic subspaces is slightly more delicate, being local.

Definition 1.1.2. *A closed subset $X \subset \mathbb{C}^n$ is called a* closed analytic subspace *if, for every point $x \in X$, there exists an open neighborhood U of x in $\mathbb{C}^n$, and finitely many holomorphic functions $f_1, f_2, \ldots, f_r$ on U, so that*

$$X \cap U = \{y \in U \mid f_i(y) = 0 \quad \forall\, 1 \leq i \leq r\}.$$

Note that, in the definition of a closed analytic subset, we do not require the existence of finitely many global holomorphic functions. The definition is local. Furthermore it is immediate, from the definitions, that any closed algebraic subspace of $\mathbb{C}^n$ must automatically be a closed analytic subspace.

It is also very clear that there are many closed analytic subspaces of $\mathbb{C}^n$ which are not algebraic. The easiest is to consider the case $n = 1$. If $f \neq 0$ is a polynomial function in one variable (that is a polynomial function on $\mathbb{C}^1$), then the set of points $\{x \in \mathbb{C}^1 \mid f(x) = 0\}$ is finite. This just says that a polynomial in one variable has finitely many roots. If X is a closed algebraic subset of $\mathbb{C}^1$, Definition 1.1.1 tells us that there is a finite set of polynomial functions $f_1, f_2, \ldots, f_r$, and X is the set of points $p \in \mathbb{C}$ at which all the f_i vanish. If the f_i are all zero then $X = \mathbb{C}^1$. Otherwise one of the f_i must be non-zero, it vanishes only at finitely many roots, and X is a subset of this finite set of roots. We conclude that a closed algebraic subspace $X \subset \mathbb{C}^1$ is either finite or is equal to $\mathbb{C}^1$.

Now consider the subset $X \subset \mathbb{C}^1$ given by

$$X = \{x \in \mathbb{C}^1 \mid \sin(x) = 0\}.$$

Since $\sin(x)$ is a holomorphic function of $x \in \mathbb{C}^1$, the set X is a closed analytic subspace of $\mathbb{C}^1$. On the other hand we know that $\sin(x)$ vanishes whenever x is a multiple of π; even better, it vanishes no place else. That is

$$X = \{n\pi \mid n \in \mathbb{Z}\}.$$

The set X is neither finite nor all of $\mathbb{C}^1$; it therefore is not algebraic.

In Definitions 1.1.1 and 1.1.2 we defined closed algebraic and analytic subspaces of $\mathbb{C}^n$. For the definitions to make sense we needed to have a clear notion of which functions on $\mathbb{C}^n$ are polynomial, and which are holomorphic. In the definitions we can replace $\mathbb{C}^n$ by any space P provided that P has, at least locally, well defined classes of polynomial and holomorphic functions.

There is a family of such spaces, which come equipped with classes of polynomial and holomorphic functions. In the first few chapters of the book we will define them and talk a little about their elementary properties. Such spaces are called *schemes of finite type over* $\mathbb{C}$. For the remainder of the introduction we ask the reader to accept, without seeing the formal definitions, that schemes of finite type over $\mathbb{C}$ are some topological spaces which have a well-defined notion of which functions are polynomial and which are holomorphic. For every P, a scheme of finite type over $\mathbb{C}$, it makes sense to ask which closed subsets $X \subset P$ are algebraic subspaces, and which are analytic. The first important corollary of GAGA is a theorem which predates Serre's paper, a theorem whose first proof was due to Chow:

Theorem 1.1.3. *Let P be a scheme of finite type over $\mathbb{C}$. Assume P is compact. Then any closed analytic subspace of P is algebraic.*

We saw above that not all closed analytic subspaces of $\mathbb{C}^n$ are algebraic. This does not contradict Theorem 1.1.3, since $\mathbb{C}^n$ is not compact. One very interesting space to which Theorem 1.1.3 applies is the complex projective space $\mathbb{CP}^n$. I do not want to define $\mathbb{CP}^n$ in the introduction;

the readers who know it do not need such a definition, the readers who do not yet know $\mathbb{CP}^n$ will learn much about it in the rest of the book.

It turns out that the general statement of Theorem 1.1.3 is easy enough to reduce to the special case where $P = \mathbb{CP}^n$. There is a lemma due to Chow that says an arbitrary compact P, which is a scheme of finite type over $\mathbb{C}$, admits a surjective, polynomial map $\pi : X \longrightarrow P$, with X a closed algebraic subset of $\mathbb{CP}^n$. Suppose $Y \subset P$ is a closed analytic subspace. The special case of Theorem 1.1.3, where the space is $\mathbb{CP}^n$, tells us that $\pi^{-1}Y \subset X \subset \mathbb{CP}^n$ is closed and algebraic, and then fairly easy, standard arguments imply that $Y = \pi(\pi^{-1}Y)$ is an algebraic subspace of P. In this book we will confine ourselves to proving the special case of Theorem 1.1.3 with $P = \mathbb{CP}^n$; the reader is expected to remember that the more general fact is an easy consequence.

The next consequence of GAGA is about vector bundles. Given a topological space P it is possible to define vector bundles on P. We will not remind the reader of the definition of a vector bundle in the introduction; there will be a great deal said about vector bundles, and more general sheaves, later in the book. The only thing we want to recall here is that one way to construct a vector bundle on P is in terms of transition functions, as follows. Take an open cover $\{U_i, i \in \mathscr{I}\}$ of the topological space P. For each pair U_i, U_j of open sets in the cover give a function

$$\varphi_{ij} : U_i \cap U_j \longrightarrow \mathrm{GL}(n, \mathbb{C}).$$

The functions φ_{ij} are called the *transition functions.* If the φ_{ij} satisfy the identities

$$\varphi_{ij}\varphi_{ji} = 1, \qquad \varphi_{ij}\varphi_{jk}\varphi_{ki} = 1$$

then the data is enough to specify a vector bundle on P. If P happens to be a scheme of finite type over $\mathbb{C}$ it makes sense to speak of polynomial functions and it makes sense to speak of holomorphic functions. It therefore makes sense to speak of vector bundles where the transition functions are polynomial, and of vector bundles where the transition functions are holomorphic. The ones with polynomial transition functions are called *algebraic vector bundles,* while the ones with holomorphic transition functions are called *analytic vector bundles.* The next theorem says

Theorem 1.1.4. *Let P be a scheme of finite type over $\mathbb{C}$. Assume P is compact. Let $\mathscr{V}$ be an analytic vector bundle on P. Then $\mathscr{V}$ is isomorphic to an algebraic vector bundle.*

And the last theorem asserts

Theorem 1.1.5. *Let P be a scheme of finite type over $\mathbb{C}$. Assume P is compact. Let $\mathscr{V}$ and $\mathscr{V}'$ be algebraic vector bundles on P, and let $\varphi : \mathscr{V} \longrightarrow \mathscr{V}'$ be an analytic map of vector bundles. Then φ is algebraic.*

Now that we have stated the theorems we should explain their import. To the extent that algebraic geometry concerns itself with schemes of finite type over $\mathbb{C}$, and with vector bundles over these schemes, we have learnt that, as long as the scheme P is compact, we can study it either by algebraic or by analytic methods. This is actually a very powerful, important observation. There are theorems we know how to prove by analytic methods without having an algebraic proof, and theorems for which the only known proof is algebraic. It is a little strange: we end up proving theorems in analysis using commutative algebra, and theorems in algebra using partial differential equations.

The most beautiful, intriguing parts of mathematics are those which lie at the confluence of different fields. Algebraic geometry is one of these. To be a good algebraic geometer one needs to be aware of both the algebraic and the analytic approaches to the subject. It also does not hurt to know some number theory; we will not, in this book, describe the interaction between algebraic geometry and number theory.

To be completely honest I should tell the reader that, even before Serre, it was known that algebraic geometry was a subject which lay at the intersection of algebra and analysis. Riemann knew this in the nineteenth century. Let me very briefly remind the reader of the relation between the algebraic and analytic approaches to elliptic curves; everything I say, in the remainder of the introduction, was known in the nineteenth century. None of the remainder of the introduction will be used in the rest of the book; the only purpose is to give the reader a very old, explicit example of the way algebra and analysis interact in algebraic geometry.

1.2 Elliptic curves

Let us give ourselves two complex numbers a and b, in general position (meaning not real multiples of each other). We want to consider functions periodic with the two periods a and b, that is

$$f(z+a) = f(z) = f(z+b).$$

These functions go by the name of *doubly periodic functions on* $\mathbb{C}$. Doubly periodic functions are determined by their values on a fundamental domain. There is a parallelogram in the complex plane with sides a and