1

Preliminaries

1.1 What is a plane curve?

There are many ways to give a precise interpretation of the informal notion of a plane curve as something drawn on a sheet of paper. Here we will explain just what will be meant in this book, and indicate why this is a fruitful choice.

We begin with thinking of the plane as consisting of points, which may be described by two coordinates (x, y). Then one basic idea is that a curve is a set of points whose coordinates satisfy some equation $f(x, y) = 0$. What kinds of function f are to be allowed? The simplest kind, where f is just a polynomial function – expressible, say, as $f(x, y) = \sum_{0 \le i,j; i+j \le d} a_{i,j} x^i y^j$ for some numerical coefficients $a_{i,j}$ – already leads to an extensive theory of curves, mostly developed in the nineteenth century. We take this as our starting point.

The next thing to decide is what sort of numbers are to be allowed for the coefficients $a_{i,j}$ and indeed for the values of x and y. For the kind of curve one may draw and picture most easily, it is natural to choose to allow arbitrary real numbers: these permit the kind of continuity one expects from a sketch of a curve. However, it turns out that a much richer structure is obtained if we allow complex coefficients, and although the geometry involved when there are two complex variables is harder to picture, there have been major advances in this type of geometry in recent years. One of our main concerns will be understanding the interrelation between these geometrical aspects of our curves with invariants defined from a more algebraic viewpoint. We will occasionally restrict to real coefficients; results over other fields will be mentioned sometimes in the notes at the end of chapters.

The topic of this book is not so much the study of entire curves, though we will obtain a number of important results valid for the curve as a whole, as a very detailed study of what may happen near a singular point of a curve. Thus we will mainly study curves just in a neighbourhood of the origin O (with coordinates $(0,0)$) in the complex plane $\mathbb{C}^2$. For this local study it is not important that the equation $f(x,y) = \sum_{0 \leq i,j} a_{i,j} x^i y^j$ defining the curve should be given by a finite sum. So we need to consider more carefully just what kind of expression f is to be: let us give some precise definitions.

A *polynomial* in one variable t is a sum of finitely many terms $a_n t^n$, where the *coefficients* a_n are complex numbers and the *exponents* n are non-negative integers. The *degree* of the polynomial is the largest number n such that a_n is non-zero. Polynomials may be added and multiplied in the usual way, and form a ring, denoted $\mathbb{C}[t]$. An immediate and important consequence of our choice of complex numbers as coefficients is that polynomials can be factorised: if $f = \sum_0^n a_r x^r$ has degree n, then the *roots* α_i of $f(x) = 0$ are such that $f(x) \equiv a_n \prod_1^n (x - \alpha_i)$.

An expression $\sum_0^\infty a_r t^r$, where infinitely many non-zero terms are allowed, is called a *formal power series*. The usual rules still allow us to add and multiply such expressions, yielding a much larger ring, denoted $\mathbb{C}[[t]]$. The *order* of the formal power series is the smallest number n such that a_n is non-zero.

It is easy to write down a formal power series such that if any non-zero complex number is substituted for t, the resulting series of complex numbers fails to converge. A simple example is $\sum_0^\infty r! t^r$. We recall from complex variable theory that if $\sum_0^\infty a_r t^r$ converges at $t = v$, then the terms $|a_n| R^n$ (where $R = |v|$) are bounded (indeed, they tend to 0), and that if conversely this condition holds then the series converges for all values of t such that $|t| < R$. If this condition holds for some value of $R > 0$, we call the series a *convergent power series*.

A function of one (or several) complex variable(s) defined on some region $U \subset \mathbb{C}$ (or $U \subset \mathbb{C}^n$) and which possesses a derivative on U is said to be *holomorphic* on U. Standard complex variable theory tells us that any function of t which is holomorphic on some neighbourhood of O can be expanded as a convergent power series (some books use the word *analytic* to describe functions defined by convergent power series, so the result can be stated as 'analytic = holomorphic'). Given two holomorphic functions, each defined on some neighbourhood of O, we can add and multiply them (on a smaller neighbourhood of O); it follows that the convergent power series form a subring of $\mathbb{C}[[t]]$. It is denoted by $\mathbb{C}\{t\}$.

Correspondingly for functions of two variables we have the polynomials, which are sums of finitely many terms $\sum a_{i,j} x^i y^j$ (with i and j non-negative integers and the coefficients $a_{i,j}$ complex numbers), and form a ring denoted $\mathbb{C}[x, y]$. We have the formal power series, which are expressions $\sum_{i=0}^{\infty}\sum_{j=0}^{\infty} a_{i,j}x^i y^j$ and form a ring, denoted $\mathbb{C}[[x, y]]$. The *degree* of a polynomial is the largest number n such that $a_{i,j}$ is non-zero for some i, j with $i + j = n$. The *order* of a power series is the smallest number n such that $a_{i,j}$ is non-zero for some i, j with $i + j = n$.

A power series in two variables is said to be *convergent* if there exist positive real numbers R, S such that the numbers $|a_{m,n}|R^m S^n$ are bounded. In this case, the series is convergent on the region $|x| < R$, $|y| < S$, and summing it defines a holomorphic function on this region. Conversely, a function which is holomorphic on a neighbourhood of O is said to be holomorphic at O and can be expanded as a convergent power series. Such series form a ring, denoted $\mathbb{C}\{x, y\}$.

Each convergent power series converges on some neighbourhood of $t = 0$. Formal power series do not in general converge at any $t \neq 0$: a compensating advantage is that they can be constructed term-by-term. Thus a sequence $\{f_k(t)\}$ of polynomials (or power series) such that, for each n, the coefficient of t^n in $f_k(t)$ is the same for all large enough values of k, defines a formal power series $f_\infty(t)$ having these coefficients. We say that f_k converges to f_∞ in the $\mathfrak{m}_t$-*adic sense*.

As a simple example, observe that a series $1 - a(t)$ with constant term 1 has an inverse, since the series $\sum_0^\infty a(t)^r$ converges in the $\mathfrak{m}_t$-adic sense, and its sum is the desired inverse. Thus any series with non-zero constant term also has an inverse in the ring. This result is also true in the ring $\mathbb{C}\{x, y\}$ for the simpler reason that if the function $f(x, y)$ is differentiable in a neighbourhood of O and $f \neq 0$ at – and hence in some neighbourhood of – O, then the usual rule allows us to differentiate also $\frac{1}{f(x,y)}$ in such a neighbourhood.

Thus in the 1-variable cases $\mathbb{C}[[x]]$ and $\mathbb{C}\{x\}$, any element of order m is equal to x^m multiplied by a power series with non-zero constant term, which has an inverse in the ring. This gives a complete description of factorisation in these rings.

Another way to approach the idea of plane curves is by parametrisations. The most familiar example is that of a graph, where y is expressed as a function of x: in general we have a parameter t, with each of x and y expressed in terms of t – say $x = \phi(t)$, $y = \psi(t)$. As above, the functions ϕ and ψ may be taken as polynomials or power series (preferably convergent).

The starting point for the analysis of singular points is the solution of a holomorphic equation $f(x, y) = 0$ to express y as a function of x. This will be discussed in the next chapter. It follows from the implicit function theorem 1.4.2 that if $\partial f/\partial y$ is non-zero at O, we can solve for

y as a holomorphic function of x near O. In this case, the curve Γ given by the equation $f(x,y)=0$ is said to be *non-singular* or (more briefly) *smooth* at the origin; its tangent there is given by $x\frac{\partial f}{\partial x}(O)+y\frac{\partial f}{\partial y}(O)=0$. The example

$$f(x,y)\equiv x+y-xy;\quad y=-\sum_{1}^{\infty}x^r$$

shows that even if f is a polynomial we will need power series, not just polynomials, to express y as a function of x. The example

$$f(x,y)\equiv x^2-y^3;\quad y=x^{\frac{2}{3}}$$

shows that in general fractional powers of x will be involved.

This last example may be represented by a parametrisation $(x,y)=(t^3,t^2)$. We also have the parametrisation $(x,y)=(u^6,u^4)$, obtained by substituting $t=u^2$. Clearly the latter is less satisfactory: each point of the curve is represented by two values of u, differing in sign. We will say that a parametrisation $(x,y)=(\phi(t),\psi(t))$ is *good* (an alternative term is 'primitive') if a general point of the curve corresponds to just one value of the parameter, i.e. the map $t\to(\phi(t),\psi(t))$ is injective on some region $|t|<\epsilon$.

Thus an equation $f(x,y)=0$ with $f\in\mathbb{C}[x,y]$, or a parametrisation with $\phi,\psi\in\mathbb{C}[t]$, defines a curve Γ as a subset of the plane $\mathbb{C}^2$. If we merely have $f\in\mathbb{C}\{x,y\}$ or, respectively, $\phi,\psi\in\mathbb{C}\{t\}$, then there is a neighbourhood U of O in $\mathbb{C}^2$ on which the series converges, so the equation $f(x,y)=0$ defines a curve in U. We will see in Chapter 2 that we obtain the same class of holomorphic curves whether we use equations or parametrisations.

It is convenient to introduce some terminology for this situation. Two functions $f_i:U_i\to\mathbb{C}$ $(i=1,2)$, defined on neighbourhoods U_1,U_2 of O in $\mathbb{C}^2$ are said to define the same *germ* at O if they coincide on some neighbourhood $U\subset U_1\cap U_2$ of O. In the case of holomorphic functions, this is the case if and only if the power series expansions of f_1 and f_2 coincide. Correspondingly, subsets $X_i\subset U_i$ of neighbourhoods U_1,U_2 of O define the same *germ* at O if for some neighbourhood $U\subset U_1\cap U_2$ of O we have $X_1\cap U=X_2\cap U$. In practice, rather than use the word 'germ', we will speak of curves defined in some neighbourhood of a point, usually O, and always be prepared to pass to smaller neighbourhoods.

The above discussion is concentrated on what happens in a small neighbourhood of the origin in the plane $\mathbb{C}^2$. Sometimes we wish to think of curves in the large, and then it is more convenient to work in the projective plane (we may refer to $\mathbb{C}^2$, when we wish to emphasise the distinction, as the *affine plane*).

In general we define n-dimensional *projective space* $P^n(\mathbb{C})$ to be the set of lines through the origin in $\mathbb{C}^{n+1}$. We may take coordinates $\mathbf{x} = (x_0, \dots, x_n)$ in $\mathbb{C}^{n+1}$; then any point other than the origin (so with coordinate vector $\mathbf{x} \neq \mathbf{0}$ in $\mathbb{C}^{n+1}$) is joined to the origin by a unique line, and if $\mathbf{x}$ and $\mathbf{y}$ lie on the same such line, then $\mathbf{y} = \lambda\mathbf{x}$ for some $\lambda \neq 0$. Thus the point is determined by the *ratios* of the coordinates: $(x_0 : x_1 : \dots : x_n)$. Note also that we may (if it is convenient to do so) change coordinates by any linear transformation of $\mathbb{C}^{n+1}$ onto itself.

If $f(x_0, \dots, x_n)$ is some function, then when we replace $\mathbf{x}$ by $\mathbf{y} = \lambda\mathbf{x}$ we will get a different function, so the condition $f(x_0, \dots, x_n) = 0$ is not in general well defined on $P^n(\mathbb{C})$. It is so is f is *homogeneous*, that is, if $f(\lambda x_0, \dots, \lambda x_n) \equiv \lambda^k f(x_0, \dots, x_n)$ for some value of k, the degree of f. In general, when working in projective space, only homogeneous polynomial functions f are considered.

Ratios are not always convenient to work with. Observe that the subsets U_r of $P^n(\mathbb{C})$ given (for $0 \leq r \leq n$) by $x_r \neq 0$ are well defined. Since we had $\mathbf{x} \neq \mathbf{0}$ above, any point in $P^n(\mathbb{C})$ lies in at least one of these subsets. In the subset U_r we may omit x_r and take the ratios $z_s = x_s/x_r \ \ (s \neq r)$ as coordinates in the usual sense, or equivalently, fix $x_r = 1$. Thus each U_r is isomorphic to the affine space $\mathbb{C}^n$. Observe that if f is any polynomial function on U_r, and d is the highest degree of any term in f, then we may define a homogeneous polynomial function on $P^n(\mathbb{C})$ by

$$F(x_0, \dots, x_n) = x_r^d f(z_0, \dots, z_{r-1}, \uparrow^r, z_{r+1}, \dots, z_n).$$

A projective algebraic variety is a subset of $P^n(\mathbb{C})$ defined by some homogeneous polynomial equations. We will be particularly interested in subsets of $P^2(\mathbb{C})$ defined by a single such equation: *projective plane curves*. The solutions of $f = 0$ coincide with the solutions of $f^2 = 0$. It will usually be convenient to insist that we consider only equations with no squared factor (these are called *reduced*). If $f = 0$ is a reduced equation, homogeneous of degree d, for a curve Γ, then Γ is said to be of degree d, and we write $d = \deg \Gamma$.

A useful background reference for the first three sections (and some later ones) of this book is the student text [99] by Kirwan, which assumes less background than we do.

1.2 Intersection numbers

We will have frequent occasion to consider intersections of curves, and to count intersection numbers. If the curves Γ_1 and Γ_2 are both smooth at the origin and have distinct tangents there, we will say that they have intersection number 1 at O. In general, we seek to deform one or both the curves a small amount so that at each point of intersection of the resulting curves, both are smooth and they have distinct tangents. Thus for example, for the curves given by $y = 0$ and $y = x^2$, we deform the latter to $y = x^2 - t^2$, giving two intersection points $(t, 0)$ and $(-t, 0)$, so the intersection number is 2.

More generally, consider the curves $y = 0$ and $y = f(x)$ in $\mathbb{C}^2$. If f is a polynomial, we may factorise it as $A\prod_1^n(x - t_i)$, where the roots t_i need not all be distinct. Since a small deformation will make them so, if just r of the t_i take a given value T, the intersection number at $(T, 0)$ is equal to r. If f is not a polynomial, but can be expressed as a power series $f(x) = \sum_0^\infty a_r x^r$ of order m, then we may write $f(x) = x^m g(x)$ with $g(0) \neq 0$, and then the intersection number at O is equal to m. Sometimes we will also refer to the order m as the *multiplicity* of 0 as a root of f.

To count intersections of $y = 0$ with a general curve Γ given by an equation $g(x, y) = 0$, we substitute $y = 0$ in the equation to obtain $g(x, 0)$ and proceed as above with $g(x, 0)$ in place of $f(x)$. Unless g has a repeated factor, Γ will intersect a general line $y = \epsilon$ in distinct points, which (for small ϵ) provide a deformation of the intersection $\Gamma \cap \{y = 0\}$ as before.

Suppose we have a curve Γ_1 given by an equation $g(x, y) = 0$ and a curve Γ_2 given by a good parametrisation $(x, y) = (\phi(t), \psi(t))$ such that $\phi(0) = \psi(0) = 0$. Then the intersection number of Γ_1 and Γ_2 at O is equal to the order of $g(\phi(t), \psi(t))$. For a first perturbation of g will ensure that g does not vanish at singular points of Γ_2, and for intersections at non-singular points we can take local coordinates in which Γ_2 is given by $y = 0$ with parameter $t = x$, and then argue as above. In future we will often find it convenient to write the parametrisation as a single map $\gamma : \mathbb{C} \to \mathbb{C}^2$, with $\gamma(0) = O$.

We will denote the intersection number of Γ_1 and Γ_2 at a point P by $(\Gamma_1.\Gamma_2)_P$ or, if P is understood, simply by $\Gamma_1.\Gamma_2$. The above unsymmetrical rule for calculating intersection numbers is very convenient, but does not cover all needs. For a general discussion of intersection numbers, including precise definitions, the reader may refer to [188], [99] or, for a more advanced treatment, [74]. We now list some key properties.

Lemma 1.2.1

(i) *Let C_1, C_2 be germs of holomorphic curves at a point $P \in \mathbb{C}^2$, with P an isolated point of $C_1 \cap C_2$, then there is a well defined intersection number $C_1.C_2$, which is a positive integer.*

(ii) *If C_1 has equation $g_1(x,y) = 0$ and C_2 has a good parametrisation $\gamma : \mathbb{C} \to \mathbb{C}^2$, with $\gamma(0) = O$, then $C_1.C_2$ is equal to the order of $g(\gamma(t))$.*

(iii) *If g_1 factorises as $g_3 g_4$ so that $C_1 = C_3 \cup C_4$, then $C_1.C_2 = C_3.C_2 + C_4.C_2$.*

(iv) *If C_i has equation $g_i(x,y) = 0$ for $i = 1, 2$ and $P = O$, then $C_1.C_2$ is equal to the dimension of the quotient ring $\mathbb{C}\{x,y\}/\langle g_1, g_2 \rangle$.*

(v) *Intersection numbers are symmetric: $C_1.C_2 = C_2.C_1$.*

(vi) *Suppose given a holomorphic map $F : \mathbb{C}^2 \to \mathbb{C}^2$, defined near O, with $F(O) = O$, a good parametrisation $\gamma_2 : \mathbb{C} \to \mathbb{C}^2$ with $\gamma(0) = O$ of a curve C_2 such that $F \circ \gamma_2$ is a good parametrisation of $F(C_2)$, and an equation g_1 of a curve C_1, so that we can define a curve $F^{-1}C_1$ by the equation $g_1 \circ F$. Then $C_1.FC_2 = F^{-1}C_1.C_2$.*

Proof We omit (or defer) the proof that (ii) and (iv) give the same result, which is independent of all choices. Then (i) follows from either version of the definition, (iii) follows from (ii) and (v) from (iv).

(vi) is an immediate consequence of (ii), since each of the intersection numbers is equal to the order of the composite function $g_1 \circ F \circ \gamma_2$. □

Consider a curve Γ of degree d in the projective plane $P^2(\mathbb{C})$. A general line in the plane intersects Γ in d points, and (unless Γ contains the whole line as a subset) any line has intersections with Γ whose intersection numbers add up to d. For we may choose coordinates $(x : y : z)$ so that the line is given by $y = 0$ and Γ does not pass through $(1, 0, 0)$. Then if $f(x, y, z) = 0$ is the equation of Γ, the intersections are given by the vanishing of $f(x, 0, 1)$, which is a polynomial of degree d, and the sum of the multiplicities of the roots of such a polynomial is equal to d.

More generally, suppose Γ_1, Γ_2 are two projective plane curves, of respective degrees d_1, d_2, and such that their intersection does not contain a curve. Then the sum of the intersection multiplicities at all points of $\Gamma_1 \cap \Gamma_2$ is equal to $d_1 d_2$. This is known as Bézout's theorem; a proof may be found in [99]; see also the following section.

1.3 Resultants and discriminants

Suppose $f(x,y), g(x,y)$ are homogeneous polynomials of respective degrees m, n. Then

$$x^{n-s-1}y^s f(x,y)\ (0 \le s < n), \quad x^{m-r-1}y^r g(x,y)\ (0 \le r < m)$$

are $m+n$ homogeneous polynomials of degree $m+n-1$, so their coefficients form a square matrix. Its determinant is denoted $R(f,g)$ and called the *resultant* of f and g. If f and g have a common factor (e.g. $(x-\lambda y)$), then this divides all the above polynomials, so these are linearly dependent, the matrix is singular, and $R(f,g)=0$. In fact we have the following well-known results: see e.g. [18]1.4.4; see also [182] Sections 26–28 for this section.

Lemma 1.3.1 *For polynomials* $f(x,y) = c\prod_1^m (x-a_i y)$, $g(x,y) = c'\prod_1^n (x-b_j y)$ *in factorised form, with* $c, c' \neq 0$, *we have:*

(i) $R(f,g) = c^n c'^m \prod_{i,j}(a_i - b_j)$;
(ii) $R(g,f) = (-1)^{mn} R(f,g)$;
(iii) $R(f,gh) = R(f,g)R(f,h)$;
(iv) $R(f,g) = c^n \prod_1^m g(a_i, 1)$;
(v) *if* $\deg\phi = \deg f - \deg g$, $R(f+g\phi, g) = R(f,g)$;
(vi) $R(f,g) = 0$ *if and only if* f *and* g *have a common factor.*

We have used the fact that we are working over the field $\mathbb{C}$ to factorise our polynomials. But the definition and the formulae (ii), (iii), (v) are valid if the coefficients are taken in any commutative ring. The conclusion (vi) also holds, provided that we work over a unique factorisation domain. General results about unique factorisation in commutative rings may be found in elementary algebra texts, e.g. in [182] Section 19.

For a single homogeneous polynomial $f(x,y)$, we can form the resultant of $\partial f/\partial x$ and $\partial f/\partial y$: the result $D(f) = R(\partial f/\partial x, \partial f/\partial y)$ is called the *discriminant* of f. This also has important properties.

Lemma 1.3.2 *If* $f(x,y) = c\prod_1^m (x-a_i y)$, *with* $c \neq 0$, *we have*

(i) $D(f) = m^{m-2}c^{2m-2}\prod_{i\neq j}(a_i - a_j)$, *and hence*
(ii) $D(f) = 0$ *if and only if* f *has a repeated factor.*

Indeed since $mf = x\partial f/\partial x + y\partial f/\partial y$, we have

$$R(mf, \partial f/\partial x) = R(y\partial f/\partial y, \partial f/\partial x) = (-1)^{m-1}R(y, \partial f/\partial x)D(f),$$

and as $R(y, \partial f/\partial x) = mc$, this reduces to $(-1)^{m-1}mcD(f)$.

By substituting $y = 1$ we can regard these as results about polynomials in a single variable. Although this version is more familiar, there is a certain ambiguity since, for example, the function $f = bx^2 + cx + d$ may be regarded as a quadratic, or as a special case of a cubic $f = ax^3 + bx^2 + cx + d$ where the coefficient a happens to be 0. This ambiguity disappears if we insist, as we often will, that the coefficient of the highest power of x is 1. Such polynomials are called *monic*. Thus for monic polynomials, $D(f) = (-1)^{m-1}m^{m-2}R(f, df/dx)$.

We will sometimes be interested in the situation where the coefficients of f and g depend on a further parameter. Consider for example two homogeneous polynomials $f(x, y, z)$, $g(x, y, z)$ of respective degrees p and q. Substituting yt for y and zt for z makes them homogeneous in the two variables x and t; forming the resultant as above gives a homogeneous polynomial $P(y, z)$ of degree pq in y and z. The roots of $P(y, 1) = 0$ are those values y_0 of y for which the polynomials $f(x, y_0, 1)$, $g(x, y_0, 1)$ have a common root, and thus correspond to the intersections of the curves $f(x, y, z) = 0 = g(x, y, z)$. One proof of Bézout's theorem consists in counting these intersections carefully to see that indeed the multiplicities correspond to those of the roots of P.

1.4 Manifolds and the Implicit Function Theorem

One of the main objectives of the book is to explore the topology of plane curve singularities, so it will be necessary from Chapter 5 on to assume that the reader knows the rudiments of topology. There are numerous textbooks on this subject: for example the beginner's text [9] and the rather detailed exposition [169]. Although we need only elementary algebraic topology, the concept of manifold is important to us, and we now recall some basic facts.

A space X is an n-dimensional manifold if every point of X has a neighbourhood U_α such that there is a homeomorphism (a 'chart') $\phi_\alpha : U_\alpha \to V_\alpha$ where V_α is an open set in $\mathbb{R}^n$. We have coordinate transformations defined on the overlaps: if all of these are differentiable, more precisely, C^∞, then X is a differentiable manifold. If Y_1 and Y_2 are subsets of X which are differentiable manifolds, and the identity map of $Y_1 \cap Y_2$ is smooth in terms of charts of Y_1 on one side and of Y_2 on the other, then the collection of all the charts gives X the structure of differentiable manifold. This construction is known as glueing.

A differentiable n-manifold X has a well defined tangent space $T_P(X)$ at each point P, which is an n-dimensional vector space; there are several equivalent definitions. The notion of differentiable, or smooth, map is defined using charts and requiring differentiability in the local coordinate systems. A smooth map $f : X \to Y$ induces a linear map $T_P f : T_P(X) \to T_{f(P)}(Y)$ of corresponding tangent spaces. In local coordinates, if $f : \mathbb{R}^m \to \mathbb{R}^n$ is defined and differentiable at O, its partial derivatives $\partial f_i/\partial x_j$ form an $m \times n$ matrix, the *Jacobian matrix* of f, which we denote Df_O, which is the matrix of the map $T_O f$.

A bijection $f : X \to Y$ such that both f and f^{-1} are smooth maps is called a diffeomorphism of X to Y. Diffeomorphism is the basic equivalence relation between smooth manifolds.

A smooth map $f : X \to Y$ is called a smooth embedding if it is injective and, for each $P \in X$, the induced map $T_P(X) \to T_{f(P)}(Y)$ is also injective (strictly speaking, if X is non-compact one adds a further requirement to ensure that at each point of $f(X)$ there is a chart of Y in which X corresponds to a linear subspace). A smooth embedding $S^1 \to S^3$ is called a (smooth) knot. A link is a finite collection of knots with disjoint images, so can be taken as a smooth embedding $A \times S^1 \to S^3$ with A a finite set.

A manifold with boundary is defined in the same way as a manifold except that charts may map to open subsets of $\mathbb{R}^n_+ := \{(x_1, \ldots, x_n) \in \mathbb{R}^n \mid x_n \geq 0\}$. The boundary is the part corresponding in these charts to the subset where $x_n = 0$: it is an $(n-1)$ dimensional manifold. All the above extend naturally to this case.

The key to discussing changes of coordinates is the Inverse Function Theorem, which is proved in many textbooks, e.g. [52].

Theorem 1.4.1 *Let U be a neighbourhood of $O \in \mathbb{R}^n$; let $f : U \to \mathbb{R}^n$ be differentiable, and suppose $T_O f$ an isomorphism. Then there is a neighbourhood $U_1 \subset U$ of O such that $f|_{U_1}$ is a bijection of U_1 with a neighbourhood V_1 of $f(O)$ and its inverse $f^{-1} : V_1 \to U_1$ is again differentiable.*

An important application is the Implicit Function Theorem, which gives a first general result for proceeding from a subset of Euclidean space defined by equations to one given by a parametrisation.