ℓ-adic Galois Representations

Kyrie Johnson

December 23, 2020

Abstract

In this paper, we introduce objects which are central to modern number theory: ℓ-adic Galois representations. These special representations correspond with both highly-symmetric meromorphic functions on the upper half-plane of $\mathbb{C}$ — so-called "automorphic forms" — and with the rational zero-sets of polynomials — so-called "algebraic varieties" over $\mathbb{Q}$. In section 1, we motivate our study by describing how this correspondence proved Fermat's Last Theorem, one of the great triumphs of twentieth century mathematics. In section 2, we lay down a foundation of essential number theory. In section 3, we precisely define ℓ-adic Galois representations and study some examples. And finally in section 4, we return to the big picture and discuss the vast frontier of current research in this area.

1 Fermat's Last Theorem as a motivating example

One of the most famous results of the twentieth century, Fermat's Last Theorem (FLT) states that for $n \geq 3$ the Fermat equation $x^n + y^n = z^n$ has no integer solutions with $xyz \neq 0$. Pierre de Fermat first conjectured FLT in 1637, but only proved the case of $n = 4$. Note that a solution

$$x^{pq} + y^{pq} = z^{pq}$$

to the Fermat equation with exponent pq yields a solution

$$(x^p)^q + (y^p)^q = (z^p)^q$$

to the Fermat equation with exponent q. Thus, that Fermat proved FLT for the exponent $n = 4$ in fact proved FLT for all exponents divisible by four, namely every even number larger than two. So one need only consider the cases of odd prime exponents ℓ. Early developments tackled the exponents $\ell = 3$, $\ell = 5$, and $\ell = 7$ in work by Euler, Dirichlet, Legendre, and Lamé.

Although the nineteenth century witnessed the development of some more general cases of FLT — in the work of Sophie Germain and later Ernst Kummer — the modern proof of Fermat's Last Theorem did not emerge until the twentieth century. The modern story begins with a 1965 conjecture of Goro Shimura and Yutaka Taniyama, which pertains to two objects seemingly unrelated to FLT: elliptic curves and modular forms. An "elliptic curve" over $\mathbb{Q}$ is

a "nice" curve given by an equation of the form $y^2 + a_1xy + a_3y = x^3 + a_2x^2 + a_4x + a_6$, with $a_i \in \mathbb{Q}$. A "modular form" is a holomorphic automorphic form, and so is a highly-symmetric holomorphic function from the upper half-plane of $\mathbb{C}$ to $\mathbb{C}$. At first glance, elliptic curves and modular forms live in entirely different worlds. Shimura and Taniyama's incredible insight was that through ℓ-adic Galois representations, there exists a precise way of associating to each elliptic curve over $\mathbb{Q}$ a family of modular forms. Even more remarkable, this bridge forms a key ingredient in the proof of FLT.

Throughout the twentieth century, the proof of FLT developed through a series of reductions which ultimately contradict the existence of an integer solution to $a^\ell + b^\ell = c^\ell$ with $abc \neq 0$ and $\ell > 7$:

1. In 1975, Yves Hellegouarch took this hypothetical solution $a^\ell + b^\ell = c^\ell$ and associated to it a curve
$$E_\ell : y^2 = x(x + a^\ell)(x - b^\ell).$$
Ten years later, Gerhard Frey determined that for $\ell > 7$, the curves E_ℓ are "semi-stable" elliptic curves with some interesting properties.

2. Because the E_ℓ define elliptic curves, the Shimura-Taniyama Conjecture (STC) applies and associates each E_ℓ with a family of modular forms via their ℓ-adic Galois representations; these modular forms have a "weight" $k = 2$ and a "level" N.

3. In 1990, Ken Ribet proved that if the family of forms which correspond to E_ℓ have weight $k = 2$ and any level N, then that family also contains forms of weight $k = 2$ and level $N = 2$. But it turns out that there are no modular forms of weight $k = 2$ and level $N = 2$, so the elliptic curves E_ℓ cannot exist and thus the solution $a^\ell + b^\ell = c^\ell$ cannot exist!

So by 1990, every link in the chain was complete except for a proof of STC. Andrew Wiles, assisted by Richard Taylor, famously proved STC (for semi-stable elliptic curves) in 1995, thereby completing the proof of FLT.

Because STC uses ℓ-adic Galois representations to associate modular forms to elliptic curves, the theory of Galois representations plays a central role; this motivates our study here. Additional reading — which we referenced in writing this section — includes [1], which offers an excellent introduction to modular forms and STC, and [2], which contains a more-detailed look at how the proof of FLT proceeds. In Section 3.3, we'll return to one aspect of STC by investigating the ℓ-adic Galois representations associated to elliptic curves; only in passing will we describe how they correspond to the Galois representations of modular forms.

2 Number theory preliminaries

As their name suggests, Galois representations represent Galois groups, so we begin with some number theoretic Galois theory. Recall that a field extension $L/\mathbb{Q}$ is Galois if every irreducible polynomial in $\mathbb{Q}[X]$ either has no roots in L or splits completely in L (that is, the extension is normal).

Now, let $\overline{\mathbb{Q}} \subset \mathbb{C}$ denote the sub-field of $\mathbb{C}$ of numbers algebraic over $\mathbb{Q}$ and fix $s \in \overline{\mathbb{Q}}$. Then, by definition, there exists a monic polynomial $p(X) \in \mathbb{Q}[X]$ with $p(s) = 0$. Every other root of p necessarily also lies in $\overline{\mathbb{Q}}$ so that the minimal polynomial of s over $\mathbb{Q}$ — which divides $p(X)$ — splits completely over $\overline{\mathbb{Q}}$. It follows that the extension $\overline{\mathbb{Q}}/\mathbb{Q}$ is normal and so defines a Galois extension of $\mathbb{Q}$.

Definition 2.1. Call the Galois group $G_{\overline{\mathbb{Q}}} := \mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q}) := \mathrm{Aut}(\overline{\mathbb{Q}}/\mathbb{Q})$ the **absolute Galois group** (of $\mathbb{Q}$). Note that the notation $G_{\overline{\mathbb{Q}}}$ is unconventional ($G_{\mathbb{Q}}$ is standard) but we adopt the former here for consistency with later notation. Note also that this notation suppresses the base field $\mathbb{Q}$; we are justified in doing this as the base field will always be $\mathbb{Q}$ here.

Understanding the structure of both $\overline{\mathbb{Q}}$ and $G_{\overline{\mathbb{Q}}}$ remains an important — and often intractable — goal of number theory. A natural method of simplifying their study is to consider finite sub-extensions.

Definition 2.2. A **(Galois) algebraic number field**, or simply a **(Galois) number field**, is a finite (Galois) extension $L/\mathbb{Q}$ such that $L \subset \overline{\mathbb{Q}}$.

We will use Galois number fields throughout the rest of the paper, and emphasise the importance of their finiteness as extensions of $\mathbb{Q}$.

2.1 Inverse limits

This section largely follows the exposition of [3], with additional references including [6] and [4]. Because we have an equality of sets

$$\overline{\mathbb{Q}} = \bigcup_{\substack{\text{Galois} \\ \text{number fields}}} L,$$

we expect that we can realise the absolute Galois group $G_{\overline{\mathbb{Q}}}$ as a sort of limit of finite Galois groups. To that end, consider a tower of field extensions $\overline{\mathbb{Q}}/L/\mathbb{Q}$ with L a Galois number field. Given an automorphism $\sigma \in G_{\overline{\mathbb{Q}}}$, define an automorphism $\sigma_L \in \mathrm{Gal}(L/\mathbb{Q})$ by restriction: $\sigma_L := \sigma|_L$. Conversely, given all such restrictions σ_L, we can reconstruct σ: for each $s \in \overline{\mathbb{Q}}$, set $\sigma(s) := \sigma_L(s)$ for some Galois number field L containing s. In this way, we correspond $\sigma \in G_{\overline{\mathbb{Q}}}$ with a collection $\{\sigma_L\}$ of automorphisms $\sigma_L \in \mathrm{Gal}(L/\mathbb{Q})$ which satisfy, for all Galois number fields L and M,

- σ_L and σ_M agree on $L \cap M$, and
- σ restricts to σ_L on L.

The first bullet asks that automorphisms in distinct Galois groups knit together in a natural way. The second bullet connects elements of $G_{\overline{\mathbb{Q}}}$ to elements of finite sub-extensions. We shall see that this motivating example in fact fits into a much more general construction.

Definition 2.3. Let $\mathcal{I}$ denote a set together with a directed partial ordering $\leq$ on its elements. Let $(G_i)_{i \in \mathcal{I}}$ be a collection of groups with maps $r_{j,i} : G_j \to G_i$ such that

- $r_{i,i} = \mathrm{id}_{G_i}$ for all i, and
- $r_{j,i} \circ r_{k,j} = r_{k,i}$ whenever $i \leq j \leq k$.

The pair $(G_i)_{i\in\mathcal{I}}$ and $(r_{j,i})_{j\geq i}$ make up an **inverse system** of groups and **bonding morphisms** over $\mathcal{I}$.

We denote the collections of groups and bonding morphisms with parentheses — rather than braces — to emphasise that indexing over $\mathcal{I}$ induces a directed partial ordering on the collections. We will want to keep track of the induced ordering to precisely work with the structure we next define.

Definition 2.4. The **inverse limit** of an inverse system $(G_i)_{i\in\mathcal{I}}$ and $(r_{j,i})_{i,j\in\mathcal{I}}$ is defined as

$$\varprojlim_{i\in\mathcal{I}} G_i := \left\{ (a_i) \in \prod_{i\in\mathcal{I}} G_i : r_{j,i}(a_j) = a_i \text{ for all } j \geq i \right\}$$

We will refer to the conditions $r_{j,i}(a_j) = a_i$ as the **compatibility conditions** of the limit. Note that $\varprojlim_{i\in\mathcal{I}} G_i$ inherits a group structure from $\prod_{i\in\mathcal{I}} G_i$ by component-wise multiplication.

Thus, the inverse limit of $(G_i)_{i\in\mathcal{I}}$ consists of those "sequences" in $\prod_{i\in\mathcal{I}} G_i$ whose elements are "compatible" in the sense that they are related by the maps $r_{j,i}$. Note also that we required that the maps $r_{j,i}$ be group homomorphisms so that the inverse limit indeed defines a subgroup of $\prod_{i\in\mathcal{I}} G_i$.

Let's digest all of this abstraction via the example $G_{\overline{\mathbb{Q}}}$. Set

$$\mathcal{I} := \{L/\mathbb{Q} : L \text{ is a Galois number field}\}$$

with partial ordering defined by $L \leq M$ whenever $L \subseteq M$ and consider the groups

$$G_L := \mathrm{Gal}(L/\mathbb{Q}).$$

For $M \geq L$ we define a bonding morphism $r_{M,L}$ by restriction,

$$\begin{aligned} r_{M,L} : G_M &\to G_L \\ \sigma &\mapsto \sigma|_L, \end{aligned}$$

so that whenever $M \geq L \geq K$, we have $r_{L,K} \circ r_{M,L} = r_{M,K}$. Thus, the groups $(G_L)_{L\in\mathcal{I}}$ and bonding morphisms $(r_{M,L})_{M\geq L}$ together constitute an — absolutely massive — inverse system. The following diagram shows an excerpt of this system, with G_L denoting $\mathrm{Gal}(L/\mathbb{Q})$ as before.

$$
\begin{tikzcd}
\cdots & \cdots & \cdots & \cdots & \\
 & G_{\mathbb{Q}(\sqrt{2},\sqrt{3})} \arrow[d] \arrow[dr] & G_{\mathbb{Q}(\sqrt{2},\sqrt{5})} \arrow[dl] \arrow[dr] & G_{\mathbb{Q}(e^{2\pi i/5})} \arrow[d] & \\
G_{\mathbb{Q}(\cos(\frac{2\pi}{7}))} \arrow[d] \arrow[drr, bend right] & G_{\mathbb{Q}(\sqrt{2})} \arrow[dr] & G_{\mathbb{Q}(\sqrt{3})} \arrow[d] & G_{\mathbb{Q}(\sqrt{5})} \arrow[dl] & \cdots \arrow[d] \arrow[dll, bend left] \\
 & & G_{\mathbb{Q}} & &
\end{tikzcd}
$$

Our earlier discussion establishes the following result regarding the inverse limit of Galois groups of algebraic extensions over $\mathbb{Q}$:

Theorem 2.5. *Let $\mathcal{I} := \{L/\mathbb{Q} : L$ is a Galois number field$\}$ have partial ordering defined by inclusion and let $G_L := \mathrm{Gal}(L/\mathbb{Q})$ define an inverse system of groups $(G_L)_{L\in\mathcal{I}}$ with bonding morphisms given by restriction. Then $\varprojlim_L (G_L) \cong G_{\overline{\mathbb{Q}}}$ by the isomorphism*

$$
\begin{aligned}
G_{\overline{\mathbb{Q}}} &\xrightarrow{\sim} \varprojlim_{L\in\mathcal{I}} G_L \\
\sigma &\mapsto (\sigma|_L)_{L\in\mathcal{I}}.
\end{aligned}
$$

To make clear the nature of this isomorphism, consider the automorphism $\sigma \in \mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$ which swaps $\sqrt{2}$ and $-\sqrt{2}$ and fixes everything without relation to $\sqrt{2}$. Then

- $\sigma|_{\mathbb{Q}(\sqrt{2})}$ is the unique non-trivial automorphism in $G_{\mathbb{Q}(\sqrt{2})}$,
- $\sigma|_{\mathbb{Q}(\sqrt{2},\sqrt{3})}$ is the unique non-trivial automorphism of $G_{\mathbb{Q}(\sqrt{2},\sqrt{3})}$ fixing $\sqrt{3}$,
- $\sigma|_{\mathbb{Q}(e^{2\pi i/5})}$ is the identity automorphism of $G_{\mathbb{Q}(e^{2\pi i/5})}$, and
- $\sigma|_{\mathbb{Q}(e^{2\pi i/8})}$ is the automorphism of $\mathbb{Q}(e^{2\pi i/8})$ which maps $e^{2\pi i/8} = \frac{\sqrt{2}}{2} + i\frac{\sqrt{2}}{2}$ to $-e^{2\pi i/8}$.

Likewise, for any Galois number field L, we have that σ restricts to an automorphism $\sigma|_L$ of $L/\mathbb{Q}$. We have previously argued that these automorphisms satisfy the compatibility conditions — which one may check explicitly for the four examples above — and so indeed define a sequence $(\sigma|_L)_{L\in\mathcal{I}} \in \varprojlim_{L\in\mathcal{I}} G_L \subset \prod_{L\in\mathcal{I}} G_L$.

So far, we have realised $G_{\overline{\mathbb{Q}}}$ as a limit of finite Galois groups so that we have an alternative — often more convenient — description of its group structure. Along the way, we defined the inverse limit of a system of groups, which we will later use to construct a family of crucial objects. But for now we endow $G_{\overline{\mathbb{Q}}}$ with a nice topology. We do so by endowing each finite Galois group with the discrete topology and then passing along the inverse limit.

2.2 The Profinite Topology

Our primary reference for this section is [6], with [4] serving as a more down-to-earth reference of the topological structure of just $G_{\overline{\mathbb{Q}}}$ (rather than an arbitrary profinite group). A topological group G is a group equipped with a topology for which the multiplication map $G \times G \to G$ and the inverse map $G \to G$ are continuous. Lie groups form an important class of topological groups as do "profinite groups".

Definition 2.6. A **profinite group** is a topological group which arises as an inverse limit — also known as a *pro*jective limit — of finite groups.

In this section we analyse the topology of profinite groups.

Let $\mathcal{I}$ denote a directed partially-ordered set for which we have an inverse system $(G_i)_{i\in\mathcal{I}}$ of finite groups (implicitly, we also have bonding morphisms $r_{j,i}$). Endow each finite group G_i with the discrete topology, so that we may endow the product group $\prod_{i\in\mathcal{I}} G_i$ — with the group operation defined component-wise — with the product topology. Recall that product topology is the smallest topology for which the natural projection maps $\pi_j : \prod_i G_i \to G_j$ are continuous. Finally, we endow the inverse limit $G := \varprojlim_{i\in\mathcal{I}} G_i$ with the subspace topology of $\prod G_i$.

Definition 2.7. Let G_i be finite groups. Then the above-constructed topology on $G := \varprojlim_{i\in\mathcal{I}} G_i$ is the **profinite topology** on the profinite group G.

While the construction of the profinite topology includes many steps, the overall idea is quite natural. The profinite group G knits together the groups G_i, so we expect that the natural projections $G \to G_i$ are continuous. By appealing to the product topology $\prod G_i$, we guarantee such continuity. Moreover, the nature of the subspace and product topologies gives a profinite group the structure of a topological group:

Theorem 2.8. *Let $G = \varprojlim_{i\in\mathcal{I}} G_i$ define a profinite group (endowed with the profinite topology). Then the inverse map $G \to G$ and the product map $G \times G \to G$ are continuous.*

Proof. By the definition of the product topology and the fact that each G_i has the discrete topology, we have that

$$\mathcal{B} := \left\{ \prod_{i\notin S} G_i \times \prod_{i\in S} \{g_i\} : g_i \in G_i \text{ and } S \subset \mathcal{I} \text{ finite} \right\}$$

is a basis of the product topology on $\prod_{i\in\mathcal{I}} G_i$. Then the definition of the subspace topology guarantees that

$$\mathcal{U} := \{G \cap B : B \in \mathcal{B}\}$$

defines a basis of the profinite topology on G. To show that the inverse map is continuous, it suffices to show that U^{-1} is open for any $U \in \mathcal{U}$. But

$$B = \prod_{i\notin S} G_i \times \prod_{i\in S} \{g_i\} \in \mathcal{B} \implies B^{-1} = \prod_{i\notin S} G_i \times \prod_{i\in S} \{g_i^{-1}\} \in \mathcal{B}$$

so that for $U := G \cap B$, we have $U^{-1} = G \cap B^{-1} \in \mathcal{U}$. Thus, U^{-1} is open and we're done.

The proof of the continuity of the product map $G \times G \to G$ relies on the same ideas — but is slightly more involved — so we leave it as an exercise for the reader to contemplate. $\square$

A topological group G comes with not only a topology but also a distinguished identity element, so one naturally wonders about how the two relate. In general, a "neighborhood base" of the identity offers a natural relationship; in the case of our motivating example $G_{\overline{\mathbb{Q}}}$, such a neighborhood base will have a relatively simple structure so this motivates what comes next.

Definition 2.9. Let X be a topological space and fix $x \in X$. A **neighbourhood base** $\mathcal{N}$ for x is a set of (open) neighbourhoods of x such that any neighborhood U of x contains some $N \in \mathcal{N}$.

Because left multiplication defines a homeomorphism, if $\mathcal{N}$ defines a neighborhood base of the identity, then $g\mathcal{N}$ defines a neighborhood base of any g in the topological group. So once one has a neighborhood base of the identity, one has a collection of arbitrarily small neighborhoods throughout G. Such a collection serves to determine properties like continuity, so we'll want to understand neighborhood bases in profinite groups.

Theorem 2.10. *Let $G = \varprojlim_{i \in \mathcal{I}} G_i$ define a profinite group and let $\pi_j : G \to G_j$ denote the projection maps. Then*

$$\mathcal{N} := \{\ker \pi_j : j \in \mathcal{I}\}$$

forms a neighborhood base of the identity.

Proof. We include this proof to establish additional familiarity with the discrete topology of each G_i, the product topology on $\prod G_i$, and the subspace topology on G, as well as show how one uses the directed-ness of $\mathcal{I}$ and the compatibility conditions of the inverse limit. However, only the statement of this theorem — not the slightly technical proof which follows — will become important, so readers should feel secure in skipping or skimming this proof for now.

Let $V \subset \prod_{i \in \mathcal{I}} G_i$ be open. Then the product topology guarantees the existence of some finite $S \subset \mathcal{I}$ for which

$$\prod_{i \notin S} G_i \times \prod_{i \in S} \{1_{G_i}\} \tag{1}$$

defines an open set inside V (note that $\{1_{G_i}\}$ is open inside G_i because we endowed each finite group G_i with the discrete topology). Because $\mathcal{I}$ is directed, we may choose some $j \in \mathcal{I}$ such that $j \geq i$ for all $i \in S$. Replacing S with $S \cup \{j\}$ in equation 1 gives an even smaller open set inside V; we therefore assume, without loss of generality, that the S in equation 1 contains a unique maximal element. In particular, we have shown that if $\mathcal{C}$ denotes the collection of subsets of $\prod_{i \in \mathcal{I}} G_i$ of the form in 1 with S a finite subset of $\mathcal{I}$ containing a unique maximal element, then $\mathcal{C}$ defines a neighborhood base of the identity inside $\prod G_i$. By the definition of the subspace topology, it follows that the collection

$$\mathcal{G} := \{G \cap C : C \in \mathcal{C}\}$$

defines a neighborhood base of the identity inside G. We complete the proof of the theorem by showing that $\mathcal{N} = \mathcal{G}$.

For a fixed j, the compatibility condition of the inverse limit implies that

$$\ker \pi_j = G \cap \left(\prod_{i \neq j} G_i \times \{1_{G_j}\} \right)$$

so that we have the containment $\mathcal{N} \subset \mathcal{G}$. For the other containment, let $C \in \mathcal{C}$, let S denote the finite set which defines C as in equation 1, and let j denote the unique maximal element of S. Then the compatibility condition of the inverse limit — together with the fact that $j \geq i$ for all $i \in S$ — implies that

$$G \cap C = G \cap \left(\prod_{i \notin S} G_i \times \prod_{i \in S} \{1_{G_i}\} \right) = G \cap \left(\prod_{i \neq j} G_i \times \{1_{G_j}\} \right) = \ker \pi_j.$$

Thus, $\mathcal{G} \subset \mathcal{N}$ and we're done. $\square$

Now that we've developed some abstract theory of profinite groups, let's return to the world of number theory and see how everything applies in our favourite example $G_{\overline{\mathbb{Q}}}$. Let $\mathcal{I}$, G_L, and $r_{M,L}$ be as in Theorem 2.5 so that $\varprojlim_{L \in \mathcal{I}} G_L \cong G_{\overline{\mathbb{Q}}}$. For each Galois number field $M/\mathbb{Q}$, we have a projection

$$\begin{aligned} \pi_M : \varprojlim_{L \in \mathcal{I}} G_L &\to G_M \\ (\sigma_L)_{L \in \mathcal{I}} &\mapsto \sigma_M. \end{aligned}$$

From this, we have that $\ker \pi_M$ consists of those sequences $(\sigma_L)_{L \in \mathcal{I}} \in G_{\overline{\mathbb{Q}}}$ for which σ_M is the identity in $G_M = \mathrm{Gal}(M/\mathbb{Q})$. Unfortunately, this perspective isn't terribly illuminating. Rather, we will want to transfer these sequences to the more familiar setting of automorphisms in $\mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$.

Under the isomorphism $\varprojlim_{L \in \mathcal{I}} G_L \cong G_{\overline{\mathbb{Q}}}$ in Theorem 2.5, we may equivalently regard π_M as a projection

$$\begin{aligned} \pi_M : G_{\overline{\mathbb{Q}}} &\to G_M \\ \sigma &\mapsto \sigma|_M. \end{aligned}$$

From this point of view, the kernel of π_M in $G_{\overline{\mathbb{Q}}}$ consists of those automorphisms of $G_{\overline{\mathbb{Q}}}$ which act trivially on $M/\mathbb{Q}$. This proves the following.

Corollary 2.11. *Equip $G_{\overline{\mathbb{Q}}}$ with the profinite topology. Then the collection*

$$\{\mathrm{Gal}(\overline{\mathbb{Q}}/M) : M/\mathbb{Q} \textit{ finite and Galois}\}$$

constitutes a neighbourhood base of the identity (note that here we regard $\sigma \in \mathrm{Gal}(\overline{\mathbb{Q}}/M)$ as living in $G_{\overline{\mathbb{Q}}}$). In particular, the subgroups $\mathrm{Gal}(\overline{\mathbb{Q}}/M)$ are open for $M/\mathbb{Q}$ finite and Galois.

Ultimately, we will use the corollary to determine the continuity of maps from $G_{\overline{\mathbb{Q}}}$. We note, however, that the profinite topology on $G_{\overline{\mathbb{Q}}}$ — often called the Krull topology — has another important use: the Fundamental Theorem of Infinite Galois Theory (FTIGC). We refer the reader to [4] for details, and simply note that the power of the FTIGC serves as yet another motivator for studying profinite topologies.

In addition to relatively simple neighborhood bases, profinite topologies wind up with some other nice properties.

Theorem 2.12. *Let G be a profinite topological group (endowed with the profinite topology). Then G is Hausdorff, compact, and totally disconnected.*

Proof. See the proof of the stronger theorem 7.1.7 in [6]. □

While we won't need the preceding theorem for future arguments, it offers some intuition about the look of profinite topologies and so the look of the profinite topology on $G_{\overline{\mathbb{Q}}}$.

2.3 ℓ-adic Numbers

In this section we define the ℓ-adic numbers — the base field of the ℓ-adic Galois representations we will ultimately study — as the fraction field of an important profinite ring $\mathbb{Z}_\ell$. Recall that we motivated the profinite construction by inspecting the Galois extension $\overline{\mathbb{Q}}/\mathbb{Q}$ and expecting a connection between $\mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$ and $\mathrm{Gal}(L/\mathbb{Q})$ for various intermediate fields L. So we started with something big and tried to craft it from a bunch of small pieces. Here, we will start with a bunch of small pieces which naturally fit together, and then we assemble them — via an inverse limit — into something new. The inverse limit construction of $\mathbb{Z}_\ell$ can be found in [1] while the classical definitions of $\mathbb{Z}_\ell$ have an excellent description in [5].

For positive integers n and m with $m \leq n$, there exists a natural map $\mathbb{Z}/n\mathbb{Z} \to \mathbb{Z}/m\mathbb{Z}$ given by reduction mod m. So we could form an inverse system on the (additive) groups $\{\mathbb{Z}/m\mathbb{Z}\}_{m\in\mathbb{Z}}$ with bonding morphisms given by reduction mod m. While this gives an interesting structure, the so-called profinite completion of $\mathbb{Z}$, we will want to focus attention on a designated prime ℓ and isolate one strand of the full inverse system.

Definition 2.13. Let $I := \mathbb{Z}_{>0}$ be the directed, partially-ordered set given by the ordering of the integers, and fix an integer prime ℓ. For $n \in \mathcal{I}$, set $G_n := \mathbb{Z}/\ell^n\mathbb{Z}$ and for $n \geq m$ define $r_{n,m} : \mathbb{Z}/\ell^n\mathbb{Z} \to \mathbb{Z}/\ell^m\mathbb{Z}$ by reduction mod ℓ^m. Then we have an inverse system

$$\mathbb{Z}/\ell\mathbb{Z} \leftarrow \mathbb{Z}/\ell^2\mathbb{Z} \leftarrow \mathbb{Z}/\ell^3\mathbb{Z} \leftarrow \cdots \leftarrow \mathbb{Z}/\ell^n\mathbb{Z} \leftarrow \cdots$$

where each leftward map denotes reduction mod some power of ℓ. Then we define the **ℓ-adic integers** by

$$\mathbb{Z}_\ell := \varprojlim_{n\in\mathcal{I}} \mathbb{Z}/\ell^n\mathbb{Z}.$$

Because each (additive) group $\mathbb{Z}/\ell^n\mathbb{Z}$ also comes with a ring structure, the profinite limit $\mathbb{Z}_\ell$ inherits the ring structure through coordinate-wise multiplication. So we often call $\mathbb{Z}_\ell$ the **ring of ℓ-adic integers**.

As we did for $G_\mathbb{Q}$, we will determine the profinite topology on the ℓ-adic integers, but we first look at some structural features of $\mathbb{Z}_\ell$. First and foremost, there exists a natural ring embedding of $\mathbb{Z}$ into $\mathbb{Z}_\ell$; define

$$\begin{aligned} \mathbb{Z} &\hookrightarrow \mathbb{Z}_\ell \\ a &\mapsto (a, a, a, a, \ldots) \end{aligned}$$

where $(a, a, a, a, \ldots)$ represents $(a \mod \ell^m)_m \in \prod_m \mathbb{Z}/\ell^m$. The map defines a ring homomorphism because addition and multiplication happen component-wise, and defines an injection because only $0 \in \mathbb{Z}$ maps to $(0, 0, 0, \ldots) \in \mathbb{Z}_\ell$.

But the ℓ-adic integers also contain interesting non-integer elements. Fix $\ell = 3$. Let $x_1 := 1$ and recursively define $x_{n+1} := 3x_n + 1 \mod 3^{n+1}$. Inductively, one sees that $x_n = \sum_{i=1}^{n} 3^{i-1}$ so that we have

$$(x_i) = (1, 4, 13, 40, 121, \ldots)$$

and $x_{n+1} \equiv x_n \mod 3^n$ for all n. Thus, (x_i) defines a 3-adic integer. To see that (x_i) is not in $\mathrm{Im}(\mathbb{Z} \hookrightarrow \mathbb{Z}_\ell)$, note that x_1 satisfies $2x_1 + 1 \equiv 0 \mod 3^1$ and, inductively,

$$2x_{n+1} + 1 = 6x_n + 3 = 3(2x_n + 1) \equiv 0 \mod 3^{n+1}$$

so $2(x_i) + (1, 1, 1, \ldots) = 0$ in $\mathbb{Z}_3$. Thus, (x_i) is a root of $2x + 1$ in $\mathbb{Z}_3$ so that (x_i) defines a 3-adic integer $-\frac{1}{2}$.

While at first the recursive definition above seems arbitrary, it in fact fits into a much more general statement known as "Hensel's Lemma". We won't need Hensel's Lemma for our purposes, but we note an interesting corollary of Hensel's lemma which the previous example demonstrates: for all k with $\ell \nmid k$, there exists an ℓ-adic integer α such that $k\alpha + 1 = 0$. So not only does $\mathbb{Z}_\ell$ possess many "non-integer" elements, but also possesses an inverse for all integers co-prime to ℓ.

For a slightly more sophisticated example, fix $\ell := 7$, define $x_1 := 3$, and recursively define $x_{n+1} = x_n + (x_n^2 - 2)$. We leave it to the reader to verify that

$$(x_i) = (3, 10, 108, 2166, \ldots)$$

defines a 7-adic integer α such that $\alpha^2 - 2 = 0$, namely a square root of two. So $\mathbb{Z}_7$ contains not only pseudo-rationals — like $\frac{1}{2}$ above — but also contains non-rationals like the square root of two.

Now we return to the profinite topology on $\mathbb{Z}_\ell$. By Theorem 2.10, the projections $\pi_n : \mathbb{Z}_\ell \to \mathbb{Z}/\ell^n\mathbb{Z}$ yield a collection of open sets $U(n) := \ker \pi_n$ which form a neighborhood base of the additive identity $(0, 0, \ldots)$. Let $\alpha = (x_1, x_2, \ldots) \in U(n) \subset \mathbb{Z}_\ell$. Then $x_n \equiv 0 \mod \ell^n$ so the compatibility of the reduction maps forces $x_m \equiv 0 \mod \ell^m$ for all $m \leq n$. Thus,

$$(x_i) = (\underbrace{0, \ldots, 0}_{n \text{ times}}, x_{n+1}, x_{n+2}, \ldots)$$

so $U(n) = \ell^n\mathbb{Z}_\ell$. By translating this neighborhood base by other group elements, we obtain a neighborhood base of any element and so obtain the following result on the topology of $\mathbb{Z}_\ell$.

Theorem 2.14. *For $\alpha \in \mathbb{Z}_\ell$ and $n \in \mathbb{Z}_{\geq 0}$, define*

$$U_\alpha(n) = \alpha + \ell^n\mathbb{Z}_\ell.$$

Then the collection of all $U_\alpha(n)$ determines the profinite topology on $\mathbb{Z}_\ell$ (as every open neighborhood of α contains some $U_\alpha(n)$). Distinguish the neighborhood base of the identity 0 by setting $U(n) := U_0(n)$ as above.

And so we have a nice topology on $\mathbb{Z}_\ell$ for which the addition map $\mathbb{Z}_\ell \times \mathbb{Z}_\ell \to \mathbb{Z}_\ell$ is continuous and the projection maps $\mathbb{Z}_\ell \to \mathbb{Z}/\ell^n\mathbb{Z}$ are continuous. Because $\mathbb{Z}_\ell$ has the structure of a ring, we would also like that the topology interacts well with multiplication.

Lemma 2.15. *The multiplication map*

$$\begin{aligned} p : \mathbb{Z}_\ell \times \mathbb{Z}_\ell &\to \mathbb{Z}_\ell \\ (\alpha, \beta) &\mapsto \alpha \cdot \beta \end{aligned}$$

is continuous (where, as always, we endow $\mathbb{Z}_\ell \times \mathbb{Z}_\ell$ with the product topology).

Proof. By the definition of continuity, it suffices to show that the inverse image of any basis set $U_\gamma(n)$, $\gamma \in \mathbb{Z}_\ell$ and $n \in \mathbb{Z}_{\geq 0}$, is itself open. To this end, fix $(\alpha, \beta) \in p^{-1}(\gamma + \ell^n \mathbb{Z}_\ell)$. Then $\alpha \cdot \beta \in \gamma + \ell^n \mathbb{Z}_\ell$ so computing

$$p(U_\alpha(n) \times U_\beta(n)) = \alpha \cdot \beta + \ell^n(\alpha + \beta)\mathbb{Z}_\ell + \ell^{2n}\mathbb{Z}_\ell \subset \gamma + \ell^n \mathbb{Z}_\ell$$

shows that the open set $U_\alpha(n) \times U_\beta(n)$ containing (α, β) lies inside $p^{-1}(U_\gamma(n))$. This completes the proof. □

Now, we finally define the ℓ-adic numbers.

Definition 2.16. The **ℓ-adic numbers**, denoted $\mathbb{Q}_\ell$, is the field of fractions of the ring $\mathbb{Z}_\ell$ of ℓ-adic integers. In analogy with the ℓ-adic integers, for $\alpha \in \mathbb{Q}_\ell$ and $n \in \mathbb{Z}_{\geq 0}$, define

$$U_\alpha(n) := \alpha + \ell^n \mathbb{Z}_\ell$$

and endow $\mathbb{Q}_\ell$ with the topology generated by the sets $U_\alpha(n)$. Note that $\mathbb{Z}_\ell$ is not a typo in the definition of $U_\alpha(n)$ as using $\alpha + \ell^n \mathbb{Q}_\ell$ gives too coarse a topology.

We make three remarks. First, note that $(\ell^m, \ell^m, \ell^m, \ldots)$ — the image of $\ell^m \in \mathbb{Z}$ under the embedding $\mathbb{Z} \hookrightarrow \mathbb{Z}_\ell$ — is not invertible for any m so that the field of fractions $\mathbb{Q}_\ell$ indeed produces something new and "larger". Second, note that there is a natural way of arriving at the given topology on $\mathbb{Q}_\ell$ from the topology on $\mathbb{Z}_\ell$ — by regarding the field of fractions as a quotient of $\mathbb{Z}_\ell \times \mathbb{Z}_\ell$ — but for brevity we simply appeal to intuition about the similarity to the profinite topology we obtained on $\mathbb{Z}_\ell$. Third and finally, the continuity of addition and multiplication in $\mathbb{Q}_\ell$ follows from arguments analogous to the one we made for multiplication in $\mathbb{Z}_\ell$.

In summary, we have a field $\mathbb{Q}_\ell$ which inverts all elements of the ℓ-adic integers together with a natural topology. Under this topology, addition and multiplication are continuous so that $\mathbb{Q}_\ell$ in fact admits the structure of a topological field.

3 ℓ-adic Galois representations

3.1 Definition and properties

This section references many resources, with [1,6,7] used for basic definitions and properties, [8] used for the example of the ℓ-adic cyclotomic character, and [1,9] used for the example of the Galois representations of elliptic curves. We may now leverage all of our setup to at last define ℓ-adic Galois representations.

Definition 3.1. A d**-dimensional** ℓ**-adic Galois representation** is a continuous homomorphism

$$\rho : G_{\overline{\mathbb{Q}}} \to \mathrm{GL}_d(\mathbb{Q}_\ell).$$

Here we endow $G_{\overline{\mathbb{Q}}}$ with the profinite toplogy and $\mathrm{GL}_d(\mathbb{Q}_\ell)$ with the subspace topology of $\mathbb{Q}_\ell^{d^2}$.

As is a theme throughout representation theory, we study representations of $G_{\overline{\mathbb{Q}}}$ as a roundabout method of getting at the structure of $G_{\overline{\mathbb{Q}}}$. Indeed, the map $\rho : G_{\overline{\mathbb{Q}}} \to \mathrm{GL}_d(\mathbb{Q}_\ell)$ transfers the study of some properties of $G_{\overline{\mathbb{Q}}}$ to the study of the simpler structure $\mathrm{GL}_d(\mathbb{Q}_\ell)$. That said, we could theoretically study representations over any field — for example, over $\mathbb{C}$ — so we ought to justify the choice of $\mathbb{Q}_\ell$.

Definition 3.2. Say that a topological group G has **no small subgroup** if some sufficiently small neighborhood of the identity element contains no non-trivial subgroups.

Theorem 3.3. *Let G be a topological group with no small subgroups. Then every continuous homomorphism $\rho : G_{\overline{\mathbb{Q}}} \to G$ has finite image.*

Proof. Let $N \subset G$ be an open neighborhood of the identity which contains no non-trivial subgroup. Then the continuity of ρ implies that $\rho^{-1}(N)$ is open in $G_{\overline{\mathbb{Q}}}$; in particular, Theorem 2.11 implies that there exists some finite Galois extension M of $\mathbb{Q}$ such that $\rho(\mathrm{Gal}(\overline{\mathbb{Q}}/M)) \subset N$. But $\mathrm{Gal}(\overline{\mathbb{Q}}/M)$ is a subgroup of $G_{\overline{\mathbb{Q}}}$ and ρ is a group homomorphism, so $\rho(\mathrm{Gal}(\overline{\mathbb{Q}}/M))$ is a subgroup of G inside N and is therefore trivial. Thus, ρ factors as

$$\begin{tikzcd} G_{\overline{\mathbb{Q}}} \arrow[rr, "\rho"] \arrow[dr, "r"'] & & G \\ & \mathrm{Gal}(M/\mathbb{Q}) \arrow[ur] & \end{tikzcd}$$

with r the natural restriction map. Because $\mathrm{Gal}(M/\mathbb{Q})$ is finite, its image in G is finite so that ρ itself has finite image. $\square$

Thus, groups with small subgroups and groups with no small subgroups detect different aspects of the structure of $G_{\overline{\mathbb{Q}}}$. But because a continuous homomorphism $\rho : G_{\overline{\mathbb{Q}}} \to G$ to a group G with no small subgroup has finite image, such groups are necessarily limited in their ability to detect the structure of the infinite Galois group $G_{\overline{\mathbb{Q}}}$. Many groups fall into this category, including $\mathrm{GL}_d(\mathbb{C})$.

Corollary 3.4. *The group of linear transformations $\mathrm{GL}_d(\mathbb{C})$ has no small subgroup. Thus, every continuous representation $\rho : G_{\overline{\mathbb{Q}}} \to \mathrm{GL}_d(\mathbb{C})$ factors through*

$$\begin{tikzcd} G_{\overline{\mathbb{Q}}} \arrow[rr, "\rho"] \arrow[dr, "r"'] & & \mathrm{GL}_d(\mathbb{C}) \\ & G_L \arrow[ur, "\overline{\rho}"'] & \end{tikzcd}$$

for L a Galois number field, r the natural restriction map, and $\overline{\rho}$ the map which makes the diagram commute.

Proof. We start by proving the "thus". That $\mathrm{GL}_d(\mathbb{C})$ has no small subgroup implies that ρ has finite image by the previous theorem. So $\ker \rho$ has finite index in $G_{\overline{\mathbb{Q}}}$ and the universal property of the quotient implies that ρ factors through a finite quotient group of $G_{\overline{\mathbb{Q}}}$. By the Fundamental Theorem of Galois Theory, the quotient equals the Galois group of some finite Galois extension $L/\mathbb{Q}$, and this gives the desired.

To achieve that $\mathrm{GL}_d(\mathbb{C})$ has no smal subgroup, we appeal to the matrix exponential function $\exp : \mathrm{M}_d(\mathbb{C}) \to \mathrm{GL}_d(\mathbb{C})$. We refer less familiar readers to [10] for an introduction to the matrix exponential sufficient for the argument which follows. Let G be a subgroup of $\mathrm{GL}_d(\mathbb{C})$ which lies inside the unit ball $\mathcal{B}$ centered at the identity I. We will show that G is trivial so $\mathrm{GL}_d(\mathbb{C})$ has no small subgroups.

Take $A \in G \subset \mathcal{B}$. That A has norm less than one implies the existence of some $M \in \mathrm{M}_d(\mathbb{C})$ for which $e^M = A$. Then for all positive integers n, we have $A^n \in G \subset \mathcal{B}$ so $nM \in \exp^{-1}(\mathcal{B})$. But exp is continuous so that $\exp^{-1}(\mathcal{B})$ is bounded. It follows that $M = 0$ so A equals the identity and we're done. $\square$

The map $\overline{\rho}$ in the theorem is an example of a Dirichlet character, an important class of Galois representations over $\mathbb{C}$. Similar ideas show that every Lie group has no small subgroup. In contrast, profinite groups — such as $\mathbb{Q}_\ell$ and $\mathrm{GL}_d(\mathbb{Q}_\ell)$ — always have arbitrarily small subgroups by Theorem 2.11. So it makes sense to study Galois representations over profinite fields. Indeed, our first example of an ℓ-adic Galois representations will have infinite image.

3.2 ℓ-adic Cyclotomic Character

In this section, we describe one of the simplest non-trivial Galois representations: the ℓ-adic cyclotomic character. Fix an integer prime ℓ and for each positive integer n, let ζ_n denote a primitive ℓ^n root of unity. Then the cyclotomic extension $\mathbb{Q}(\zeta_n)$ is a Galois number field, so we have restriction maps $r_n : G_{\overline{\mathbb{Q}}} \to G_{\mathbb{Q}(\zeta_n)}$. Now, we recall that automorphisms in $G_{\mathbb{Q}(\zeta_n)}$ — likewise, automorphisms of $G_{\overline{\mathbb{Q}}}$ — permute the roots of $x^{\ell^n} - 1$ and so permute the ℓ^n roots of unity. In fact, the automorphisms not only permute the roots of unity, but also permute the subset of primitive roots of unity. So for each n it makes sense to define a map $\chi_{n,\ell} : G_{\mathbb{Q}(\zeta_n)} \hookrightarrow (\mathbb{Z}/\ell^n\mathbb{Z})^\times$ by the equation

$$\sigma(\zeta_n) = \zeta_n^{\chi_{n,\ell}(\sigma)}.$$

Altogether, this discussion yields a commuting diagram of homomorphisms

$$\begin{tikzcd}
 & G_{\mathbb{Q}(\zeta_1)} \arrow[r, hook, "\chi_{1,\ell}"] & (\mathbb{Z}/\ell^1\mathbb{Z})^\times \\
G_{\overline{\mathbb{Q}}} \arrow[ur, "r_1"] \arrow[r, "r_2"] \arrow[dr, "r_3"] \arrow[ddr] & G_{\mathbb{Q}(\zeta_2)} \arrow[u] \arrow[r, hook, "\chi_{2,\ell}"] & (\mathbb{Z}/\ell^2\mathbb{Z})^\times \arrow[u] \\
 & G_{\mathbb{Q}(\zeta_3)} \arrow[u] \arrow[r, hook, "\chi_{3,\ell}"] & (\mathbb{Z}/\ell^3\mathbb{Z})^\times \arrow[u] \\
 & \vdots \arrow[u] \arrow[r, hook] & \vdots \arrow[u]
\end{tikzcd}$$

where the vertical arrows are given by the bonding morphisms for the respective inverse systems. Applying the inverse limit, we obtain a homomorphism

$$\chi_\ell : G_{\overline{\mathbb{Q}}} \to \mathbb{Z}_l^\times \subset \mathbb{Q}_\ell = \mathrm{GL}_1(\mathbb{Q}_\ell),$$

the ℓ-adic cyclotomic character. Written out precisely, we have

$$\begin{aligned} \chi_\ell : G_{\overline{\mathbb{Q}}} &\to \mathbb{Z}_\ell^\times \subset \mathbb{Q}_\ell \\ \sigma &\mapsto (\chi_{n,\ell} \circ r_n(\sigma))_{n \in \mathbb{Z}_{\geq 0}}. \end{aligned}$$

The first definition hints at how one might go about constructing an ℓ-adic Galois representations in general, while the latter more-detailed description unravels the commutative diagram and will prove useful for computations with χ_ℓ.

Theorem 3.5. *The ℓ-adic cyclotomic character χ_ℓ is continuous and so defines a 1-dimensional ℓ-adic Galois representation.*

Proof. As we've used before, it suffices to show that the inverse image of any basis set in $\mathbb{Q}_\ell$ is open in $G_{\overline{\mathbb{Q}}}$. In fact, that $\mathbb{Q}_\ell$ has the structure of a topological field — so multiplication is continuous — means we need only show that the inverse image of $U_1(m)$ is open for all m (as these constitute a neighborhood base of the identity). We compute

$$\begin{aligned} \chi_\ell^{-1}(U_1(n)) &= \{\sigma \in G_{\overline{\mathbb{Q}}} : (\chi_{m,\ell} \circ r_m(\sigma))_m \in U_1(n)\} \\ &= \{\sigma \in G_{\overline{\mathbb{Q}}} : \chi_{m,\ell} \circ r_m(\sigma) = 1 \in \mathbb{Z}/\ell^m\mathbb{Z} \text{ for } m \leq n\} \\ &= \{\sigma \in G_{\overline{\mathbb{Q}}} : \sigma \text{ fixes } \mathbb{Q}(\zeta_n)\} \end{aligned}$$

to see that $\chi_\ell^{-1}(U_1(n)) = \mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q}(\zeta_n))$. Because $\mathbb{Q}(\zeta_n)$ defines a Galois extension of $\mathbb{Q}$, Corollary 2.11 shows that $\mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q}(\zeta_n))$ is open in $G_{\overline{\mathbb{Q}}}$ and we're done. □

Now that we've defined the ℓ-adic cyclotomic character — our first ℓ-adic Galois representation! — one naturally wonders about the utility of such a representation: in what ways does χ_ℓ offer insight into the structure of $G_{\overline{\mathbb{Q}}}$? Of particular interest in $G_{\overline{\mathbb{Q}}}$ are the so-called "Frobenius elements" $f_\mathfrak{p}$, where $\mathfrak{p}$ denotes a prime ideal in some $\mathbb{Z}$-submodule of $\overline{\mathbb{Q}}$ (see a text on algebraic number theory for the precise definition). Then there exists an integer prime p for which $\mathfrak{p} \cap \mathbb{Z} = (p)$ and it turns out that $\chi_\ell(\mathrm{Frob}_\mathfrak{p}) = p$ as long as $p \neq \ell$. So the cyclotomic character indeed detects something of the structure of $\mathfrak{p} \subset \overline{\mathbb{Q}}$ via its image on $f_\mathfrak{p} \in G_{\overline{\mathbb{Q}}}$.

This example also illustrates another general phenomenon: ℓ-adic Galois representations often "break down" around finitely-many "bad" primes, where the precise notion of bad depends on the nature of the structures involved in crafting the representation. We see this in the condition $p \neq \ell$ above. For this reason, it is quite useful to study ℓ-adic Galois representations for ranging primes ℓ, rather than always considering a fixed prime.

3.3 ℓ-adic Galois representation of an elliptic curve

In Section 1, we defined an elliptic curve E over $\mathbb{Q}$ is a "nice" curve defined by an equation of the form

$$y^2 + a_1xy + a_3y = x^3 + a_2x^2 + a_4x + a_6 \qquad a_i \in \mathbb{Q}. \tag{2}$$

While we often picture these curves as living in the plane $\mathbb{R} \times \mathbb{R}$ — as in Figure 1 — graphs don't capture the full picture. We will regard any point $(x, y) \in \overline{\mathbb{Q}} \times \overline{\mathbb{Q}}$ which satisfies equation 2 as lying on the elliptic curve E; many such points have non-real coordinates and so don't appear in graphs of the real plane. For example, the non-real points $(0, \pm i) \in \overline{\mathbb{Q}}^2$ lie on the curve $y^2 = x^3 + x$ in Figure 1.

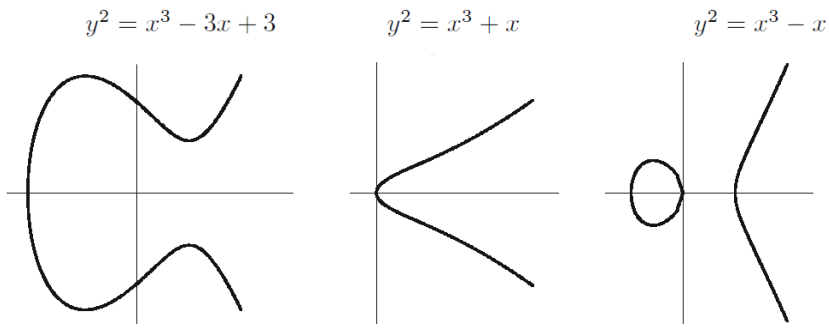

Figure 1: The real points on three elliptic curves. Image adapted from [9].

Elliptic curves possess a wealth of applications in both number theory and cryptography, such as in the proof of Fermat's Last Theorem as in Section 1. Their utility results from the following beautiful result.

Theorem 3.6. *Let E be an elliptic curve defined by some equation 2. Let*

$$E(\overline{\mathbb{Q}}) = \{(x, y) \in \overline{\mathbb{Q}}^2 : (x, y) \text{ lies on the curve } E\} \cup \{\infty\}$$

denote the set of points on E together with a "point out at infinity". Then there is an abelian group structure on $E(\overline{\mathbb{Q}})$ with identity element ∞.

Proof. Many excellent sources detail the group law of an elliptic curve, such as [9], so we refer the reader there. □

For our purposes — defining the ℓ-adic Galois representation of an elliptic curve — we can black-box the details of group law.

The presence of a group law means that there may exist torsion points in $E(\overline{\mathbb{Q}})$: points P such that for some positive integer n,

$$nP = \underbrace{P + \cdots + P}_{n \text{ times}} = \infty.$$

As when we defined the ℓ-adic integers, rather than consider all possible n we will consider powers of a fixed prime ℓ. Indeed, notate the ℓ^n torsion points by

$$E[\ell^n] := E(\overline{\mathbb{Q}})[\ell^n] := \{P \in E(\overline{\mathbb{Q}}) : \ell^n P = \infty\}$$

so that we have an inverse system

$$E[\ell] \leftarrow E[\ell^2] \leftarrow E[\ell^3] \leftarrow \cdots E[\ell^n] \leftarrow E[\ell^{n+1}] \cdots$$

where each leftward map is given by multiplication by ℓ. This works because if P is an ℓ^{n+1}-torsion point, then ℓP is an ℓ^n-torsion point. Recall that we defined the ℓ-adic cyclotomic character via a group action of $G_{\overline{\mathbb{Q}}}$ which we then transferred to $\mathbb{Q}_\ell$ via homomorphisms to $(\mathbb{Z}/\ell^n\mathbb{Z})^\times < \mathbb{Z}/\ell^n\mathbb{Z}$. An analogous strategy works here by using the ℓ^n-torsion subgroups.

There is a natural component-wise action of $G_{\overline{\mathbb{Q}}}$ on the points of $E(\overline{\mathbb{Q}})$: $\sigma \cdot (x, y) = (\sigma(x), \sigma(y))$ and $\sigma \cdot \infty = \infty$. This assignment makes sense because the fact that (x, y) lies on E implies

$$\begin{aligned} y^2 + a_1xy + a_3y &= x^3 + a_2x^2 + a_4x + a_6 \\ \sigma(y^2 + a_1xy + a_3y) &= \sigma(x^3 + a_2x^2 + a_4x + a_6) \\ \sigma(y)^2 + a_1\sigma(x)\sigma(y) + a_3\sigma(y) &= \sigma(x)^3 + a_2\sigma(x)^2 + a_4\sigma(x) + a_6 \end{aligned}$$

so $(\sigma(x), \sigma(y))$ also lies on E. Because we didn't explicitly define the group law on E, we assert without proof that for all points P and Q on E,

$$\sigma \cdot (P + Q) = \sigma(P) + \sigma(Q).$$

Lastly, we note that $\sigma^{-1} \cdot (\sigma(x), \sigma(y)) = (x, y)$ so we may regard the action of σ on $E(\overline{\mathbb{Q}})$ as a (group) automorphism of $E(\overline{\mathbb{Q}})$ and so we have a map

$$G_{\overline{\mathbb{Q}}} \to \mathrm{Aut}(E(\overline{\mathbb{Q}})).$$

A group automorphism restricts to torsion subgroups, so for all n we have a map $G_{\overline{\mathbb{Q}}} \to \mathrm{Aut}(E[\ell^n])$ and thus a commutative diagram

$$\begin{tikzcd} & \mathrm{Aut}(E[\ell^1]) & \\ G_{\overline{\mathbb{Q}}} \arrow[ur] \arrow[r] \arrow[dr] \arrow[ddr] & \mathrm{Aut}(E[\ell^2]) \arrow[u] & \\ & \mathrm{Aut}(E[\ell^3]) \arrow[u] & \\ & \vdots \arrow[u] & \end{tikzcd}$$

with the upward arrows induced by multiplication by ℓ. To get a Galois representation out of this commutative diagram, we need to relate the above to $\mathbb{Z}/\ell^n\mathbb{Z}$ as we did with the cyclotomic characters; for this we apply the following theorem.

Theorem 3.7. *That $\overline{\mathbb{Q}}$ is algebraically closed implies $E[\ell^n] \cong (\mathbb{Z}/\ell^n\mathbb{Z})^2$.*

Proof. While we haven't developed the background required to prove the theorem in general — for this, see [9] — we can offer some explanation for the case of $\ell^n = 2$. It turns out that the two-torsion points of an elliptic curve are precisely the zeroes of the curve. So that $\overline{\mathbb{Q}}$ is

algebraically closed implies that the cubic equation defining E, which has coefficients in $\mathbb{Q}$, has three roots $P_1, P_2, P_3 \in E(\overline{\mathbb{Q}})$. Then

$$E[2] = \{\infty, P_1, P_2, P_3\}$$

is a group of order four each of whose non-identity elements has order two, so $E[2]$ is isomorphic to the Klein four group $(\mathbb{Z}/2\mathbb{Z})^2$ as desired. $\square$

The theorem implies $\mathrm{Aut}(E[\ell^n]) \cong \mathrm{GL}_2(\mathbb{Z}/\ell^n\mathbb{Z})$. Because the modding-by-ℓ inverse system on $\{\mathbb{Z}/\ell^n\mathbb{Z}\}_n$ induces a natural inverse system on $\{\mathrm{GL}_2(\mathbb{Z}/\ell^n\mathbb{Z})\}_n$, we may augment our earlier diagram to obtain

$$\begin{tikzcd}
 & \mathrm{Aut}(E[\ell^1])] \arrow[r, "\sim"] & \mathrm{GL}_2(\mathbb{Z}/\ell^1\mathbb{Z}) \\
G_{\overline{\mathbb{Q}}} \arrow[ur] \arrow[r] \arrow[dr] \arrow[ddr] & \mathrm{Aut}(E[\ell^2]) \arrow[u] \arrow[r, "\sim"] & \mathrm{GL}_2(\mathbb{Z}/\ell^2\mathbb{Z}) \arrow[u] \\
 & \mathrm{Aut}(E[\ell^3]) \arrow[u] \arrow[r, "\sim"] & \mathrm{GL}_2(\mathbb{Z}/\ell^3\mathbb{Z}) \arrow[u] \\
 & \vdots \arrow[u] \arrow[r] & \vdots \arrow[u]
\end{tikzcd}$$

where the vertical sections form two inverse systems. Apply the inverse limit as we did for the cyclotomic characters to obtain a homomorphism

$$\rho_{E,\ell} : G_{\overline{\mathbb{Q}}} \to \mathrm{GL}_2(\mathbb{Z}_\ell) \subset \mathrm{GL}_2(\mathbb{Q}_\ell).$$

Theorem 3.8. *Let E be an elliptic curve over $\mathbb{Q}$. The map $\rho_{E,\ell}$ is continuous and so defines a 2-dimensional ℓ-adic Galois representation, the ℓ-adic Galois representation of E.*

Proof. The proof proceeds rather similarly to that of the ℓ-adic cyclotomic character. We encourage the reader to think about it as an exercise and reference [1] as necessary. $\square$

The representations $\rho_{E,\ell}$ are precisely the ℓ-adic Galois representations in the Shimura-Taniyama Conjecture (STC) of Section 1. A similar construction — with a fair bit more technical background — gives ℓ-adic Galois representations of "modular forms". STC then corresponds the resulting ℓ-adic Galois representations with modular forms (away from finitely-many "bad" primes ℓ) and in doing so associates to each elliptic curve a family of modular forms. So ℓ-adic Galois representations form the bridge at the heart of Fermat's Last Theorem, and here we've seen the bridge where it leaves the world of elliptic curves.

Before concluding this section, we remark on how the construction of $\rho_{E,\ell}$ generalises to broader classes of algebraic curves. We required three facts about elliptic curves:

1. there is a group structure on the points of $E(\overline{\mathbb{Q}})$,

2. there is an action of $G_{\overline{\mathbb{Q}}}$ on the torsion subgroups, and

3. the torsion subgroups relate to $\mathbb{Z}/\ell^n\mathbb{Z}$ for some n (in the case of elliptic curves, $n = 2$).

So in principal, if we are given some "algebraic curve" over $\mathbb{Q}$ with an associated group on which $G_{\overline{\mathbb{Q}}}$ acts, and whose torsion relates to $\mathbb{Z}/\ell^n\mathbb{Z}$, we may go through an analogous construction to obtain an ℓ-adic Galois representation associated to the given curve. That this construction may proceed in general comes from the theory of "etalé cohomology".

4 Concluding Remarks and Big Picture

We've discussed the way in which one can transfer from elliptic curves to Galois representations (more or less with precision) and then from Galois representations to modular forms (only in passing). As we hinted in the introduction and in Section 1, a much broader picture continues to emerge. There are three worlds at play: "algebraic varieties" over $\mathbb{Q}$, "geometric representations" of $G_{\overline{\mathbb{Q}}}$, and "automorphic forms". Elliptic curves live in the first, ℓ-adic Galois representations in the second, and modular forms in the third, so we have at least mentioned an object from each of the three worlds. Figure 2 summarises the known — and unknown — relationships between the three.

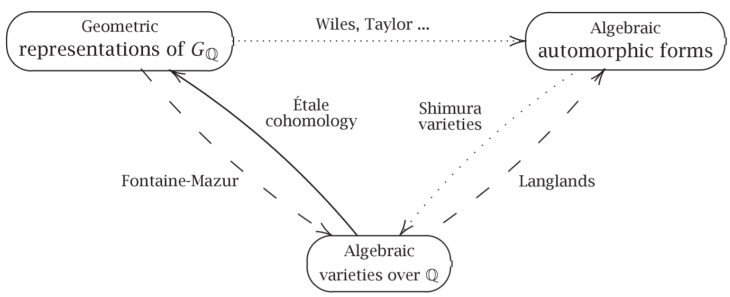

Figure 2: Solid arrows indicate complete results, dotted arrows indicate partial progress, and dashed arrows indicate conjectures. Note that $G_{\mathbb{Q}}$ denotes $\mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$ here. Image taken from [11].

The figure demonstrates that much remains unknown! The "Langlands program" — so-called because it is based in an array of conjectures from Robert Langlands — forms the bridge from algebraic varieties to automorphic forms. These conjectures lie at the forefront of modern number theory as most remain barely-scratched. Indeed, note the absence of an arrow from automorphic forms to geometric representations. So although the Shimura-Taniyama Conjecture (and later developments) allow one to go from geometric representations to automorphic forms, the lack of a bridge the other way remains an unsolved mystery.

Nevertheless, the twentieth and now twenty-first centuries have witnessed great progress with some astounding results. We discussed the diagram's sole solid arrow in Section 3.3: elliptic curves to Galois representations, known in its more general form as "etalé cohomology". This correspondence bridges the world of curves over $\mathbb{Q}$ and the world of Galois representations. By then following the work of Wiles, Taylor, and others over the world of automorphic forms, one may pass from pass from Galois representations all the way over to

modular forms and in doing so prove Fermat's Last Theorem. So this vast theory also yields remarkable, down-to-earth results. As another example, the dotted arrow labelled "Shimura varieties" served to prove the "Ramanujan conjecture". The conjecture outlines some crucial properties of the Ramanujan τ function which has since found a wide variety of applications. As such, Figure 2 captures a wealth of mathematical progress and undoubtedly hides a trove of future results.

References

[1] Fred Diamond and Jerry Shurman *A First Course in Modular Forms*. Springer Science and Business Media, New York, 2005, pp. 371-412.

[2] Henri Darmon and Fred Diamond and Richard Taylor *Fermat's Last Theorem*. Online, September 2007, pp. 3-14.

[3] Kyrie Johnson *Modular Forms, Elliptic Curves, and their Connection to Fermat's Last Theorem*. Online, 2020, pp. 7-9.

[4] J.S. Milne *Fields and Galois Theory*. Online, University of Michigan, September 2018, pp. 91-100.

[5] Paul Garrett *Classical Definitions of $\mathbb{Z}_p$ and $\mathbb{A}$*. Online, University of Minnesota, September 2010, pp. 1-19.

[6] Jonatan Lindell *Profinite Groups and Infinite Galois Extensions*. Online, Uppsala University, November 2019, pp. 29-43.

[7] Andrew Gleason *Groups without Small Subgroups*. Annals of Mathematics, Vol. 56, No. 2, September 1952.

[8] Joel Bellaiche *Galois Representations*. Online, Brandeis University, Spring 2009, pp. 9-10.

[9] Joseph Silverman *The Arithmetic of Elliptic Curves*. Springer Science and Business Media, New York, 1994, pp. 41-114.

[10] Andrew Baker *An introduction to matrix groups and their applications*. Online, University of Glasgow, July 2000, pp. 1-16.

[11] Mark Kisin *What is a Galois Representation?*. Notices of the AMS, 2007.