0. Introduction

These notes are meant as a gentle introduction to *algebraic geometry*, a combination of *linear algebra* and *algebra*:

(a) In linear algebra (as e. g. in the "Foundations of Mathematics" class [G2]), we study systems of linear equations in several variables over a fixed ground field K.

(b) In algebra (as e. g. in the "Algebraic Structures" or "Introduction to Algebra" classes [G1, G3]), a central topic are polynomials in one variable over K.

Algebraic geometry combines this by studying systems of polynomial equations in several variables over K. Of course, such polynomials in several variables occur in many places both in pure mathematics and in applications. Consequently, algebraic geometry has become a very large and active field of mathematics with deep connections to many other areas, such as commutative algebra, computer algebra, number theory, topology, and complex analysis, just to name a few.

On the one hand, all these connections make algebraic geometry into a very interesting field to study — but on the other hand they may also make it hard for the beginner to get started. So to keep everything digestible, we will restrict ourselves here to the first case that is covered by neither (a) nor (b) above: *one polynomial equation in two variables*. Its set of solutions in K^2 can then be thought of as a curve in the plane, we can draw it (at least in the case $K = \mathbb{R}$), ask geometric questions about it, and try to answer them with algebraic methods. This restriction will significantly reduce the required theoretical background, but still leads to many interesting results that we will discuss in these notes.

To get a feeling for the kind of problems that one may ask about plane curves, we will now mention a few of them in this introductory chapter. Their flavor differs a bit depending on the chosen ground field K.

Example 0.1 (Curves over $\mathbb{R}$). The following picture shows some real plane curves. Note that they can have many different "shapes": The curve (a) lies in a bounded region of the plane, whereas the others do not. The curve (b) consists of two components in the sense that it can be decomposed into two subsets (given by $x + y = 0$ and $x - y = 0$) that are given by polynomial equations themselves. The curve (c) has a so-called singularity at the origin, i. e. a point where it does not locally look like a smoothly deformed real line (in fact, (b) has a singularity at the origin as well). Finally, the image in (d) consists of two disconnected parts, but these parts are *not* given by separate polynomial equations themselves, as we will see in Exercise 1.8.

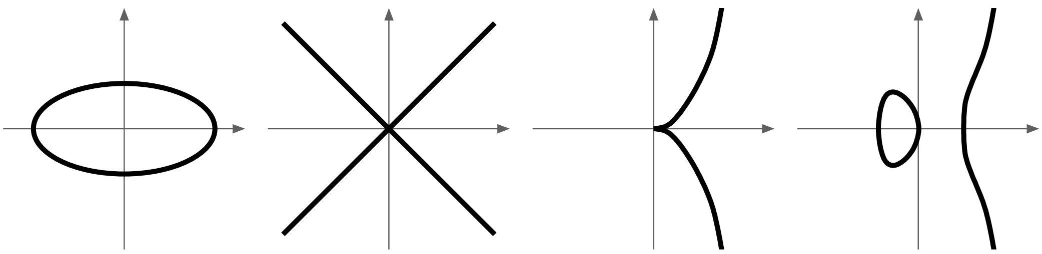

$\frac{x^2}{4} + y^2 - 1 = 0$ (a) $\quad x^2 - y^2 = 0$ (b) $\quad y^2 - x^3 = 0$ (c) $\quad y^2 + x - x^3 = 0$ (d)

It is a main goal of algebraic geometry to prove such properties of curves just from looking at the polynomials, i. e. without drawing and referring to a picture (which would not be an exact proof anyway). Other related questions we might ask are: In how many points can two curves intersect? How many singularities can a curve have?

Example 0.2 (Curves over $\mathbb{C}$)**.** Over the complex numbers, the pictures of curves will look different, since a 1-dimensional complex object is real 2-dimensional, i. e. a surface. Note that we cannot draw such a surface as a subset of $K^2 = \mathbb{C}^2 = \mathbb{R}^4$ since we would need four dimensions for that. But we can still get a correct topological picture of the curve itself if we disregard this embedding. Let us show informally how to do this for the curve with the equation $y^2 + x - x^3 = 0$ as in Example 0.1 (d) above; for more details see 5.16.

Note that in this case it is actually possible to write down all the points of the curve explicitly, because the given equation

$$y^2 = x^3 - x = x(x-1)(x+1)$$

is (almost) solved for y already: We can pick x to be any complex number, and then get two values for y, namely the two square roots of $x(x-1)(x+1)$ — unless $x \in \{-1,0,1\}$, in which case there is only one value for y (namely 0).

So one might think that the curve looks like two copies of the complex plane, glued together at the three points $-1,0,1$: The complex plane parametrizes the values for x, and the two copies of it correspond to the two possible values for y, i. e. to the two roots of the number $x(x-1)(x+1)$.

This is not the correct topological picture however, because a non-zero complex 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 we run around the origin once: If we consider a closed path

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

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 x runs around one of the points $-1,0,1$ (i. e. around a point at which y is the square root of 0), we go from one copy of the plane to the other. One way to draw this topologically is to cut the two planes along the real intervals $(-1,0)$ and $(1,\infty)$, and to glue the two planes along these edges as in the following picture on the left, where edges with the same letter are meant to be identified. The gluing itself is then visualized best by first turning one of the planes upside down; this is shown in the picture on the right.

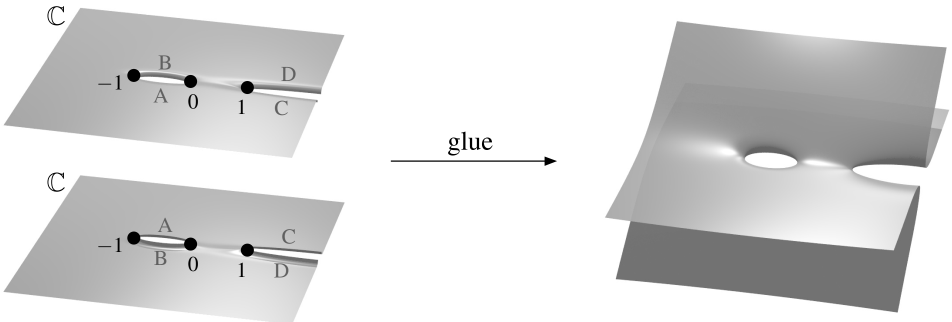

This is now actually a topologically correct picture of the given curve. To make the situation a little nicer, we can compactify it by adding a point at infinity, which corresponds to identifying the two planes at their infinitely far points as well (the precise construction will be described in Chapter 3). This is shown in the picture below, and leads topologically to a torus.

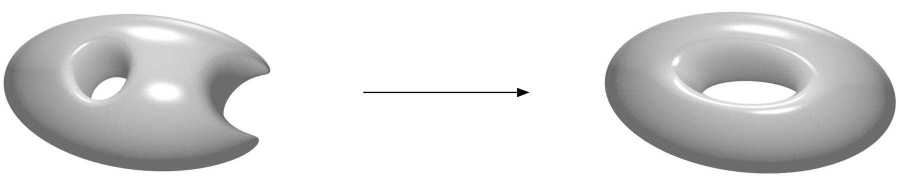

We will show in Proposition 5.16 how such topological pictures can be obtained immediately from the given equation of the curve.

Example 0.3 (Curves over $\mathbb{Q}$)**.** The most famous application of algebraic geometry to ground fields other than the real or complex numbers is certainly Fermat's Last Theorem: This is just the statement that, for $n \in \mathbb{N}_{\geq 3}$, the curve given by the equation $x^n + y^n - 1 = 0$ *over the rational numbers* has only the trivial solutions where $x = 0$ or $y = 0$, or equivalently (by setting $x = \frac{a}{c}$ and $y = \frac{b}{c}$ for $a, b, c \in \mathbb{Z}$ with $c \neq 0$), that the equation $a^n + b^n = c^n$ has no non-trivial solutions over $\mathbb{Z}$. Note that this picture is very different from the case of the ground fields $\mathbb{R}$ or $\mathbb{C}$ above. But a large part of the theory of algebraic curves applies to the rational numbers as well, and in fact the proof of Fermat's Last Theorem uses concepts of the theory of algebraic curves in many places. So, in some sense, we can view (algebraic) number theory as a part of algebraic geometry.

Example 0.4 (Relations to complex analysis)**.** We have just seen in the examples above that algebraic geometry has deep relations to topology and number theory, and it should not come as a surprise that there are many relations to algebraic fields of mathematics such as commutative algebra and computer algebra as well. Although it is not within the scope of these notes, let us finish this introductory chapter by showing interesting relations to complex analysis as well.

Consider a (sufficiently nice) compactified complex curve, such as a torus as in Example 0.2. Of course, in algebraic geometry one does not only study curves for themselves but also maps between them; and hence we will have to consider "nice" functions on such curves (where "nice" will translate into "locally a quotient of polynomials"). What do such functions f look like if they are defined globally on the whole curve? As the curve is compact, note that the image of f must be a compact subset of the complex plane, which means that the absolute value $|f|$ must take a maximum somewhere. But locally the curve just looks like the complex plane, and by the Maximum Modulus Principle [G4, Proposition 6.14] the absolute value of a nice (read: holomorphic) function on the complex plane cannot have a local maximum unless it is constant. So we conclude that f must be a constant function: There are actually no non-trivial nice global functions on a compact curve.

In fact, we will prove this statement in Corollary 6.29 using only algebraic methods, and hence over arbitrary (algebraically closed) ground fields. In a similar way, many interesting results over the ground field $\mathbb{C}$ can be obtained using both algebraic geometry and complex analysis, with completely different methods, and thus give a close relation between these two branches of mathematics as well.

But let us now start with our study of plane curves. In order to keep these notes as accessible as possible, we will only assume a basic knowledge of groups, rings, and fields as about to the extent of the "Algebraic Structures" class [G1], but a little more experience in dealing with these structures would certainly be advantageous. Very occasionally we will need to assume results from commutative algebra that go beyond these prerequisites (marked as "Facts" in the notes), but they will always be clearly stated and motivated, and provided with a reference. However, in order not to lose this very interesting part of the subject we will nevertheless quite frequently explore the relations of our results to other fields of mathematics in side remarks and excursions (that will then not be needed afterwards to follow the remaining parts of the notes).