The Galois correspondence

John Galindo,
advised by Professor Conidis and Professor Kofman

Spring 2024

1 Introduction

In the subject of algebraic topology, there is the notion of a "covering space." A **covering space** of a topological space X is a space $\tilde{X}$ with a map $p : \tilde{X} \to X$ with the property that every point of X belongs to an open neighborhood U such that $p^{-1}(U)$ is the disjoint union of homeomorphic copies of U in $\tilde{X}$. The typical example of a covering space is the projection of $\mathbb{R}$ onto the circle S^1, described by the composition of maps $\mathbb{R} \to \mathbb{R}^3 \to S^1 \hookrightarrow \mathbb{R}^2$ given by $t \mapsto (cos(2\pi t), sin(2\pi t), t) \mapsto (cos(2\pi t), sin(2\pi t))$. Visually, these maps describe a helix in $\mathbb{R}^3$ projecting downwards onto the circle.

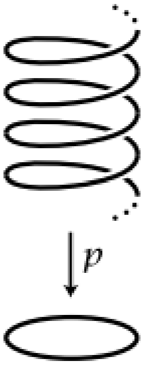

The covering space $p : \mathbb{R} \to S^1$, as illustrated in "Algebraic Topology" by Allen Hatcher

In the subject of field theory, one is often concerned with **field extensions** of a field F. K is an extension of F if K is a field containing F, and this is denoted K/F. In particular, one is usually concerned with **algebraic extensions** of F, extensions that only add elements that are the roots of polynomials with coefficients in F. An important type of algebraic extension is the **splitting field** of a polynomial $f(x)$ over F, or the minimum extension of F that contains every root of $f(x)$.

A natural first project when given a new mathematical object is to classify all of its instances. Given a space X, one would want to know what covering spaces of X exist, what criteria must be satisfied for them to exist, and how different covering spaces are related. While most fields clearly have an infinite amount of algebraic extensions, similar classification questions can be asked about the fields that lie between a giving splitting field and its base field. Surprisingly, these ostensibly unrelated objects obey a very similar structure in their classification. This structure is the titular Galois correspondence, which can be briefly described as a one-to-one but inverse correspondence between the subgroups of a group of isomorphisms from the object (covering space or splitting field) to itself and the subobjects (covers, fields) lying between the object and its base (the covered space X or the extended field F).

This paper will briefly introduce the point-set topology, algebraic topology, and field theory needed to motivate and understand the Galois correspondence as it pertains to covering spaces and to Galois extensions.

2 The Topological Story

A **topological space** (X, τ) is a pair of sets X and $\tau \subseteq 2^X$ such that the elements of τ behave similarly to open intervals on $\mathbb{R}$. Specifically, (X, τ) must obey

1. $X \in \tau$ and $\emptyset \in \tau$
2. The union of any elements in τ is an element of τ
3. The finite intersection of any elements in τ is an element of τ

If (X, τ) is a topological space, the elements of τ are referred to as **open sets**, and τ itself as a **topology**. The purpose of a topology is to give structure to a set. For example, $\mathbb{Z}$ and $\mathbb{Q}$ have the same cardinality, but the standard topology given to each is different: $\mathbb{Z}$ is treated as a series of isolated points (it has the "discrete topology" $\tau = 2^{\mathbb{Z}}$), whereas $\mathbb{Q}$ is considered "dense" as every open interval contains infinitely many points. Most of the spaces in this paper have the topology induced by intersecting the set with open n-balls. For example, the topology on a figure X lying in 3-space will usually be the intersection of X with "open spheres" in 3-space, and any unions or intersections that arise thereafter. Additionally, by abuse of notation this paper will simply refer to spaces (X, τ) by X, when the topology is clear.

Two spaces X and Y are considered **topologically equivalent**, or **homeomorphic**, if there is a **homeomorphism** $f : X \to Y$ between them. A homeomorphism is a bijection with the property that $U \subset Y$ is an open set of Y if and only if $f^{-1}(U)$ is an open set of X. An obvious example of two homeomorphic spaces are two circles with different (nonzero) radii. The radius is a geometric property and has no bearing on the broader topological structure of the circle: a homeomorphism can simply be the appropriate translation and dilation in $\mathbb{R}^2$, for example. Similarly, the perimeter of a square is also homeomorphic to a circle.

Less obviously, take a solid square ($[0, 1] \times [0, 1]$) and "glue" its opposite edges straight across. That is, consider the point $(0, y)$ to be the same as point $(1, y)$, and consider the point $(x, 0)$ to be the same as $(x, 1)$. The topology on this set is the standard topology, with the caveat that an open set can't contain one of these "glued" points on its boundary and must extend past these points slightly on both sides of the square if needed (This is an example of a **quotient space**, of which the precise workings are outside the scope of this paper). This space is homeomorphic to the torus, or "donut shape."

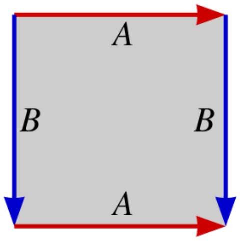

These associations on a square form a torus.

A tool to determine that two spaces are equivalent demands a tool to determine that they are not. The principal technique is to find a "topological property" that would be preserved by homeomorphism, and show that one space has it while the other does not. One such property is **connectedness**. If a space can be fully covered by two non-overlapping open sets, that space is said to be disconnected, otherwise it is connected. This should reassure the reader that a circle is not homeomorphic to two circles. Even if both spaces X and Y are connected, it might be possible in X but not Y to *create* a disconected space by removing a single point. This technique proves that $\mathbb{R}$ is not homeomorphic to S^1. Another property familiar to students of analysis is that of **compactness**. In Euclidean space, this is the property of being closed and bounded, so the open interval (0,1) is topologically distinct from the closed interval [0,1].

Much of covering space theory hinges on a stronger topological property, the **fundamental group** of the space, $\pi_1(X, x_0)$. Given a space X and a basepoint x_0, this is the set of homotopy classes of loops in X that start and end at x_0. Two loops $l, k : I \to X$ are said to be **homotopic** if there is a continuous function $f : I^2 \to X$ with $f(0, t) = l(t)$, $f(1, t) = k(t)$, and $f(s, 0) = f(s, 1) = x_0$ for all s. This corresponds to the idea of a "continuous deformation" of one loop into another, and in the context of the fundamental group means that any two loops that can be deformed into each other are considered to be the same element. It's worth mentioning that while the choice of basepoint is an important detail, every space presented for the rest of this paper will be **path-connected**, meaning that any two points can be connected by a continuous path, so any two points $x_0, x_1 \in X$ satisfy $\pi_1(X, x_0) \cong \pi_1(X, x_1)$.

It is possible that every loop in X can be contracted to the constant path x_0. In that case, X is said to be **simply connected** (assuming it is also path-connected). $\mathbb{R}^n$ is connected for any choice of n, as is S^n for any $n \geq 2$. S^1, however, is not. Without delving into computation, notice that the loop that travels around the circle once clockwise cannot be deformed into a loop that does so twice—or clockwise, or not at all—without breaking the loop or leaving the circle. The fundamental group therefore corresponds exactly to the number of times the loop winds around the circle, with clockwise winding taken to be ”positive” and counterclockwise taken to be ”negative”. So, $\pi_1(S^1, x_0) \cong \mathbb{Z}$. The circle provides a good example of the group operation of π_1: concatenating two loops by doing one after the other. Here, traversing the circle five times clockwise and thrice counterclockwise is the same as traversing the circle twice clockwise.

More complicated spaces may imply more complicated algebra. The torus is the Cartesian product of two circles, so it's not terribly surprising that its fundamental group is $\mathbb{Z} \oplus \mathbb{Z}$. However, it is interesting that the fundamental group of this space is commutative: traversing three times around a longitude and twice around a meridian is the same as traversing twice around the meridian and three times around the longitude, which is the same still as traversing along a ”three-two curve” that does both before returning to x_0. This is specifically because of the ”connective tissue” between the meridian and longitude, and is not true in general.

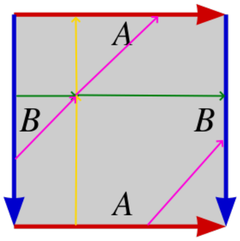

The pink loop is homotopic to the concatenation of the green and yellow loops.

To contrast, consider $S_1 \vee S_1$, the ”wedge space” of two circles joined at a point. The fundamental group of this space is *not* $\mathbb{Z} \oplus \mathbb{Z}$ because it is not commutative: the loop a does not commute with the loop b. The fundamental group is instead $\mathbb{Z} * \mathbb{Z}$, the ”free product” on two generators. This is a group in which the only algebraic relations are $aa^{-1} = a^{-1}a = e = bb^{-1} = b^{-1}b$, such that the elements $b^{-2}a^3b$, $a^2b^{-1}a$, and $aba^{-1}b^{-1}a^2b^{-1}$ are all distinct.

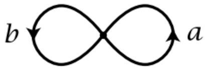

The wedge of two circles, $S_1 \vee S_1$. Illustration by Hatcher.

Let this paper now return to the topic of covering spaces. To reiterate, a **covering space** of X is a space $\tilde{X}$ with a map $p : \tilde{X} \to X$ with the property that every point of X belongs to an open neighborhood U such that $p^{-1}(U)$ is the disjoint union of homeomorphic copies of U in $\tilde{X}$. If a covering space is simply-connected, it can be further called the **universal cover** of X. That $\mathbb{R}$ is the universal cover of S^1 has already been shown, so now consider the torus. The universal cover of the torus can be constructed in the following way: let $\tilde{X} = \mathbb{R}^2$, and tile $\tilde{X}$ in unit squares such that the origin is a vertex. The covering map $p : \mathbb{R}^2 \to T^2$ is then the map that takes each point of each unit square to the corresponding point on the square representation of the torus. $\mathbb{R}^2$ is simply-connected, so this is the universal cover of the torus.

To begin classifying these covering spaces, one needs a notion of when two covering spaces are equivalent. A **covering space isomorphism** between two covering spaces $p : \tilde{X_1} \to X$ and $q : \tilde{X_2} \to X$ is a homeomorphism $f : \tilde{X_1} \to \tilde{X_2}$ satisfying $p = qf$. That is, f is a homeomorphism that preserves the covering of X, and thus preserves all information about the covering space. An isomorphism $\tilde{X} \to \tilde{X}$ is called a **deck transformation**, and the set of all deck transformations of $\tilde{X}$ forms a group $G(\tilde{X})$.

One more definition is needed: if $\tilde{x}$ and $\tilde{x}'$ are any two points in $\tilde{X}$ that have the same image in X, then $\tilde{X}$ is **normal** if there is a deck transformation of $\tilde{X}$ with $\tilde{x} \mapsto \tilde{x}'$.

With these definitions, this paper will now present the most important theorems classifying the covering spaces of X (given a few assumptions about X).

Theorem 1. *The map $p_* : \pi_1(\tilde{X}, \tilde{x_0}) \to \pi_1(X, x_0)$ induced by a covering map $p : \tilde{X} \to X$ is injective, and the image subgroup $p_*(\pi_1(\tilde{X}, \tilde{x_0}))$ in $\pi_1(X, x_0)$ consists of loops in X that ”lift” to loops in $\tilde{X}$ (whose preimages in $\tilde{X}$ are loops).*

Briefly, this is because of the ”homotopy lifting property” of covering spaces: the kernel of p_* consists of loops in $\tilde{X}$ whose image can be deformed to a point in X, but the homotopy lifting property states that this homotopy can be ”lifted” to $\tilde{X}$ and thus the original loop must also deform to a point in $\tilde{X}$, so the kernel is trivial.

The significance of this result is that the fundamental group of a covering space must be a subgroup of the fundamental group of the space being covered. This is the first of a few significant results classifying covering spaces, and it places a restriction on what covering spaces are *not* possible: for instance, the circle cannot be a covering space of the real line.

Theorem 2. *Suppose X is path-connected, locally path-connected, and semilocally simply-connected. Then for every subgroup H of $\pi_1(X, x_0)$, there is a covering space with $p_*(\pi_1(\tilde{X}, \tilde{x_0})) = H$.*

First, some vocabulary: **locally path-connected** spaces are those in which every point belongs to small open sets that are path-connected, and **semilocally simply-connected** spaces are those in which every point belongs to a small open set whose loops can be contracted to a point in the space as a whole.

More importantly, Theorem 2 establishes what covering spaces *are* possible for a nicely-behaved space X: there's one for every subgroup of $\pi_1(X, x_0)$. This confirms that every X meeting this criteria has a simply-connected covering space, and more generally implies a staggering amount of possible covering spaces for spaces with large (in the sense of subgroups) π_1. The wedge of two circles, for instance, admits many covering spaces because of its free fundamental group (diagrams from Allen Hatcher's textbook *Algebraic Topology* on the next page).

Theorem 3. *If X is path-connected, locally path-connected, and semilocally simply-connected, then there is a bijection between the path-connected covering spaces of X (up to basepoint-preserving isomorphism) and the subgroups of $\pi_1(X, x_0)$. This bijection is given by $p_*(\pi_1(\tilde{X}, \tilde{x_0})) \leftrightarrow (\tilde{X}, \tilde{x_0})$. If basepoints are ignored, then this bijection instead exists between the path-connected covering spaces of X and the conjugacy classes of these subgroups, and the conjugacy action amounts to a change of basepoint.*

Theorem 4. *Suppose X is path-connected and locally path-connected, let $p : (\tilde{X}, \tilde{x_0}) \to (X, x_0)$ be a path-connected covering space, and let H be the image subgroup $p_*(\pi_1(\tilde{X}, \tilde{x_0}))$. Then*

a) This covering space is normal if and only if H is a normal subgroup of $\pi_1(X, x_0)$.

b) $G(\tilde{X}) \cong N(H)/H$, where $N(H)$ is the normalizer of H in $\pi_1(X, x_0)$. If this is a normal covering space, then $G(\tilde{X}) \cong \pi_1(X, x_0)/H$.

These last two conditions completely characterize the path-connected covering spaces of X: Given a well-behaved space X, the covering spaces correspond exactly to the conjugacy classes of the subgroups of $\pi_1(X, x_0)$, and fixing the basepoint creates an even more pleasant correspondence between these covering spaces and the subgroups themselves. Further, there is an inverse relationship between the elements of $G(\tilde{X})$ and the elements of $\pi_1(\tilde{X}, \tilde{x_0}$. That is, a covering space with a greater fundamental group will have fewer deck transformations, and vice versa.

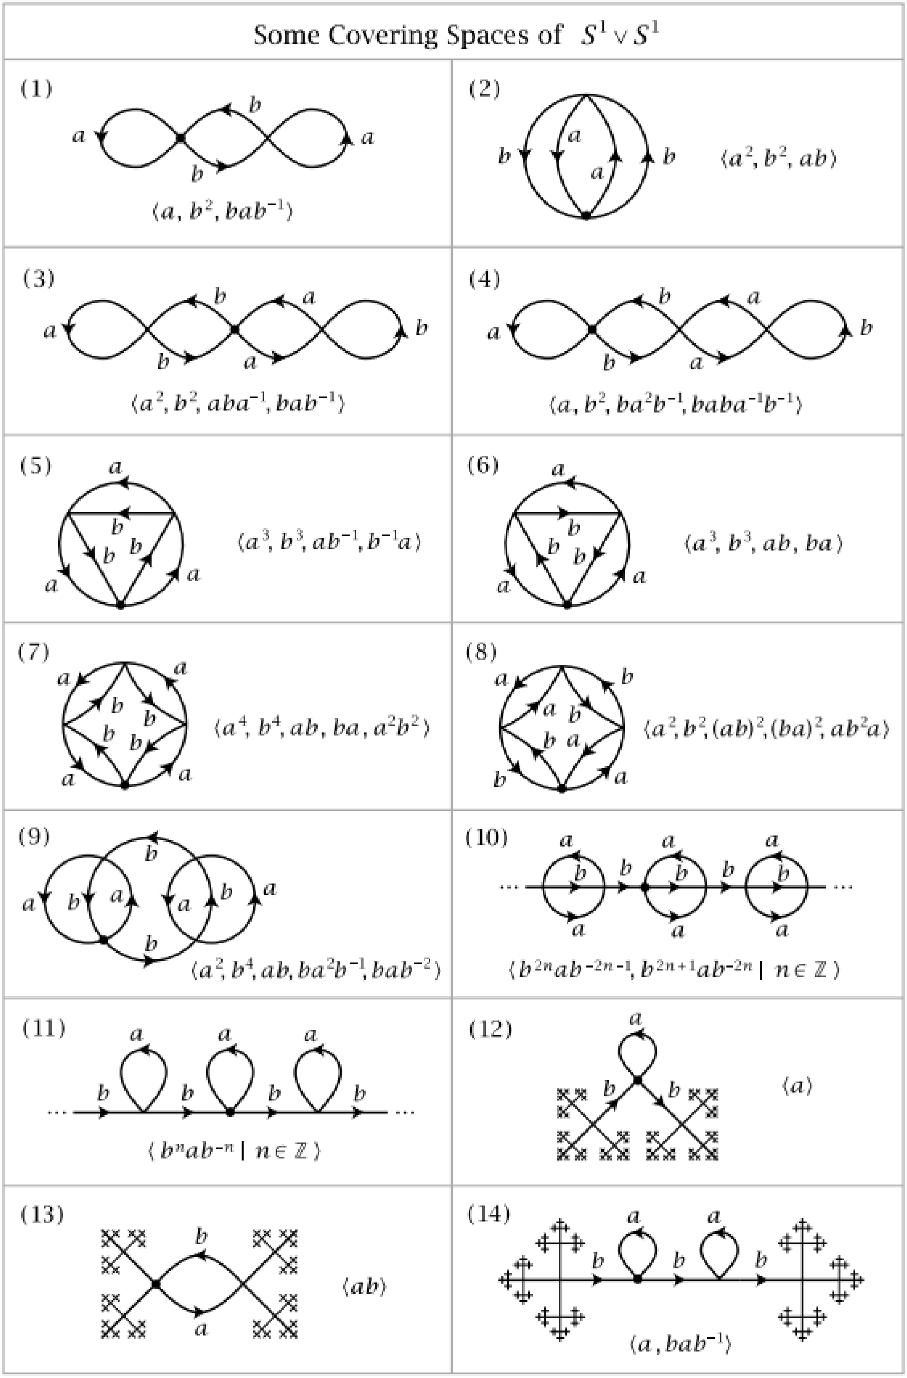
Some Covering Spaces of $S^1 \vee S^1$
(1)
$\langle a, b^2, bab^{-1} \rangle$
(2)
$\langle a^2, b^2, ab \rangle$
(3)
$\langle a^2, b^2, aba^{-1}, bab^{-1} \rangle$
(4)
$\langle a, b^2, ba^2b^{-1}, baba^{-1}b^{-1} \rangle$
(5)
$\langle a^3, b^3, ab^{-1}, b^{-1}a \rangle$
(6)
$\langle a^3, b^3, ab, ba \rangle$
(7)
$\langle a^4, b^4, ab, ba, a^2b^2 \rangle$
(8)
$\langle a^2, b^2, (ab)^2, (ba)^2, ab^2a \rangle$
(9)
$\langle a^2, b^4, ab, ba^2b^{-1}, bab^{-2} \rangle$
(10)
$\langle b^{2n}ab^{-2n-1}, b^{2n+1}ab^{-2n} \mid n \in \mathbb{Z} \rangle$
(11)
$\langle b^nab^{-n} \mid n \in \mathbb{Z} \rangle$
(12)
$\langle a \rangle$
(13)
$\langle ab \rangle$
(14)
$\langle a, bab^{-1} \rangle$
a
b

Assorted covering spaces of $S_1 \vee S_1$. Illustrations by Hatcher.

]

3 The Algebraic Story

Recall that a field F is a ring with the property that $F \setminus \{0\}$ is an Abelian group under multiplication. K is said to be an **extension field** over F if K is a field containing F. This is denoted by K/F. A field extension can be thought of as a vector space over the base field: for example, $\mathbb{C}$ is an extension of $\mathbb{R}$, and the elements of $\mathbb{C}$ are linear combinations of 1 and i with coefficients in $\mathbb{R}$. $\mathbb{C}$ is not unique in this regard: one can similarly define the extension $\mathbb{Q}(\sqrt[3]{2})/\mathbb{Q}$ as the smallest field containing both $\mathbb{Q}$ and $\sqrt[3]{2}$. The elements of $\mathbb{Q}(\sqrt[3]{2})$ then take the form $a + b\sqrt[3]{2} + c\sqrt[3]{4}$, where $a, b, c \in \mathbb{Q}$.

$\mathbb{C}$ is a **degree two** extension over $\mathbb{R}$, and $\mathbb{Q}(\sqrt[3]{2})/\mathbb{Q}$ is similarly a degree three extension. More generally, the **degree** $[K : F]$ of a field extension is its dimension as a vector space over the base field. This paper is mostly concerned with finite extensions, for reasons that will become clear. For finite extensions, the degree of an extension is multiplicative: If $F \subseteq K \subseteq L$ are fields, then $[L : F] = [L : K][K : F]$, and so $[K : F]$ divides $[L : F]$.

If an element α of K is the root of a polynomial with coefficients in F, then α is said to be **algebraic**. If every element of K is algebraic, then K is an **algebraic extension**. $\mathbb{Q}(i)$ is an algebraic extension, but $\mathbb{Q}(\pi)$ is not, for example. If α is algebraic, then the **minimal polynomial** for α over F is the unique monic irreducible polynomial with coefficients in F that has α as a root. This polynomial may be denoted $m_\alpha(x)$.

Theorem 5. *α is algebraic over F if and only if $F(\alpha)/F$ is a finite extension, and every finite extension is algebraic.*

Proof: If α is algebraic, then the degree of the extension is the degree of $m_\alpha(x)$. If $F(\alpha)/F$ is a finite extension, then the elements $1, \alpha, \alpha^2, ..., \alpha^n$ is a collection of $n + 1$ elements of a dimension n vector space, so they are linearly dependent. Therefore, there is a linear combination $b_0 + b_1\alpha + ... + b_n\alpha^n = 0$, proving that α is the root of a polynomial in F. Finally, if K/F is a finite extension and $\alpha \in K$, then $[F(\alpha) : F] \leq [K : F]$, so $F(\alpha)/F$ is a finite extension and thus α is algebraic. □

The fields this paper are most concerned with are the **splitting fields** of polynomials with coefficients in F. The splitting field of a polynomial $f(x) \in F[x]$ is the minimum field containing both F and every root of $f(x)$. For example, $\mathbb{Q}(\sqrt{2})/\mathbb{Q}$ is the splitting field of $x^2 - 2$, as it contains both roots of the polynomial and $\mathbb{Q}$. $\mathbb{Q}(\sqrt{2}, \sqrt{3})/\mathbb{Q}$ is not the splitting field of $x^2 - 2$, because it has a strict subfield that contains every root of $x^2 - 2$. On the other hand, $\mathbb{Q}$ is not the splitting field of $x^4 - 1$, because it contains some but not all of the roots of $x^4 - 1$. Any two splitting fields of the same polynomial are isomorphic, hence the splitting field will be referred to in the singular.

A polynomial is **separable** if it has no repeated roots, and an extension K/F is a **separable extension** if every element of K is the root of a separable polynomial in $F[x]$. If F is finite, or has characteristic 0, then F is a **perfect** field. Most fields students commonly encounter will be perfect, and it takes some work to find a field that is not. A counterexample is the field of rational functions with coefficients in the finite field $\mathbb{F}_2$, $\mathbb{F}_2(p)$. The significance of this definition are the following results tying separable and irreducible polynomials:

Theorem 6. *Every irreducible polynomial over a perfect field is separable.*

Theorem 7. *Every finite extension of a separable field is separable.*

An isomorphism from a field to itself is an **automorphism**, and as with the deck transformations, the automorphisms of a field K form a group $Aut(K)$ under composition. When K is a field extension, $Aut(K/F)$ is the subgroup of automorphisms that fix F, such that $\sigma(x) = x$ for any $x \in F$. Additionally, these automorphisms preserve the roots of polynomials: if $f(\alpha) = 0$, then $f(\sigma(\alpha)) = 0$. This restricts the action of these automorphisms to permuting the roots of polynomials that are elements of K and not F.

Consider the field extension $\mathbb{Q}(\sqrt{2})/\mathbb{Q}$. $Aut(\mathbb{Q}(\sqrt{2})/\mathbb{Q})$ contains only the identity map and the map $\sqrt{2} \mapsto -\sqrt{2}$. The extension $\mathbb{Q}(\sqrt{2}, \sqrt{3})/\mathbb{Q}$ additionally contains the map $\sqrt{3} \mapsto -\sqrt{3}$, and the map that negates both square roots. It's also possible to have a nontrivial extension with a trivial $Aut(K/F)$: $\mathbb{Q}(\sqrt[3]{2})$ is a degree 3 extension over $\mathbb{Q}$, but $\sqrt[3]{2}$ is the only root of $x^3 - 2$ in the field, so any automorphism of $\mathbb{Q}(\sqrt[3]{2})/\mathbb{Q}$ must fix both $\sqrt[3]{2}$ and all of $\mathbb{Q}$, and thus is constant.

To avoid this issue, one needs to consider fields that have "enough" automorphisms to be interesting. To this end, define a **Galois extension** to be an extension K/F with the property $|Aut(K/F)| = [K : F]$. If K/F is a Galois extension, then denote $Aut(K/F)$ as the **Galois group** $Gal(K/F)$. The earlier comment about automorphisms permuting the roots of polynomials might suggest that Galois extensions are connected to splitting fields. In fact, this connection is very strong:

Theorem 8. *K/F is Galois if and only if K is the splitting field of a separable polynomial over F. If this is the case, then K/F is a separable extension.*

By definition, every automorphism in $Aut(K/F)$ fixes F. However, notice in the earlier examples that these automorphisms can fix larger subfields of K as well. The constant map obviously fixes K/F in its entirety, but the map $\sqrt{2} \mapsto -\sqrt{2}$ fixes $\mathbb{Q}(\sqrt{3}) \subseteq \mathbb{Q}(\sqrt{2}, \sqrt{3})$ as well. Define the **fixed field** of a subgroup $H \subseteq Aut(K/F)$ to be the subfield of K that is fixed by all automorphisms in H. In this case, $\mathbb{Q}(\sqrt{3})$ is the fixed field of the subgroup $\{\mathbb{1}, \sqrt{2} \mapsto -\sqrt{2}\}$

If K/F is a Galois extension, then there is a one-to-one correspondence between the subfields E with $F \subseteq E \subseteq K$ and the subgroups of the Galois group $Gal(K/F)$, given by the association of each subgroup with its fixed field. As with the deck transformations of a covering space, this is an inverse relationship: larger subfields are associated with smaller subgroups. The finiteness of the typical examples in Galois theory actually reveals an even stronger result in this case: if H fixes E, then the degree of K/E is exactly $|H|$, so the diagrams one draws of the field extensions between K and F are precisely the diagrams one draws of the subgroups of $Gal(K/F)$, even if one numbers the edges to denote the degrees of extension. Further, K/E is always Galois with $Gal(K/E) = H$, E

is Galois over F if and only if H is normal in $Gal(K/F)$ (recall the similar result from the covering space material), and if H is normal then $Gal(E/F) \cong G/H$.

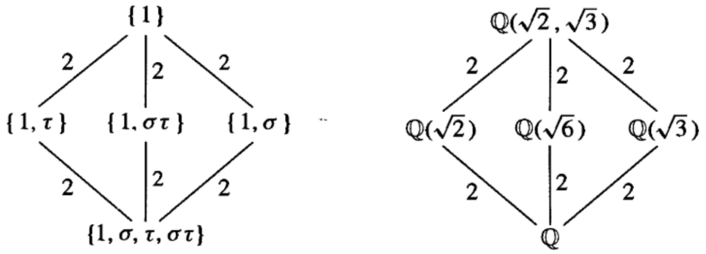

Hasse diagram of the subgroups of $Gal(\mathbb{Q}(\sqrt{2}, \sqrt{3})/\mathbb{Q})$ *and subfields of* $\mathbb{Q}(\sqrt{2}, \sqrt{3})/\mathbb{Q}$*. Illustration by David S. Dummit and Richard M. Foote in their textbook "Abstract Algebra"*

4 The Galois correspondence

All of these results are analogous to covering space theory, despite the ostensible lack in overlap between the very geometric topic of coverings and the focus on solving polynomials found in Galois/field theory. The most apparent difference between the two theories lies, if anything, in the typical examples used in the two subjects. The simplest examples in algebraic topology are composed of circles, planes, lines, and their quotients: with the exception of some quotient spaces like the projective plane, these mostly have either trivial or infinite fundamental groups. Conversely, the simplest examples in field theory are splitting fields of $\mathbb{Q}$, the degree of which are finite. So, this section will construct both a field extension with Galois group $\mathbb{Z}/6\mathbb{Z}$ and a topological space with fundamental group $\mathbb{Z}/6\mathbb{Z}$, to make concrete the analogy.

Finding a field extension is very straightforward: when p is a prime number, $\mathbb{Q}(\zeta_p)/\mathbb{Q}$ is a Galois extension with Galois group $\mathbb{Z}/(p-1)\mathbb{Z}$, where ζ_p is any p^{th} root of unity (a solution to $x^p - 1$ other than 1). So, consider the extension $\mathbb{Q}(\zeta_7)/\mathbb{Q}$. Obviously, the constant subgroup corresponds to the full extension $\mathbb{Q}(\zeta_7)/\mathbb{Q}$, and the full Galois group corresponds to $\mathbb{Q}$. $\mathbb{Z}/3\mathbb{Z}$ corresponds to some degree 2 extension of $\mathbb{Q}$ given by $\mathbb{Q}(\zeta_7 + \zeta_7^2 + \zeta_7^4)$. This is simple enough to compute: letting $\alpha = \zeta_7 + \zeta_7^2 + \zeta_7^4$, note that $1 + \alpha + ... + \alpha^6 = 0$. Then,

$$\begin{aligned}
\alpha^2 &= \zeta_7 + \zeta_7^2 + \zeta_7^4 + 2\zeta_7^3 + 2\zeta_7^5 + 2\zeta_7^6, \\
\alpha^2 + \alpha + 2 &= 2(1 + \zeta_7 + \zeta_7^2 + ... + \zeta_7^6), \\
\alpha^2 + \alpha + 2 &= 0, \\
\implies \alpha &= -(1/2) \pm i\sqrt{7}/2 \quad \text{(Quadratic Formula)}
\end{aligned}$$

So the extension $\mathbb{Q}(\zeta_7 + \zeta_7^2 + \zeta_7^4)/\mathbb{Q} \cong \mathbb{Q}(i\sqrt{7})/\mathbb{Q}$. The subgroup $\mathbb{Z}/2\mathbb{Z}$ corresponds to some cubic extension given by $\mathbb{Q}(\zeta_7 + \zeta_7^6)$, which is substantially trickier to compute.

For a corresponding topological space, consider the lens space $L(r;q)$. The lens space is defined by "gluing" two r-faced solid pyramids (not including the base) at their base, and then further gluing each face of the top pyramid to the face on the bottom pyramid q faces to the right. r and q must be coprime, but this issue can be handwaved by picking $q = 1$. The lens space $L(r;q)$ has fundamental group $\mathbb{Z}/r\mathbb{Z}$, so pick $r = 6$.

It's also worth noting that the covering space has an equivalent construction created by applying a quotient map ("gluing") to the four-dimensional hypersphere S^3. The exact details of this map are hard to visualize, but the fact that such a construction exists helps explain the coverings of the lens space.

So, the chosen space is the lens space $L(6;1)$. The covering space with $G(\tilde{X}) = 0$ is $L(6;1)$, and the covering space with $G(\tilde{X}) = \mathbb{Z}/2\mathbb{Z}$ is $L(3;1)$. The simply-connected covering space is S^3. Intuitively, these covering spaces represent "ungluings" of the lens space: in the case of S^3, the quotient map is totally undone, and so every point in the lens space is covered by the six points on S^3 that were glued to form it. The challenging cover to understand is that with $G(\tilde{X}) = \mathbb{Z}/3\mathbb{Z}$. This is not $L(2;1)$, because there is no pyramid with just two faces and so this lens space does not exist. However, this paper conjectures that this covering space is RP^3, the space obtained by attaching antipodal points of S^3 to each other. This is difficult to visualize even in the three-dimensional case of attaching antipodal points on a conventional sphere, but nevertheless RP^3 has the correct fundamental group and is a quotient of the same space as $L(6;1)$.

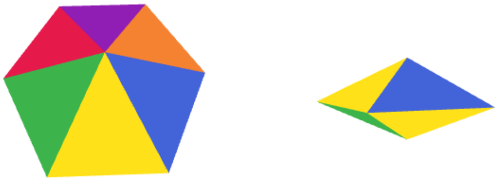

Top and side view of $L(6;1)$. Same-colored faces are considered the same face. Generated on https://vinequai.com/lensspace

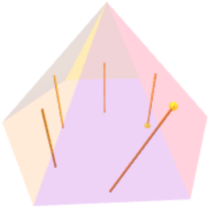

An example of a "5-loop" in $L(6;1)$. If one more segment is added, all six segments of the loop can be homotopically moved to the peak of the pyramid, and then the loop can be deformed to a point. Hence, the fundamental group is $\mathbb{Z}/6\mathbb{Z}$. Generated on https://vinequai.com/lensspace

5 Works Cited

Dummit, David Steven, and Richard M. Foote. *Abstract Algebra*. John Wiley and Sons, 1990.
Hatcher, Allen. *Algebraic Topology*. 2001.
"Lens Space." *Vinequai*. https://vinequai.com/lensspace. Accessed 2024-05-15.