0. Introduction

In this introductory chapter we will explain in a very rough sketch what algebraic geometry is about and what it can be used for. We will stress the many correlations with other fields of research, such as complex analysis, topology, differential geometry, singularity theory, computer algebra, commutative algebra, number theory, enumerative geometry, and even theoretical physics. The goal of this chapter is just motivational; you will not find definitions or proofs here (and maybe not even a mathematically precise statement). In the same way, the exercises in this chapter are not designed to be solved in a mathematically precise way. Rather, they are just given as some “food for thought” if you want to think a little further about the examples presented here.

To start from something that you probably know, we can say that algebraic geometry is the combination of *linear algebra* and *algebra*:

- In linear algebra (as e. g. in the “Foundations of Mathematics” class [G2]), we study systems of linear equations in several variables.
- In algebra (as e. g. in the “Introduction to Algebra” class [G3]), one of the main topics is the study of polynomial equations in one variable.

Algebraic geometry combines these two fields of mathematics by studying systems of polynomial equations in several variables.

Given such a system of polynomial equations, what sort of questions can we ask? Note that we cannot expect in general to write down explicitly all the solutions: we know from algebra that even a single complex polynomial equation of degree $d > 4$ in one variable can in general not be solved exactly [G3, Problem 0.2]. So we are more interested in statements about the geometric structure of the set of solutions. For example, in the case of a complex polynomial equation of degree d, even if we cannot compute the solutions we know that there are exactly d of them (if we count them with the correct multiplicities). Let us now see what sort of “geometric structure” we can find in polynomial equations in several variables.

Example 0.1. Probably the easiest example that is covered neither in linear algebra nor in algebra is that of a single polynomial equation in two variables. Let us consider the example

$$C_n = \{(x_1,x_2) \in \mathbb{C}^2 : x_2^2 = (x_1-1)(x_1-2)\cdots(x_1-2n)\} \quad \subset \mathbb{C}^2,$$

where $n \in \mathbb{N}_{>0}$. Note that in this case it is actually possible to write down all the solutions, because the equation is (almost) solved for x_2 already: we can pick x_1 to be any complex number, and then get two values for x_2 — unless $x_1 \in \{1,\dots,2n\}$, in which case there is only one value for x_2 (namely 0).

So it seems that C_n looks like two copies of the complex plane, with the two copies of each point $1,\dots,2n$ identified: the complex plane parametrizes the values for x_1, and the two copies of it correspond to the two possible values for x_2, i. e. the two roots of the number $(x_1-1)\cdots(x_1-2n)$.

This is not the correct topological picture however, because a complex non-zero number does not have a distinguished first and second root that could correspond to the first and second copy of the complex plane. Rather, the two roots of a complex number get exchanged if you run around the origin once: if we consider a closed path

$$z = r\,e^{i\varphi} \qquad \text{for } 0 \le \varphi \le 2\pi \text{ and fixed } r > 0$$

around the complex origin, the square root of this number would have to be defined by

$$\sqrt{z} = \sqrt{r}\,e^{\frac{i\varphi}{2}}$$

which gives opposite values at $\varphi = 0$ and $\varphi = 2\pi$. In other words, if in C_n we run around one of the points $1, \dots, 2n$ (i. e. around a point at which x_2 is the square root of 0), we go from one copy of the plane to the other. The way to draw this topologically is to cut the two planes along the real intervals $[1,2], \dots, [2n-1, 2n]$, and to glue the two planes along these lines as in this picture for $n = 3$ (lines marked with the same letter are to be identified):

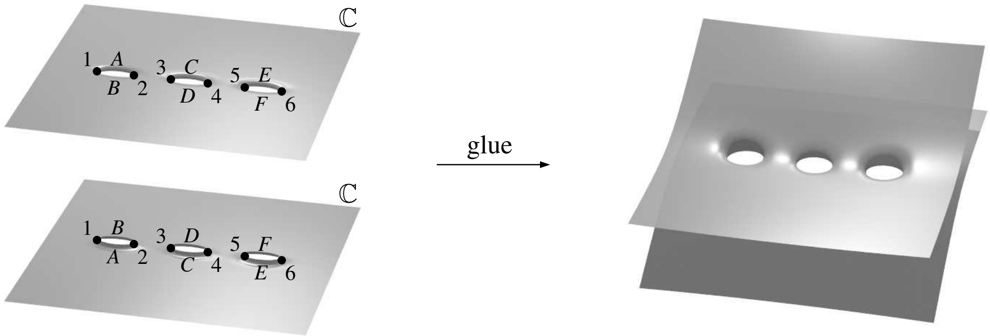

To make the picture a little nicer, we can compactify our set by adding two points at infinity — one for each copy of the plane — in the same way as we can compactify the complex plane $\mathbb{C}$ by adding a point ∞. The precise construction of this compactification will be given in Example 5.6. If we do this here, we end up with a compact surface with $n-1$ handles:

Such an object is called a surface *of genus* $n-1$; the example above shows a surface of genus 2.

Example 0.2. What happens in the previous Example 0.1 if we move the points $1, \dots, 2n$ for x_1 at which we have only one value for x_2, i. e. if we consider

$$C_n = \{(x_1, x_2) \in \mathbb{C}^2 : x_2^2 = f(x_1)\} \quad \subset \mathbb{C}^2$$

with f some polynomial in x_1 of degree $2n$? Obviously, as long as the $2n$ roots of f are still distinct, the topological picture above does not change. But if two of the roots approach each other and finally coincide, this has the effect of shrinking one of the tubes connecting the two planes until it reduces to a "singular point" (also called a *node*), as in the following picture on the left:

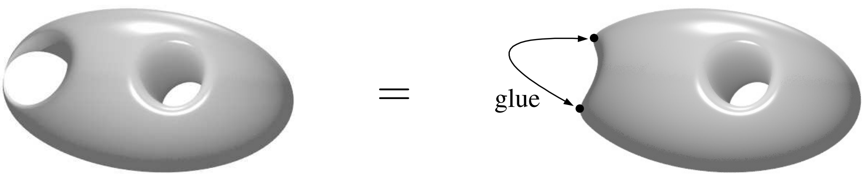

Obviously, we can view this as a surface with one handle less, where in addition we identify two of the points (as illustrated in the picture on the right). Note that we can still see the "handles" when we draw the surface like this, just that one of the handles results from the gluing of the two points.

Note that our examples so far were a little "cheated" because we said before that we want to figure out the geometric structure of equations that we cannot solve explicitly. In the examples above however, the polynomial equation was chosen so that we could solve it, and in fact we used this solution to construct the geometric picture. Let us now see what we can still do if we make the polynomial more complicated.

Example 0.3. Let $d \in \mathbb{N}_{>0}$, and consider

$$C_d = \{(x_1,x_2) \in \mathbb{C}^2 : f(x_1,x_2) = 0\} \quad \subset \mathbb{C}^2,$$

where f is an arbitrary polynomial of degree d. This is an equation that we certainly cannot solve directly if f is sufficiently general. Can we still deduce the geometric structure of C_d?

In fact, we can do this with the idea of Example 0.2. We saw there that the genus of the surface does not change if we perturb the polynomial equation, even if the surface acquires singular points (provided that we know how to compute the genus of such a singular surface). So why not deform the polynomial f to something singular that is easier to analyze? Probably the easiest thing that comes into mind is to degenerate the polynomial f of degree d into a product of d linear equations $l_1,\dots,l_d$: consider

$$C'_d = \{(x_1,x_2) \in \mathbb{C}^2 : l_1(x_1,x_2) \cdot \cdots \cdot l_d(x_1,x_2) = 0\} \quad \subset \mathbb{C}^2,$$

which should have the same "genus" as the original C_d.

It is easy to see what C'_d looks like: of course it is just a union of d complex planes. Any two of them intersect in a point, and we can certainly choose them so that no three of them intersect in a point. The picture below shows C'_d for $d = 3$ (note that every complex plane is — after compactifying it with a point at infinity — just a sphere).

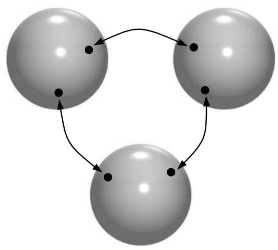

What is the genus of this surface? In the picture above it is obvious that we have one loop; so if $d = 3$ we get a surface of genus 1. In the general case we have d spheres, and every two of them connect in a pair of points, so in total we have $\binom{d}{2}$ connections. But $d-1$ of them are needed to glue the d spheres to a connected chain without loops; only the remaining ones then add a handle each. So the genus of C'_d (and hence of C_d) is

$$\binom{d}{2} - (d-1) = \binom{d-1}{2}.$$

This is commonly called the *degree-genus formula* for complex plane curves. We will show it in Proposition 13.17.

Remark 0.4 (Real vs. complex dimension)**.** One of the trivial but common sources for misunderstandings is whether we count dimensions over $\mathbb{C}$ or over $\mathbb{R}$. The examples considered above are *real surfaces* (the dimension over $\mathbb{R}$ is 2), but *complex curves* (the dimension over $\mathbb{C}$ is 1). We have used the word "surface" so far as this fitted best to the pictures that we have drawn. When looking at the theory however, it is usually more natural to call these objects (complex) curves. In what follows, we always mean the dimension over $\mathbb{C}$ unless stated otherwise.

Exercise 0.5. What do we get in Example 0.1 if we consider the equation

$$C'_n = \{(x_1,x_2) \in \mathbb{C}^2 : x_2^2 = (x_1-1)(x_1-2)\cdots(x_1-(2n-1))\} \subset \mathbb{C}^2$$

for $n \in \mathbb{N}_{>0}$ instead?

Exercise 0.6. In Example 0.3, we argued that a polynomial of degree d in two complex variables gives rise to a surface of genus $\binom{d-1}{2}$. In Example 0.1 however, a polynomial of degree $2n$ gave us a surface of genus $n-1$. Can you see why these two results do not contradict each other?

Remark 0.7. Here is what we should learn from the examples considered so far:

- Algebraic geometry can make statements about the topological structure of objects defined by polynomial equations. It is therefore related to *topology* and *differential geometry* (where similar statements are deduced using analytic methods).
- The geometric objects considered in algebraic geometry need not be "smooth" (i. e. they need not be *manifolds*). Even if our primary interest is in smooth objects, degenerations to singular objects can greatly simplify a problem (as in Example 0.3). This is a main point that distinguishes algebraic geometry from other geometric theories (e. g. differential or symplectic geometry). Of course, this comes at a price: our theory must be strong enough to include such singular objects and make statements how things vary when we pass from smooth to singular objects. In this regard, algebraic geometry is related to *singularity theory* which studies precisely these questions.

Remark 0.8. Maybe it looks a bit restrictive to allow only algebraic (polynomial) equations to describe our geometric objects. But in fact it is a deep theorem that for *compact* objects, we would not get anything different if we allowed *holomorphic* equations too. In this respect, algebraic geometry is very much related (and in certain cases identical) to *complex (analytic) geometry*. The easiest example of this correspondence is that a holomorphic map from the compactified complex plane $\mathbb{C}\cup\{\infty\}$ to itself must in fact be a rational map, i. e. a quotient of two polynomials.

Example 0.9. Let us now turn our attention to the next more complicated objects, namely complex surfaces in 3-dimensional space. We just want to give one example here: let X be the *cubic surface*

$$X = \{(x_1,x_2,x_3) : 1+x_1^3+x_2^3+x_3^3-(1+x_1+x_2+x_3)^3 = 0\} \quad \subset \mathbb{C}^3.$$

As X has real dimension 4, it is impossible to draw pictures of it that reflect its topological properties correctly. Usually, we overcome this problem by just drawing the *real* part, i. e. we look for solutions of the equation over the real numbers. This then gives a real surface in $\mathbb{R}^3$ that we can draw. We should just be careful about which statements we can claim to see from this incomplete geometric picture.

The following picture shows the real part of the surface X.

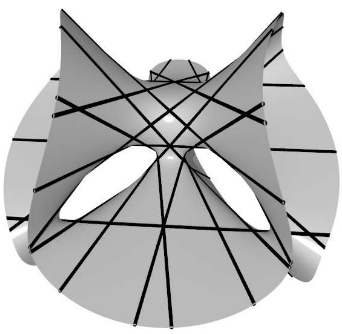

In contrast to our previous examples, we have now used a *linear* projection to map the real 3-dimensional space onto the drawing plane (and not just a topologically correct picture).

We see that there are some lines contained in X. In fact, one can show that (after a suitable compactification) *every* smooth cubic surface has exactly 27 lines on it, see Chapter 11. This is another sort of question that one can ask about the solutions of polynomial equations, and that is not of topological nature: do they contain curves with special properties (in this case lines), and if so, how many? This branch of algebraic geometry is usually called *enumerative geometry*.

Remark 0.10. It is probably surprising that algebraic geometry, in particular enumerative geometry, is very much related to *theoretical physics*. In fact, many results in enumerative geometry have been found by physicists first.

Why are physicists interested e. g. in the number of lines on the cubic surface? We try to give a short answer to this (that is necessarily vague and incomplete): There is a branch of theoretical physics called *string theory* whose underlying idea is that the elementary particles (electrons, quarks, ...) might not be point-like, but rather 1-dimensional objects (the so-called strings), that are just so small that their 1-dimensional structure cannot be observed directly by any sort of physical measurement. When these particles move in time, they sweep out a surface in space-time. For some reason this surface has a natural complex structure coming from the underlying physical theory.

Now the same idea applies to space-time in general: string theorists believe that space-time is not 4-dimensional as we observe it, but rather has some extra dimensions that are again so small in size that we cannot observe them directly. (Think e. g. of a long tube with a very small diameter — of course this is a 2-dimensional object, but if you look at this tube from very far away you cannot see the small diameter any more, and the object looks like a 1-dimensional line.) These extra dimensions are parametrized by a space that sometimes has a complex structure too; it might for example be the complex cubic surface that we looked at above.

So in this case we are in fact looking at complex curves in a complex surface. A priori, these curves can sit in the surface in any way. But there are equations of motion that tell you how these curves will sit in the ambient space, just as in classical mechanics it follows from the equations of motion that a particle will move on a straight line if no forces apply to it. In our case, the equations of motion say that the curve must map *holomorphically* to the ambient space. As we said in Remark 0.8 above, this is equivalent to saying that we must have algebraic equations that describe the curve. So we are looking at exactly the same type of questions as we did in Example 0.9 above.

Example 0.11. Let us now have a brief look at curves in 3-dimensional space. Consider the example

$$C = \{(x_1,x_2,x_3) = (t^3,t^4,t^5) : t \in \mathbb{C}\} \quad \subset \mathbb{C}^3.$$

We have given this curve parametrically, but it is in fact easy to see that we can describe it equally well in terms of polynomial equations:

$$C = \{(x_1,x_2,x_3) : x_1^3 = x_2x_3,\ x_2^2 = x_1x_3,\ x_3^2 = x_1^2x_2\}.$$

What is striking here is that we have *three* equations, although we would expect that a 1-dimensional object in 3-dimensional space should be given by two equations. But in fact, if you leave out any of the above three equations, you are changing the set that it describes: if you leave out e. g. the last equation $x_3^2 = x_1^2x_2$, you would get the whole x_3-axis $\{(x_1,x_2,x_3) : x_1 = x_2 = 0\}$ as additional points that do satisfy the first two equations, but not the last one.

So we see another important difference to linear algebra: it is not clear whether a given object of codimension d can be given by d equations — in any case we have just seen that it is in general not possible to choose d defining equations from a given set of such equations. Even worse, for a given set of equations it is in general a difficult task to figure out what dimension their solution has. There do exist algorithms to find this out for any given set of polynomials, but they are so complicated that you will in general want to use a computer program to do that for you. This is a simple example of an application of *computer algebra* to algebraic geometry.

Exercise 0.12. Show that if you replace the three equations defining the curve C in Example 0.11 by

$$x_1^3 = x_2x_3,\ x_2^2 = x_1x_3,\ x_3^2 = x_1^2x_2 + \varepsilon$$

for a (small) non-zero number $\varepsilon \in \mathbb{C}$, the resulting set of solutions is in fact 0-dimensional, as you would expect from three equations in 3-dimensional space. So we see that very small changes in the equations can make a very big difference in the resulting solution set. Hence we usually cannot apply numerical methods to our problems: very small rounding errors can change the result completely.

Remark 0.13. Especially the previous Example 0.11 is already very algebraic in nature: the question that we asked there does not depend at all on the ground field being the complex numbers. In fact, this is a general philosophy: even if algebraic geometry describes geometric objects (when viewed over the complex numbers), most methods do not rely on this, and therefore should be established in purely algebraic terms. For example, the genus of a curve (that we introduced topologically in Example 0.1) can be defined in purely algebraic terms in such a way that all the statements from complex geometry (e. g. the degree-genus formula of Example 0.3) extend to this more general setting. Many geometric questions then reduce to pure *commutative algebra*, which is in some sense the foundation of algebraic geometry.

Example 0.14. The most famous application of algebraic geometry to ground fields other than the complex numbers is certainly Fermat's Last Theorem: this is just the statement that the algebraic curve *over the rational numbers*

$$C = \{(x_1, x_2) \in \mathbb{Q}^2 : x_1^n + x_2^n = 1\} \quad \subset \mathbb{Q}^2$$

contains only the trivial points where $x_1 = 0$ or $x_2 = 0$. Note that this is very different from the case of the ground field $\mathbb{C}$, where we have seen in Example 0.3 that C is a curve of genus $\binom{n-1}{2}$. But a large part of the theory of algebraic geometry applies to the rational numbers (and related fields) as well, so if you look at the proof of Fermat's Theorem you will notice that it uses e. g. the concepts of algebraic curves and their genus at many places, although the corresponding point set C contains only some trivial points. So, in some sense, we can view *(algebraic) number theory* as a part of algebraic geometry.

With this many relations to other fields of mathematics (and physics), it is obvious that we have to restrict our attention in this class to quite a small subset of the possible applications. Although we will develop the general theory of algebraic geometry, our focus will mainly be on geometric questions, neglecting number-theoretic aspects most of the time. So, for example, if we say "let K be an algebraically closed field", feel free to read this as "let K be the complex numbers" and think about geometry rather than algebra.

Every now and then we will quote results from or give applications to other fields of mathematics. This applies in particular to commutative algebra, which provides some of the basic foundations of algebraic geometry. So to fully understand algebraic geometry, you will need to get some background in commutative algebra as well, to the extent as covered e. g. in [AM] or [G5]. However, we will not assume this here — although this is probably not the standard approach it is perfectly possible to follow these notes without any knowledge of commutative algebra. To make this easier, all commutative algebra results that we will need will be stated clearly (and easy to understand), and you can learn the algebraic techniques to prove them afterwards. The only algebraic prerequisite needed for this class is some basic knowledge on groups, rings, fields, and vector spaces as e. g. taught in the "Algebraic Structures" and "Foundations of Mathematics" courses [G1, G2].