Number fields

Avi Zeff
Seminar on Fermat's last theorem

The goal of today's talks is to give an introduction to some key objects in algebraic number theory. In particular we want to understand algebraic extensions of the rational numbers $\mathbb{Q}$, meaning everything between $\mathbb{Q}$ and its algebraic closure $\overline{\mathbb{Q}}$ (the algebraic numbers). Via Galois theory, such extensions are often studied via their automorphism groups (Galois groups), and in turn such groups are often studied via their representations (Galois representations); so we'll give brief introductions to Galois theory and representation theory along the way. This will be useful when we talk about elliptic curves, and ultimately modularity: one of the most important properties of elliptic curves from a number-theoretic point of view is that they give rise to interesting Galois representations.

1 Absolute values and Ostrowski's theorem

The rational numbers $\mathbb{Q}$ have the structure of a field: we can add, subtract, and multiply any two elements, and divide by any element other than 0. They have a further structure though which is very useful for talking about number theory: an absolute value $|\cdot| : \mathbb{Q} \to \mathbb{R}$, satisfying various good properties.

We might ask: is this absolute value unique? What other absolute values could we potentially put on $\mathbb{Q}$? To make sense of this, we first need to say what exactly we mean by an absolute value in an abstract way.

Let R be any ring (so we can add, subtract, and multiply, but not necessarily divide). An absolute value $|\cdot|$ on R is a map $R \to \mathbb{R}$ such that

- $|x| \geq 0$ for all $x \in R$;
- $|x| = 0$ if and only if $x = 0$;
- $|x| \cdot |y| = |xy|$;
- $|x + y| \leq |x| + |y|$.

(Note that this definition differs slightly from that of the textbook.)

The last condition is called the triangle inequality. Some absolute values satisfy a stronger inequality, called the strong triangle inequality or the ultrametric inequality $|x + y| \leq \max(|x|, |y|)$ (this is stronger than the triangle inequality by the nonnegativity condition). Absolute values satisfying the ultrametric inequality are called ultrametric or nonarchimedean; sometimes absolute values which satisfy the triangle inequality but not the ultrametric inequality are called archimedean.

There is always a trivial absolute value given by $|x| = 1$ for every $x \neq 0$ and $|0| = 0$, which also satisfies the ultrametric inequality. However it is usually not interesting and we typically exclude it from consideration.

One can check that the standard absolute value on $\mathbb{Q}$ satisfies the axioms above, but not the ultrametric inequality, so it is an archimedean absolute value. One can now ask: other

than the standard absolute value and the trivial one, are there any other absolute values on $\mathbb{Q}$?

The answer is yes for rather trivial reasons: e.g. $x \mapsto |x|^{1/2}$ is also an absolute value (all the conditions are clear except the triangle inequality, where one has to check that $|x+y|^{1/2} \leq |x|^{1/2}+|y|^{1/2}$, which follows from the fact that $f(x) = \sqrt{x}$ has sublinear growth). More generally $x \mapsto |x|^\alpha$ is an absolute value for each $0 \leq \alpha \leq 1$, with $\alpha = 0$ recovering the trivial absolute value; so we have an infinite family of absolute values. However these don't meaningfully differ from each other (for example they all induce the same topology and completion, for those familiar with the notions; we'll come back to completion later) so we don't really want to consider them new. Therefore we say that two absolute values $|\cdot|_1$ and $|\cdot|_2$ are equivalent if for every x we have $|x|_1 = |x|_2^\alpha$ for some $\alpha > 0$ (not depending on x), and consider absolute values up to equivalence. These equivalence classes (of nontrivial absolute values) are sometimes called places of R.

So thus far we only know about one place of $\mathbb{Q}$. However we can construct some more. First, observe that giving an absolute value of $\mathbb{Q}$ is really the same as giving an absolute value of $\mathbb{Z}$, since by the multiplicative property of absolute values for $a/b \in \mathbb{Q}$ we have $|a/b| = |a|/|b|$. So we want to define a new notion of "closeness" on $\mathbb{Z}$. Fixing a prime p, we're going to say that two integers x and y are close if they're congruent modulo p; even closer if they're congruent modulo p^2; and so on. To formalize this, we can look at how many factors of p divide $x-y$. Writing $v_p(x-y)$ for this number, we observe that it satisfies $v_p(xy) = v_p(x)+v_p(y)$, so to get an absolute value we need to exponentiate; thus we consider the absolute value $|x| = b^{v_p(x)}$ for some real number $b > 0$ for x nonzero, and $|0| = 0$.

We need to check that this satisfies the triangle inequality. If x and y are both divisible by p^e, then so is $|x+y|$; therefore $v_p(x) \geq \min(v_p(x), v_p(y))$. Therefore so long as $b < 1$ we have $|x+y| \leq \max(|x|, |y|)$, i.e. this is not just an absolute value but also an ultrametric one. Any choice of $0 < b < 1$ gives an equivalent absolute value; for concreteness we often make a standard choice of $b = p^{-1}$, giving the p-adic absolute value $|x|_p = p^{-v_p(x)}$. This formula also holds for all rational numbers: so if $x = a/b$ is in lowest terms with neither a nor b divisible by p, then $|x|_p = 1$; if (say) a is divisible by p once, then $|x|_p = p^{-1}$, while if b is divisible by p once then $|x|_p = p$. For example, $|p|_p = p^{-1}$.

It is easy to see that $|\cdot|_p$ is not equivalent to the standard absolute value on $\mathbb{Q}$, which we sometimes write as $|\cdot|_\infty$ to avoid confusion. Thus we have found infinitely many places of $\mathbb{Q}$: one for each prime number p (the p-adic places, with nonarchimedean absolute values) and one more (the standard archimedean absolute value). This last place is sometimes called the "prime at infinity," by analogy with the fact that all the other places are primes; one of the key ideas in a lot of modern number theory is that philosophically, for many purposes we should try and think of this "infinite place" as analogous to the finite ones, putting all the places on the same level, even though literally speaking they look quite different.

Finally we can ask if we've now found all the places; Ostrowski's theorem states that the answer is yes, i.e. every nontrivial absolute value on $\mathbb{Q}$ is equivalent to either a p-adic absolute value for some prime p or the archimedean absolute value.

These absolute values are related to each other in an interesting way: the archimedean absolute value measures the "size" of each element in a way we're used to, while the p-adic absolute values measure the size of the p-part (in an inverted way). We can formalize this

idea via the product formula: for any rational number x, we have

$$|x|_\infty \cdot \prod_p |x|_p = 1.$$

Indeed, let's think about the case where x is a positive integer; multiplying by -1 doesn't change the formula, and the extension to rational numbers is then easy as above. Each integer x can be written uniquely as the product of prime numbers with some exponent; we could think of this as

$$x = \prod_p p^{v_p(x)},$$

with $v_p(x) = 0$ for all but finitely many primes p (those which divide x). For x a positive integer, $|x_\infty| = x$, while $p^{v_p(x)} = |x|_p^{-1}$, so this is

$$|x|_\infty = \prod_p |x|_p^{-1},$$

and rearranging gives the product formula above.

There is also a version of this story for every finite extension $K/\mathbb{Q}$, with the primes replaced by prime ideals of $\mathcal{O}_K$. One can also do a version over function fields, which we won't define now but give a characteristic p analogue of number fields; the simplest example is $\mathbb{F}_p(t)$. There, all the places are nonarchimedean, but there is still a place "at infinity" given by $f \mapsto q^{-\deg f}$, and P-adic places for each irreducible polynomial P. Function fields together with number fields (i.e. finite extensions of $\mathbb{Q}$) make up *global fields*, which have similar properties despite living in two different settings; this motivates many analogies between the number field world and the function field world, and many results can be proven uniformly for both.

2 Completions

For those who have seen some topology, the following notion may be clear: if X is a metric space, i.e. it has a notion of distance $d(x, y) \in \mathbb{R}$ for points $x, y \in X$ satisfying various properties (in particular $d(x, y) = |x - y|$ for an absolute value $|\cdot|$ on a ring will work), we say that a sequence $\{x_n\}$ in X is a Cauchy sequence if for every $\epsilon > 0$, there exists an integer N such that for every $m, n \geq N$ we have $d(x_m - x_n) < \epsilon$. Note that this depends on the metric d, and so on the absolute value. We say that X is complete if every Cauchy sequence converges to some limit in X; for example, $\mathbb{R}$ (with the standard absolute value) is complete, every such sequence of real numbers will converge to a real number, but $\mathbb{Q}$ (with the standard absolute value) is not, e.g. the sequence of decimal approximations to $\sqrt{2}$, i.e. 1, 1.4, 1.41, ..., although a Cauchy sequence, will not converge to any rational number.

We are interested in the case where $X = K$ is a field equipped with an absolute value $|\cdot|$. As the example of $(\mathbb{Q}, |\cdot|_\infty)$ shows, this need not be complete. Our goal is to construct a *completion* $\widehat{K}$, which contains K as a subfield and is complete for some absolute value extending the one on K.

Consider the ring $S = \text{Functions}(\mathbb{N}, K)$ of sequences in K, with pointwise addition and multiplication. This contains a subring $C \subset S$ of Cauchy sequences (one can check that sums, differences, and products of Cauchy sequences are again Cauchy), which further contains a subset $N \subset C$ of "null sequences," i.e. sequences converging to 0 in K. One can show that in fact N is an *ideal* of C (and actually a maximal ideal): multiplying a Cauchy sequence by a null sequence gives another null sequence, and the null sequences are closed under sums and differences. Therefore we can take the quotient C/N, which is a ring and, by some algebra, actually a field; we think of this as taking Cauchy sequences and discarding the different ways that they can go to whatever their limit would be, if it existed, leaving only the data of this hypothetical limit. Thus this field $\widehat{K} := C/N$ is the *completion* of K in the sense of adding in all the limits that Cauchy sequences should have. It contains K as a subfield by looking at constant sequences in K, which are certainly Cauchy.

What about the absolute value on $\widehat{K}$? If $\{a_n\}$ is a Cauchy sequence in K, then $\{|a_n\}$ is a Cauchy sequence in $\mathbb{R}$, and so since $\mathbb{R}$ is complete it has a limit, which we define to be $|\{a_n\}|$. This gives an absolute value on $\widehat{K}$, which restricted to the constant sequences $\{x\}$ gives $|\{x\}| = |x|$, so this does extend the absolute value on $\widehat{K}$. In fact, the absolute value on K induces a topology on it (the metric topology), for which the map given by the absolute value to $\mathbb{R}$ (also with the metric topology) is continuous; and K is dense in $\widehat{K}$, so this is the unique continuous extension of the absolute value with respect to these topologies. Finally $\widehat{K}$ is now complete for this topology; this is an exercise in showing that Cauchy sequences of Cauchy sequences reduce to usual Cauchy sequences, which therefore have limits in $\widehat{K}$.

This might even be a process you've seen before: it's one way of constructing the real numbers. That is, the completion of $\mathbb{Q}$ with respect to the standard absolute value $|\cdot|_\infty$ is exactly the real numbers. Completing for different absolute values in the same equivalence class still gives the real numbers; the completion depends only on the equivalence class, i.e. on the place. This suggests the following question: what are the completions of $\mathbb{Q}$ with respect to the p-adic topologies? These are the p-adic numbers $\mathbb{Q}_p$, which we discuss next. We sometimes write $\mathbb{R}$ as $\mathbb{Q}_\infty$, to reflect that we should think of it as the completion at the infinite place, parallel to the completions $\mathbb{Q}_p$ at finite places.

3 p-adic numbers

We can now define the p-adic numbers $\mathbb{Q}_p$ as above, via completion. However while we have a lot of intuition around the real numbers, we have very little for this strange p-adic world; so let's take a moment to think about them.

The first thing to observe is that, unlike the real numbers, the p-adics are nonarchimedean, i.e. the absolute value (induced from $|\cdot|_p$ on $\mathbb{Q}_p$) satisfies the ultrametric inequality. This has some nice consequences. For example, while convergence of series is very subtle over the real numbers, in the p-adics it is very simple:

$$\sum_{n \geq 0} a_n$$

converges if and only if $\lim_n a_n = 0$. Indeed, we know $\mathbb{Q}_p$ is complete so it suffices to show that the sequence of partial sums $b_n = \sum_{k=0}^{n} a_k$ is Cauchy; and (for say $m \geq n$ without loss

of generality)

$$|b_m - b_n| = \left| \sum_{k=n+1}^{m} a_k \right| \leq \max_{n<k\leq m} |a_k|,$$

so the conditions for Cauchy sequences are equivalent to those for a_n to converge to 0.

Another nice property is that the subset of $\mathbb{Q}_p$ consisting of x with $|x|_p \leq 1$, which in $\mathbb{R}$ would be the interval $[-1, 1]$ which has no especially nice algebraic properties, in $\mathbb{Q}_p$ has a ring structure! Indeed, for $x \in \mathbb{Q}$ embedded into $\mathbb{Q}_p$, $|x|_p \leq 1$ if and only if, writing $x = a/b$ in lowest terms, we have $v_p(b) = 0$, i.e. x has no powers of p in its denominator. Since the p-adic absolute value only sees powers of p, these are the "p-adic integers in $\mathbb{Q}$," and indeed contain all the usual integers; more formally this is the localization $\mathbb{Z}_{(p)}$. Indeed we could take the completion of $\mathbb{Z}$ with respect to the p-adic absolute value just as well as $\mathbb{Q}$, and would get a ring $\mathbb{Z}_p \supset \mathbb{Z}$ with absolute value extending the p-adic one on $\mathbb{Z}$, which has $|x|_p \leq 1$ for every $x \in \mathbb{Z}_p$. The inclusion $\mathbb{Z} \subset \mathbb{Q}$ induces an inclusion $\mathbb{Z}_p \subset \mathbb{Q}_p$, and the completeness of both means that in fact $\mathbb{Z}_p$ is exactly the subring of $\mathbb{Q}_p$ with $|x|_p \leq 1$.

We can give an independent description of $\mathbb{Z}_p$, which will also shed more light on $\mathbb{Q}_p$. First though let's look at some properties of $\mathbb{Z}_p$. Topologically, as we've seen, we can think of it as the ball of radius 1 inside of $\mathbb{Q}_p$, centered at 0. Its ideals are pleasantly simple: the element $p \in \mathbb{Z}_p$ (since $|p|_p = p^{-1} \leq 1$) generates an ideal (p), as do all its powers (p^n), and as always there's the zero ideal (0) and the unit ideal $(1) = \mathbb{Z}_p$; and that's it. In particular this means that $(p) = p\mathbb{Z}_p$, equivalently the ball of radius p^{-1}, is the unique maximal ideal of $\mathbb{Z}_p$, since it contains (0) as well as each (p^n), and so $\mathbb{Z}_p/p\mathbb{Z}_p$ is a field; one can show that it is actually isomorphic to the finite field $\mathbb{F}_p$.

We can look at the units of $\mathbb{Z}_p$, which must have absolute value 1; in fact the units $\mathbb{Z}_p^\times$ are exactly the elements of $\mathbb{Q}_p$ with absolute value 1. Unlike $\mathbb{R}$, where there are only two units of absolute value 1, namely ± 1, here $\mathbb{Z}_p^\times$ is infinite.

Finally, we want to give a different, more algebraic description of $\mathbb{Z}_p$. We have a quotient map $\mathbb{Z}_p \to \mathbb{Z}_p/p\mathbb{Z}_p \simeq \mathbb{Z}/p\mathbb{Z}$. We can choose lifts of elements of $\mathbb{Z}/p\mathbb{Z}$ to $\mathbb{Z}$, say by restricting to $\{0, 1, \ldots, p-1\}$; for $x \in \mathbb{Z}_p$, write x_0 for such a lift of the image of x modulo p. Via the embedding $\mathbb{Z} \subset \mathbb{Z}_p$, we can view x_0 as an element of $\mathbb{Z}_p$, such that $x - x_0 \in p\mathbb{Z}_p$.

Next, we could quotient by p^2: this gives a map $p\mathbb{Z}_p \to p\mathbb{Z}_p/p^2\mathbb{Z}_p \simeq p \cdot (\mathbb{Z}/p^2\mathbb{Z})$. Similarly choosing a lift px_1 of the image of $x - x_1$ in $p(\mathbb{Z}/p^2\mathbb{Z})$, we get $x - x_0 - px_1 \in p^2\mathbb{Z}_p$. Repeating this gives us a sequence $\{x_i\}$ of integers, with x_i only really defined modulo p^{i+1}, such that

$$x = \sum_{n\geq 0} x_n p^n,$$

i.e. we can think of p-adic integers as power series "in p." More formally, the algebraic statement is

$$\mathbb{Z}_p \simeq \varprojlim_n \mathbb{Z}/p^n\mathbb{Z},$$

so a p-adic integer is a compatible system of elements of $\mathbb{Z}/p^n\mathbb{Z}$ for every n. Just like for usual power series, the units are the elements for which x_0 is invertible as an element of $\mathbb{F}_p$, i.e. $x_0 \neq 0$. More generally, x is in (p^n) if $x_i = 0$ for $i < n$. By dividing by the leading power of p, we can write every element of $\mathbb{Z}_p$ as up^n for some $n \geq 0$, where $u \in \mathbb{Z}_p^\times$.

An example of such a series expansion which is not obvious is -1, viewed as a p-adic number. Modulo p, this is $p-1$, so $x_0 = p-1$; $-1-(p-1) = -p$, so $x_1 \equiv -1 \equiv p-1$ as well, and so on. This gives

$$-1 = (p-1) + (p-1)p + (p-1)p^2 + \cdots .$$

Indeed, formally summing this series via the geometric series expansion gives $(p-1)\cdot\frac{1}{1-p} = -1$, and one can check that over the p-adics, unlike over the reals, this series actually converges: $|(p-1)p^n|_p = |p-1|_p \cdot |p^n|_p = 1 \cdot p^{-n} = p^{-n}$, which tends to 0 as $n \to \infty$ so the series converges as above.

This in fact lets us describe $\mathbb{Q}_p$ as well: $\mathbb{Q}_p \simeq \mathbb{Z}_p[1/p]$. We can think of p-adic numbers as Laurent series in p, i.e. we allow finitely many terms with negative powers, or as everything of the form up^n for $u \in \mathbb{Z}_p^\times$ and *any* integer n, rather than only nonnegative integers.

Again, there are analogues over any finite extension of $\mathbb{Q}$ by completing at the place corresponding to any prime ideal $\mathfrak{p}$. These are finite extensions of $\mathbb{Q}_p$ if $\mathfrak{p}$ lies over p, and have similar properties and descriptions. There are also analogues in the function field world: there, every completion is of the form $\mathbb{F}_q((t))$, i.e. Laurent series over a finite field. These, together with $\mathbb{R}$ and $\mathbb{C}$, are the local fields; they are given by completion of global fields at their places.

4 Algebraic extensions and closures

We've referred a few times to finite extensions of $\mathbb{Q}$ or related fields; some examples we've seen before include $\mathbb{Q}(i)$, $\mathbb{Q}(\sqrt{-3})$, and the cyclotomic fields $\mathbb{Q}(\zeta_p)$. More generally, we can study *algebraic* extensions. For a field extension L/K, an element $x \in L$ is algebraic over K if there exists some polynomial f with coefficients in K such that $f(x) = 0$ (e.g. for i over $\mathbb{Q}$, we have $f(x) = x^2 + 1$); we say that L/K is algebraic if every element of L is algebraic over K. An example of a non-algebraic extension is given by $\mathbb{Q}(t)/\mathbb{Q}$, since the element t is transcendental (i.e. not algebraic) over $\mathbb{Q}$; another is given by $\mathbb{C}/\mathbb{Q}$, since $\mathbb{C}$ includes transcendental elements such as π.

All finite extensions are algebraic; but the converse is not necessarily true. An example is given by the algebraic numbers $\overline{\mathbb{Q}}/\mathbb{Q}$: this contains e.g. roots of the irreducible polynomial $x^n - 2$ for every $n \geq 1$, and so must have infinite degree, but by definition every element is algebraic over $\mathbb{Q}$.

This is a special case of a more general idea: the algebraic closure $\overline{K}$ of a field K, which can loosely be thought of as the collection of all elements which are algebraic over K. More precisely, a field is algebraically closed if it has no nontrivial algebraic extensions (e.g. $\mathbb{C}$); an algebraic closure of K is an algebraic extension which is algebraically closed. It is not at all trivial to show that every field has an algebraic closure, but it is true (assuming the axiom of choice); there is also the question of to what extent the algebraic closure is unique. It turns out that any two algebraic closures of a field are isomorphic, but not necessarily in a canonical way, so sometimes people prefer to say that we choose an algebraic closure, though we will not generally concern ourselves too much with such subtleties.

It is interesting to ask how algebraic closures interact with completions. The answer is that the algebraic closure of a complete field is sometimes, but not always, complete. In

particular, the algebraic closure of $\mathbb{R}$ is $\mathbb{C}$, which is also complete. However, the algebraic closure $\overline{\mathbb{Q}}_p$ of $\mathbb{Q}_p$ is no longer complete. One can then take its completion $\widehat{\overline{\mathbb{Q}}}_p$, and ask if this is still algebraically closed; and the answer is yes, so we have now arrived at a (quite large) complete and algebraically closed field, analogous to $\mathbb{C}$.

At this point, before moving on to discuss Galois representations, we can make an important and interesting remark about Fermat's last theorem: it is an essentially global statement. That is: we want to show that the equation $x^p + y^p = z^p$ has no solutions in nonzero integers, or equivalently nonzero rational numbers (a global field). One way we could hope to show this is to say: well, we've now found infinitely many embeddings $\mathbb{Q} \subset \mathbb{Q}_\ell$ over various primes ℓ (since we've already used p), as well as $\mathbb{Q} \subset \mathbb{R}$; so if we could find some place v, finite or infinite, such that the Fermat equation has no nontrivial solution in $\mathbb{Q}_v$, then it would follow that it has no nontrivial solutions in $\mathbb{Q}$. This would be a *local* obstruction to the existence of solutions.

However, there are no such local obstructions: it is easy to see (via drawing some graphs) that there do exist solutions in $\mathbb{R}$. We claim that in fact for every prime ℓ, there exist solutions in $\mathbb{Q}_\ell$ as well, so whatever it is that (we hope) causes the Fermat equation to not have rational solutions must be a fundamentally global phenomenon.

We start with the case $\ell = p$. Taking $x = 1$ and $z = p$, we can find a solution in $\overline{\mathbb{Q}}_p$ by taking $y = -(1 + p^p)^{1/p}$. The function $(1 + X)^{1/p}$ has a binomial series expansion

$$(1 + X)^{1/p} = 1 - \sum_{n \geq 1} \frac{1}{n!}(p - 1)(2p - 1) \cdots ((n - 1)p - 1)(X/p)^n$$

and in particular has rational coefficients; for $X = p^p$ it can be shown to converge, giving an element of $\mathbb{Q}_p$.

For $\ell \neq p$, we similarly take $x = 1$ and $z = \ell$; in this case the binomial series for $(1 + X)^{1/\ell}$, we haven't introduced any new factors of p and so the series converges more straightforwardly to an element of $\mathbb{Q}_\ell$.

5 Introduction to representation theory

We will want to study number fields via studying their automorphisms, i.e. their Galois groups. Before we worry about studying these, we first worry about a much broader question: how should we study groups in general? The principle we will use is that groups should be understood via their actions: all the information of a group can be recovered from the data of how it acts on various different objects, but these are often easier to study.

The first claim, that we can recover all the information about groups from their actions, is a little bit subtle to make precise; a high-level keyword here is "the Tannakian formalism." In a simpler sense though it's almost trivial: every group acts on itself by translation, so if we understand this action we understand the group.

To justify the second part, we need to say more about what we mean. First of all, what actions should we take? The simplest group actions are on sets, but this is not much better than just studying groups; we'd like something with more and simpler algebraic structure. The next simplest thing we might then think to study is group actions on other groups, or

perhaps on abelian groups for simplicity; but there are various phenomena such as torsion in even abelian groups that makes this less than ideal. We could then think to ask about actions on torsion-free abelian groups, or perhaps free abelian groups, i.e. free $\mathbb{Z}$-modules (don't worry if this terminology is unfamiliar); but it's convenient to be able to scale by elements of a field rather than just $\mathbb{Z}$ so we can have inverses, suggesting studying group actions on vector spaces over some field k. Although one can take any field for k, it's often convenient to have k algebraically closed, and for definiteness we often specialize to the case $k = \mathbb{C}$.

Now, what do we mean by a group action on a complex vector space? You may recall the definition of a group action from algebra: for our vector space V and each g in our group G, we should get a linear map $V \to V$ induced by g satisfying various properties. A more concise way of saying this is that we can all agree that linear transformations $V \to V$ act on V, and since we're interested in groups we can restrict to the invertible ones, $\mathrm{GL}(V)$; and so an action of G on V is the same thing as a group homomorphism $\rho : G \to \mathrm{GL}(V)$, i.e. $\rho(g_1 g_2) = \rho(g_1)\rho(g_2)$. This is called a representation of G (on V).

For example, consider the unique group of order 3, which is the cyclic group $C_3 = \{1, \sigma, \sigma^2\}$ with generator σ. Let $V \simeq \mathbb{C}^2$ be a 2-dimensional vector space, and suppose that σ acts on V by the matrix

$$\rho(\sigma) = \begin{pmatrix} 1 & -1 \\ 3 & -2 \end{pmatrix}.$$

Since ρ is a homomorphism and G is cyclic, this determines the whole representation: we must have $\rho(1) = I_2$ and

$$\rho(\sigma^2) = \rho(\sigma)^2 = \begin{pmatrix} -2 & 1 \\ -3 & 1 \end{pmatrix}.$$

A simpler representation is on the one-dimensional representation via $\chi(\sigma) = e^{2\pi i/3}$, so $\chi(\sigma^2) = \chi(\sigma)^2 = e^{4\pi i/3}$, and again $\chi(\sigma) = 1$, acting by multiplication on $\mathbb{C}$. Simpler yet is the trivial representation $\chi_0(g) = 1$ for all g; every group has this representation.

The one-dimensional representations are always the simplest, and it's easier to think of them as just functions (in fact group homomorphisms) $G \to \mathbb{C}^\times$ instead of the linear algebraic data we need in general. In fact, we can get a function $G \to \mathbb{C}$ from any (finite-dimensional) representation, simply by taking its trace: we call this the character of a representation ρ, given by $\chi(g) = \mathrm{Tr}(\rho(g))$. Naively, it seems like this loses a lot of information, but in fact this turns out not to be the case: two representations with the same character are isomorphic.

Some of the information about a representation can be read off easily from its character: for example, if it has dimension d, then $\chi(1) = d$, since $\rho(1) = I_d$ which has trace d. Other information is less obvious. Since the trace is invariant under base change, although the matrices of the representation will be different in different bases the character will be the same, so it can be viewed as retaining the base-invariant information. More precisely, it is conjugation-invariant:

$$\chi(g^{-1}hg) = \mathrm{Tr}(\rho(g)^{-1}\rho(h)\rho(g)) = \mathrm{Tr}(\rho(g)\rho(g)^{-1}\rho(h)) = \mathrm{Tr}(\rho(h)) = \chi(h)$$

for all $g, h \in G$.

If we are interested in classifying the representations of a group, we immediately run into a problem that, given some representations, we can generate more by taking their direct

sum, which gives infinitely many representations in a rather trivial way (this corresponds to adding various characters together). Thus we say that a representation is irreducible if it cannot be written as the direct sum of two other representations; so we're interested in classifying irreducible representations of a group. (More precisely we say that a representation is irreducible if it has no proper nontrivial subspaces preserved by G; it is a nontrivial theorem that every representation of a finite group over $\mathbb{C}$ decomposes as a direct sum of irreducible representations.)

If G is an abelian group, its irreducible representations are relatively simple: they are all one-dimensional. (In practice enumerating them may still be challenging for some groups, but this is simple compared to the general case.) In fact the converse is also true, at least for finite groups.

Notably, this means that the example of C_3 acting on $\mathbb{C}^2$ above is reducible, although it doesn't naively look like a direct sum. However, the matrices in question are diagonalizable, and changing basis one finds that ρ decomposes as the direct sum of the one-dimensional representation χ (sending σ to multiplication by $e^{2\pi i/3}$ on $\mathbb{C}$) and its conjugate $\overline{\chi}$ (sending σ to $e^{4\pi i/3}$).

There is quite a lot more to say about group representations, even for finite groups over $\mathbb{C}$ (for example we have not mentioned character orthogonality relations), but this should give us the basics we'll need. To see why this sort of thing might be useful, let's look into some Galois theory.

6 Introduction to Galois theory

For a field K, its automorphisms are the ring homomorphisms $f : K \to K$, i.e. maps which satisfy $f(x+y) = f(x)+f(y)$ and $f(xy) = f(x)f(y)$ for all $x, y \in K$. If K contains $\mathbb{Q}$, it is not too hard to verify that every automorphism of K must fix $\mathbb{Q}$, i.e. for $x \in \mathbb{Q}$ we have $f(x) = x$ (it can be checked for integers using additivity and the property $f(1) = 1$ and then extended to the rationals by multiplicativity). More generally, we say that for an algebraic extension L/K, an automorphism f of L is a K-automorphism if it fixes K, i.e. for every $x \in K$ we have $f(x) = x$. An example is the finite extension $\mathbb{C}/\mathbb{R}$, with the automorphism given by complex conjugation: for every real number x we have $\overline{x} = x$, so complex conjugation is an $\mathbb{R}$-automorphism of $\mathbb{C}$. For any extension L/K the K-automorphisms of L form a group under composition, which we call the Galois group $\mathrm{Gal}(L/K)$.

There exist various subfields between L and K, i.e. subfields of L containing K; there also exist various subgroups of the Galois group $\mathrm{Gal}(L/K)$. The goal of Galois theory is to provide a dictionary between these.

Suppose M is an intermediate extension, i.e. we have a tower of extensions $L/M/K$. Any M-automorphism of L is also a K-automorphism of L, so we get a subgroup $\mathrm{Gal}(L/M) \subset \mathrm{Gal}(L/K)$.

Conversely, say we have a subgroup $H \subset \mathrm{Gal}(L/K)$. This acts on L and fixes the subfield K; since it is not necessarily the whole Galois group, it may in fact fix other elements not in K as well (e.g. the trivial subgroup $\{1\} \subset \mathrm{Gal}(L/K)$ fixes all of L!). The subset of L consisting of elements fixed by every element of H is a subfield containing K, i.e. we get an intermediate extension given by the H-fixed points, sometimes written as L^H. If the subfield

fixed by the whole group $\mathrm{Gal}(L/K)$ is the base field K, then we say that L/K is a Galois extension. In this case these two constructions give an order-reversing bijection

$$\{\text{intermediate extensions of } L/K\} \leftrightarrow \{\text{subgroups of } \mathrm{Gal}(L/K)\}.$$

This maps Galois extensions of K to normal subgroups of $\mathrm{Gal}(L/K)$, i.e. L^H is Galois over K if and only if $H = \mathrm{Gal}(L/L^H)$ is a normal subgroup of $\mathrm{Gal}(L/K)$. In this case we can study the quotient group $\mathrm{Gal}(L/K)/\mathrm{Gal}(L/L^H)$, which is isomorphic to $\mathrm{Gal}(L^H/K)$. It also maps the degree of the extension to the size of the Galois group, i.e. $[L:K] = |\mathrm{Gal}(L/K)|$. In fact, for L/K a finite extension it is Galois if and only if this equality holds.

For example, the algebraic closure $\overline{K}/K$ is always Galois for K of characteristic 0 (i.e. containing $\mathbb{Q}$). On the other hand, the real algebraic numbers $\mathbb{R} \cap \overline{\mathbb{Q}}$ are not Galois over $\mathbb{Q}$ (in fact, one can show that its automorphism group is trivial!).

For the cyclotomic fields $\mathbb{Q}(\zeta_p)$ that we saw a few weeks ago, their Galois group over $\mathbb{Q}$ is given by $(\mathbb{Z}/p\mathbb{Z})^\times$. Indeed, they are generated by the pth roots of unity, which as a group are isomorphic to $\mathbb{Z}/p\mathbb{Z}$; and the automorphisms of this group are just multiplication by invertible elements. We do indeed have $|(\mathbb{Z}/p\mathbb{Z})^\times| = |\mathbb{F}_p^\times| = p-1 = [\mathbb{Q}(\zeta_p):\mathbb{Q}]$ (recall that the defining irreducible polynomial is $1 + x + \cdots + x^{p-1}$, of degree $p-1$, and that $\mathbb{Q}(\zeta_p)$ has $\mathbb{Q}$-basis given by $1, \zeta_p, \zeta_p^2, \ldots, \zeta_p^{p-2}$ of order $p-1$), so these are Galois extensions of $\mathbb{Q}$.

In many of the cases we've seen so far, the Galois groups are abelian. In general, though, this is very far from the case. In particular if we study the largest and most interesting possible Galois group over $\mathbb{Q}$, namely the *absolute Galois group* $\mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$, this is extremely nonabelian. By Galois theory, its subgroups correspond to algebraic extensions of $\mathbb{Q}$; in principle, one has to be fairly careful here because $\mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$ is infinite (and in fact a topological group), and so one should put some topological conditions on which subgroups we allow. However, in practice we are interested in finite extensions, and use the single group $\mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$ as a gadget to handle them all at once; but we will generally focus on its finite quotients, corresponding to finite Galois extensions of $\mathbb{Q}$.

Indeed, since by our above principle we should study the Galois group via its representations, we make the following definition: a representation of a topological group is continuous if the map $G \to \mathrm{GL}(V)$ is continuous (using the natural topology on $\mathrm{GL}(V)$), which in our case of $G = \mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$ (or similarly the absolute Galois group of other field) is equivalent to saying that the kernel of the representation is a finite index subgroup, so the representation factors through the quotient, which is the Galois group of a finite extension of $\mathbb{Q}$.

It is also worth noting that the inclusions $\mathbb{Q} \subset \mathbb{Q}_v$ and $\overline{\mathbb{Q}} \subset \overline{\mathbb{Q}}_v$ for each place v give maps $\mathrm{Gal}(\overline{\mathbb{Q}}_v/\mathbb{Q}_v) \to \mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$, by restricting the automorphism of $\overline{\mathbb{Q}}_v$ to $\overline{\mathbb{Q}}$ and observing that if it fixes $\mathbb{Q}_v$ it also fixes $\mathbb{Q}$. (Really this depends on a choice of algebraic closure and embedding of algebraic closures; a different choice will result in a map differing by conjugation.) The local Galois groups are also generally complicated (except for that of $\mathbb{Q}_\infty = \mathbb{R}$, which is $\mathrm{Gal}(\mathbb{C}/\mathbb{R}) \simeq C_2 = \{1, c\}$ where c is complex conjugation), but less so than the global one.

7 Abel–Ruffini theorem

A common application of Galois theory is the following result: while for polynomial equations in one variable of degrees 2, 3, and 4 over $\mathbb{Q}$ there exist (increasingly complicated) formulas

for finding their solutions in terms of addition, subtraction, multiplication, division, and radicals (i.e. taking nth roots for some n), there exists no such formula for degree 5 or higher, no matter how complicated.

This is a corollary of the following more precise result: let $f(x)$ be a polynomial with rational coefficients, whose zeros generate the finite Galois extension $K/\mathbb{Q}$. Then the equation $f(x) = 0$ is solvable by radicals if and only $\mathrm{Gal}(K/\mathbb{Q})$ is *solvable* (a group-theoretic condition which we will return to shortly). Analogues of this result over fields other than $\mathbb{Q}$ exist as well, though one has to be careful in positive characteristic.

To see this, we think about the operation of solving an equation by radicals as follows: starting with the coefficients of the equation, we make some computation (either addition, subtraction, multiplication, division, or taking a radical). We get some new number; either we stop there and produce that number or we repeat the process, now using the coefficients and this new number as input, and so on for finitely many steps until we get the result.

By assumption, we are starting with values in $\mathbb{Q}$, and we end up with a solution to $f(x) = 0$, which is in K (and in fact generates K). After each computation step, we can look at the field generated by all of the inputs: after zero steps, this is just the field generated by the coefficients of the polynomial, which are rational, so $K_0 = \mathbb{Q}$. Addition, subtraction, multiplication, and division don't change the field, so if the computation is one of these then $K_{i+1} = K_i$; taking a radical does change it, with $K_{i+1} = K_i(x_i^{1/n_i})$ for some $x_i \in K_i$ and integer $n_i \geq 2$. (To make sure everything is Galois, we sometimes need to add n_ith roots of unity to K_{i+1} as well; this is harmless since the eventual overall extension $K/\mathbb{Q}$ is Galois.) So we get a finite series of extensions

$$\mathbb{Q} = K_0 \subset K_1 \subset \cdots \subset K_r = K,$$

with each K_{i+1}/K_i Galois and $\mathrm{Gal}(K_{i+1}/K_i)$ either trivial or cyclic, and in particular abelian. Via Galois theory, this corresponds to a descending series of subgroups

$$\mathrm{Gal}(K/\mathbb{Q}) = \mathrm{Gal}(K/K_0) \supset \mathrm{Gal}(K/K_1) \supset \cdots \supset \mathrm{Gal}(K/K_r) = \mathrm{Gal}(K/K) = \{1\}$$

with $\mathrm{Gal}(K/K_{i+1})$ normal in $\mathrm{Gal}(K/K_i)$ and each $\mathrm{Gal}(K/K_i)/\mathrm{Gal}(K/K_{i+1}) \simeq \mathrm{Gal}(K_{i+1}/K_i)$ abelian.

Now we define a solvable group: a group G is solvable if there exists a finite series of subgroups

$$G = G_0 \supset G_1 \supset \cdots \supset G_r = \{1\}$$

such that G_{i+1} is normal in G_i and G_i/G_{i+1} is abelian for every i. These are the groups that can be "built up" from abelian groups in a certain sense.

Now the above translates into the following: if $f(x) = 0$ is solvable by radicals, then $\mathrm{Gal}(K/\mathbb{Q})$ is solvable, with $G_i = \mathrm{Gal}(K/K_i)$. With a little more care one can prove the converse similarly, though it isn't actually needed for the Abel–Ruffini theorem.

For low-degree equations, this isn't an issue since all sufficiently small groups are solvable. Indeed, for $f(x)$ of degree d we have $|\mathrm{Gal}(K/\mathbb{Q})| \leq d!$ since the Galois group acts transitively on the d roots and so embeds into the symmetric group S_d, and for $d \leq 4$ one can check that all the groups up to order 24 are solvable, so all such equations of degree at most 4 are solvable by radicals. The smallest non-solvable group turns out to be the alternating

group A_5, of order 60; this is an index 2 subgroup of S_5, and does occur as the Galois group of degree 5 polynomials, and non-solvable groups become more and more common as the size increases. Therefore for any $\deg f \geq 5$ we cannot solve every polynomial equation by radicals, and so there can exist no such formula.