

RELATIVE INVARIANTS IN PERMUTATION GROUPS WITH APPLICATIONS TO GALOIS THEORY

MITCHELL WADAS
CONSULTANT: PROF. HARM DERKSEN

April 21, 2024

ABSTRACT. We describe an algorithmic solution for the computation of Galois groups which employs *relative invariants* for pairs of groups $H \leq G$ – multivariate polynomials whose stabilizer in G is H . We then introduce two new methods for the construction of such invariants (Algorithms 4.7 and 4.13), and give a new construction (Equation 5) for the projective semi-linear group $\mathrm{PTL}(2, 25)$ and its index two subgroup $\mathrm{PGL}(2, 25)$.

CONTENTS

1. Introduction	2
2. Preliminaries	2
2.1. Invariants	2
2.2. Fields & Galois Theory	5
3. Stauduhar’s Method	6
4. Generic Invariants	8
4.1. The Reynolds Operator	8
4.2. The Linear Algebra Method	10
4.3. The Molien Series	11
4.4. Computing a Basis	12
5. Special Invariants	13
5.1. Projective Geometry & Transformations	13
5.2. The Pair $\mathrm{PGL}(2, 25) \triangleleft \mathrm{PTL}(2, 25)$	16
5.3. Constructing an Invariant	17
6. Results & Future Work	19
6.1. Results	19
6.2. Future Work	20
References	21
Glossary	21

1. INTRODUCTION

Determining the Galois group of a polynomial is an old problem. Currently, the best algorithmic solution is an extension of R. P. Stauduhar's method [12]. Stauduhar's method begins with the symmetric group on n letters, denoted S_n , and uses relative invariants (definition 2.5) to traverse the lattice of subgroups down to the Galois group. C. Fieker and J. Klüners [7] developed an algorithmic approach to computing relative invariants of minimal degree for any pair of permutation groups, resulting in a practical, degree independent algorithm.

Unfortunately, there are cases where the relative invariants produced by this generic approach (generic meaning applicable to any pair of permutation groups) are difficult to evaluate, creating a bottleneck in the algorithm. In these cases, we may hope to use the specific group structure to inform the construction of a nicer invariant; invariants constructed for a specific pair of groups, or a family of pairs sharing some common characteristic, are called special invariants. The collection of special methods described by A.-S. Elsenhans in [5] effectively treats a sufficiently broad class of pairs to make Stauduhar's method an efficient and widely used algorithm.

However, there remain some pairs of groups for which none of the special methods described in [5] apply, and one must resort to the generic method of Fieker and Klüners. One example is given by the projective semi-linear group on \mathbb{F}_{25}^2 and its index two subgroup, the projective linear group (see definitions 5.11 and 5.3 for details).

The goal of this report is to describe Stauduhar's method for computing the Galois group of a polynomial, and introduce two new generic algorithms for computing relative invariants (algorithms 4.7 and 4.13). Then we will turn our attention to the particularly challenging case $\text{PGL}(2, 25) \triangleleft \text{PTL}(2, 25)$ and construct a new special invariant (given in equation 5).

2. PRELIMINARIES

This report will study the action of finite groups on polynomial rings with the goal of understanding how polynomials invariant under this action can be constructed, and applied to problems in Galois theory. To facilitate this discussion, we will begin by reviewing some background material, and introducing notation that will be used throughout.

2.1. Invariants.

Definition 2.1. An algebra A over a field K is a set equipped with two binary operations, addition and multiplication, so that $(A, +, \cdot)$ is a ring, and $(A, +)$ is a K -vector space with a scalar multiplication satisfying $k(a \cdot b) = (ka) \cdot b = a \cdot (kb)$ for all $k \in K$ and $a, b \in A$.

Notably, we have defined our algebras to be associative and unital. Some authors define an algebra differently, and would call ours an associative algebra. If $(A, +, \cdot)$ is a commutative ring, then we say that A is a commutative algebra. Unless otherwise mentioned, we will assume that our algebras are commutative.

We will fix a field K of characteristic zero, and denote the set of polynomials in indeterminates x_1, \dots, x_n with coefficients from K by $K[x_1, \dots, x_n]$, abbreviated $K[X]$. The set $K[X]$ forms a commutative K -algebra under the usual addition and multiplication of polynomials. Moreover,

every polynomial in $K[X]$ can be written uniquely as a sum of homogeneous polynomials (definition 6.1), so we can decompose the additive group of $K[X]$ as the direct sum of homogeneous pieces

$$K[X] = \bigoplus_{d=0}^{\infty} K[X]_d$$

where $K[X]_d$ is the subgroup containing all homogeneous polynomials of degree d (and the identity). In addition to the decomposition of the additive group of $K[X]$, we have

$$K[X]_i K[X]_j = \{f \cdot g : f \in K[X]_i, g \in K[X]_j\} \subseteq K[X]_{i+j}.$$

Each homogeneous piece of $K[X]$ has a basis comprised of all monomials of degree d . We can count the number of monomials of degree d in n indeterminates by imagining that we have d ‘powers’ to distribute across the n variables. If we represent the powers by *’s, we can assign exponents to the indeterminates by placing $n - 1$ bars, which divide the d stars into n groups. For example, the monomial $x_1^2 x_3 \in K[x_1, x_2, x_3]$ is specified by $** || *$. Therefore, the total number of monomials is the same as the number of ways to place $n - 1$ bars into $n - 1 + d$ spaces. In other words,

$$\dim K[X]_d = \binom{n + d - 1}{n - 1} = \binom{n + d - 1}{d}.$$

In particular, the dimension of each $K[X]_d$ is finite. We summarize these properties of $K[X]$ with the next definition.

Definition 2.2. An algebra A over a field K is said to be naturally graded if the additive group admits the representation

$$A = \bigoplus_{n \in \mathbb{N}} A_n$$

and the homogeneous pieces satisfy $A_i A_j = \{a \cdot a' : a \in A_i, a' \in A_j\} \subseteq A_{i+j}$. If the dimension of each homogeneous piece as a K -vector space is finite, we say A has finite type.

We now want to consider the action of a finite group G induced by a linear representation $\rho : G \rightarrow \text{GL}_n(K)$ (see definition 6.2 for details about linear representations). The homomorphism ρ defines a G -action on the vector space $V = K^n$ by $\sigma * v = \rho(g)v$ for vectors $v \in V$ and group elements $\sigma \in G$. This induces an action on the dual space V^* (definition 6.3) by $\sigma * l(v) = l(\sigma^{-1} * v)$ for linear functionals $l \in V^*$. We extend this action to polynomials $f \in K[X]$ by defining

$$\sigma * f(v) = f(\sigma^{-1} * v).$$

Lemma 2.3. *The G -action on $K[X]$ induced by linear representation $\rho : G \rightarrow \text{GL}_n(K)$ respects addition, multiplication, and scalar multiplication.*

Proof. Let σ belong to G . For polynomials $f, g \in K[X]$ and a vector $v \in K^n$, we have

$$\sigma * (f + g)(v) = (f + g)(\sigma^{-1} * v) = f(\sigma^{-1} * v) + g(\sigma^{-1} * v) = (\sigma * f + \sigma * g)(v).$$

Similarly, using \cdot to denote polynomial multiplication,

$$\sigma * (f \cdot g)(v) = (f \cdot g)(\sigma^{-1} * v) = f(\sigma^{-1} * v) \cdot g(\sigma^{-1} * v) = (\sigma * f \cdot \sigma * g)(v).$$

Finally, if $k \in K$ is a scalar,

$$\sigma * (kf)(v) = kf(\sigma^{-1} * v) = k(\sigma * f(v)).$$

□

An important example of a linear representation is the defining representation of a permutation group, defined to be the homomorphism $\rho : G \rightarrow \text{GL}_n(K)$ which maps each permutation σ to the matrix which sends the i th basis vector to the $\sigma(i)$ th basis vector. In essence, ρ is defined by $\rho(\sigma)(\mathbf{e}_i) = \mathbf{e}_{\sigma(i)}$. When G is a permutation group of degree n , and ρ is the defining representation, the induced action of G on $K[X]$ is equivalent to

$$\sigma * f(x_1, \dots, x_n) = f(x_{\sigma(1)}, \dots, x_{\sigma(n)}).$$

To see this, let $v \in K^n$ be the sum $v = c_1 \mathbf{e}_1 + \dots + c_n \mathbf{e}_n$, and observe that

$$\sigma * f(v) = f(\sigma^{-1} * v) = f(c_1 \mathbf{e}_{\sigma^{-1}(1)} + \dots + c_n \mathbf{e}_{\sigma^{-1}(n)}).$$

It is clear that the $\sigma(m)$ -th basis vector will be mapped to \mathbf{e}_m , so the sum is equivalent to

$$\sigma^{-1} * v = c_{\sigma(1)} \mathbf{e}_1 + \dots + c_{\sigma(n)} \mathbf{e}_n.$$

Therefore, the action of σ^{-1} on vector $v = (v_1, \dots, v_n)$ is $\sigma^{-1} * v = (v_{\sigma(1)}, \dots, v_{\sigma(n)})$. Returning to the action of σ on the polynomial f we see that G acts on $K[X]$ by permuting the indeterminates. We will focus primarily on this representation for the rest of the report.

Given a group G and a linear representation ρ , we are interested in polynomials f in $K[X]$ which are invariant under the G -action induced by ρ . That is, polynomials $f \in K[X]$ such that

$$\sigma * f = f \quad \text{for all } \sigma \in G.$$

We denote the set of polynomials invariant under this action by $K[X]^G$. Although standard, we admit that this notation is imprecise, since G may have different representations yielding different sets of invariants. However, the context should make it clear what representation is being considered.

Lemma 2.4. *The set of invariants $K[X]^G$ forms a subalgebra of $K[X]$. Moreover, $K[X]^G$ inherits the gradation of $K[X]$ and forms a naturally graded vector space of finite type.*

Proof. Certainly scalars are invariant, in particular, 0 and 1. If f and g are G -invariant then, by linearity, $f + g$ and $-f$ are also invariant. Similarly, the G -action satisfies $\sigma * (fg) = (\sigma * f)(\sigma * g)$. Hence $K[X]^G$ is closed under both operations and the taking of additive inverses, and contains the additive and multiplicative identities. Finally, by virtue of scalars being invariant and $K[X]^G$ being closed under products, a valid scalar multiplication is inherited. Therefore, $K[X]^G$ forms a subalgebra. \square

As mentioned in the introduction, this report will focus on relative invariants – a type of invariant which arises when we consider a subgroup $H \leq G$. If f is a polynomial in $K[X]^H$ such that $\sigma * f \neq f$ for all $\sigma \in G \setminus H$, then we say f is a G -relative H -invariant, or simply, a relative invariant.

Definition 2.5. If $f \in K[X]$ is a polynomial whose stabilizer in G is H , we say that f is a G -relative H -invariant.

Constructing relative invariants for pairs of permutation groups (with the action given by the defining representation) will be the central focus of this report, but to understand their significance in Stauduhar's method, we will need to discuss some Galois theory first.

2.2. Fields & Galois Theory. If L is a field, and $K \subset L$ is a field under the same operations, we say L is an extension field of K , denoted L/K . One way of forming field extensions is to adjoin elements $u_1, \dots, u_k \in L$ to the subfield K by defining the generated field extension $K(u_1, \dots, u_k)$ to be the smallest subfield of L containing K and all u_1, \dots, u_k .

Given a field extension L/K , we may regard L as a K -vector space (since L is an Abelian group under addition and has a valid K multiplication). From this perspective, we define the index, or degree, of the extension to be $[L : K] = \dim_K L$, the dimension of L as a K -vector space. If the index of the extension is finite, we say L/K is a finite extension.

The automorphisms of L form a group under composition, which we denote $\text{Aut}(L)$ – see glossary entry 6.5 for the definition of a field automorphism. We define the automorphism group of the extension L/K to be the set of automorphisms of L which leave K fixed. That is

$$\text{Aut}(L/K) = \{\alpha \in \text{Aut}(L) : \alpha(k) = k \text{ for all } k \in K\}.$$

We call the set of elements which are fixed by every automorphism in $\text{Aut}(L/K)$ the fixed field of $\text{Aut}(L/K)$. Note that the fixed field contains K but is not necessarily equal to K .

A special type of extension is obtained by adjoining the roots $r_1, \dots, r_n \in L$ of a polynomial $p \in K[x]$ to construct $N = K(r_1, \dots, r_n)$, which we call the splitting field of p over K . Regarding N as a K -vector space it is clear that $\{1, r_1, \dots, r_n\}$ forms a spanning set for N . Therefore, the automorphisms in $\text{Aut}(N/K)$ are determined by their action on the roots of p .

We call polynomial of degree n separable if it has n distinct roots in its splitting field. For example, take $q(x) = x^p - x \in \mathbb{F}_p[x]$. The multiplicative group of \mathbb{F}_p has order $p - 1$, so $x^p = x$ for every $x \in \mathbb{F}_p$. Hence, q has p distinct roots in \mathbb{F}_p .

Extensions of the form N/K where N is the splitting field of a separable polynomial in $K[x]$ are called Galois extensions, and they have a number of equivalent properties.

Theorem 2.6. *If N/K is finite, then the following statements are equivalent to N/K is Galois.*

- (1) N is the splitting field of a separable polynomial over K ,
- (2) The fixed field of $\text{Aut}(N/K)$ is K ,
- (3) $|\text{Aut}(N/K)| = [N : K]$.

Proof. See theorem 16.6.4 in Artin’s “Algebra” [1] □

If N/K is a Galois extension, we call $\text{Aut}(N/K)$ the Galois group of the field extension, and we use the notation $\text{Gal}(N/K)$.

Lemma 2.7. *Let $p = a_n x^n + \dots + a_1 x + a_0 \in K[X]$ be a separable polynomial with splitting field N over K . Suppose $R \subset N$ is the set of roots of p , and define $\phi : \text{Gal}(N/K) \rightarrow \text{Sym}(R)$ by $\phi(\gamma)(r) = \gamma(r)$. Then ϕ is an injective group homomorphism.*

Proof. First, we will show that the stated co-domain is correct. Observe that for all $\gamma \in \text{Gal}(N/K)$

$$\gamma(p(x)) = \gamma(a_n x^n + \dots + a_1 x + a_0) = a_n \gamma(x)^n + \dots + a_1 \gamma(x) + a_0 = p(\gamma(x)).$$

Applying this property to a root $r \in R$, we have $0 = \gamma(0) = \gamma(p(r)) = p(\gamma(r))$. Therefore, γ maps roots to roots. Being an automorphism of N , we know γ is injective on the finite subset $R \subset N$, so γ is a bijection.

Now suppose $\phi(\gamma) = \phi(\alpha)$ for some $\gamma, \alpha \in \text{Gal}(N/K)$. Then $\gamma(r) = \alpha(r)$ for all $r \in R$. Since the automorphism in $\text{Gal}(N/K)$ are determined by their action on the roots, we must have $\gamma = \alpha$. Finally, a direct computation shows that ϕ is a group homomorphism:

$$\phi(\gamma \circ \alpha)(r) = \gamma(\alpha(r)) = \phi(\gamma)(\phi(\alpha)(r)) = (\phi(\gamma)\phi(\alpha))(r).$$

□

Given the natural correspondence between $\text{Sym}(R)$ and S_n (obtained by identifying each r_i with its index i), we may regard ϕ as a map from $\text{Gal}(N/K)$ into S_n . Since ϕ is an injective group homomorphism, the image of $\text{Gal}(N/K)$ under ϕ is a subgroup of S_n isomorphic to $\text{Gal}(N/K)$, which we define to be the Galois group of the polynomial p , denoted $\text{Gal}(p)$.

Definition 2.8. The Galois group of a separable polynomial $p \in K[x]$ with splitting field N over K is the image of $\text{Gal}(N/K)$ under ϕ defined in lemma 2.7.

We can now construct the defining representation of $\text{Gal}(p)$ and induce an action on the ring of polynomials $K[x_1, \dots, x_n]$ by

$$\sigma * f(x_1, \dots, x_n) = f(x_{\sigma(1)}, \dots, x_{\sigma(n)})$$

for all $\sigma \in \text{Gal}(p)$ and $f \in K[x_1, \dots, x_n]$. The content of the next lemma relates the behavior of the automorphisms of $\text{Gal}(N/K)$ with the action of $\text{Gal}(p)$ on $K[x_1, \dots, x_n]$.

Lemma 2.9. Let $p \in K[x]$ be a polynomial with splitting field N/K and roots r_1, \dots, r_n . Let f be a polynomial in $K[x_1, \dots, x_n]$, and let γ belong to $\text{Gal}(N/K)$. Using ϕ to denote the representation from lemma 2.7,

$$\gamma(f(r_1, \dots, r_n)) = \phi(\gamma) * f(r_1, \dots, r_n).$$

Proof. Suppose without loss of generality that f is the sum of monomials $f = \sum a(x_1^{i_1} \dots x_n^{i_n})$. Then, if we write $\sigma = \phi(\gamma)$, a direct computation shows that

$$\begin{aligned} \gamma(f(r_1, \dots, r_n)) &= \sum a(\gamma(r_1)^{i_1} \dots \gamma(r_n)^{i_n}) \\ &= \sum a(r_{\sigma(1)}^{i_1} \dots r_{\sigma(n)}^{i_n}) \\ &= f(r_{\sigma(1)}, \dots, r_{\sigma(n)}) \\ &= \sigma * f(r_1, \dots, r_n). \end{aligned}$$

Hence, the action of $\text{Gal}(p)$ on polynomials $f \in K[X]$ evaluated at the roots of p is equivalent to the action of $\text{Gal}(N/K)$ on the field element $f(r_1, \dots, r_n)$. □

3. STAUDUHAR'S METHOD

In this section, we will introduce Stauduhar's method – a widely used algorithm for computing the Galois group of a polynomial, introduced by R. P. Stauduhar in [12] – and discuss the central ideas in a more general setting. If $p \in K[x]$ is a separable polynomial over an infinite field K and we know $\text{Gal}(p) \leq G$ (for example take $G = S_n$), we can determine whether the Galois group is contained in some proper subgroup $H \leq G$ by constructing a relative invariant and applying the next proposition.

Proposition 3.1. *Let $p \in K[x]$ be a separable polynomial with roots r_1, \dots, r_n in its splitting field N . Let $H \leq G$ be subgroups of S_n with G containing $\text{Gal}(p)$, and suppose $f \in K[x_1, \dots, x_n]$ is a G -relative H -invariant such that*

$$\sigma * f(r_1, \dots, r_n) \neq f(r_1, \dots, r_n) \quad \text{for all } \sigma \in G \setminus H. \quad (1)$$

*Then the Galois group $\text{Gal}(p)$ is contained in $\sigma H \sigma^{-1}$ if and only if $\sigma * f(r_1, \dots, r_n)$ lies in K .*

Proof. Let $\lambda = \sigma * f(r_1, \dots, r_n) \in N$. If $\text{Gal}(p)$ is a subgroup of $\sigma H \sigma^{-1}$, then for an automorphism $\gamma \in \text{Gal}(N/K)$, we have $\phi(\gamma) = \sigma h \sigma^{-1}$ for some $h \in H$. Since h stabilizes f , this implies

$$\phi(\gamma) * (\sigma * f(r_1, \dots, r_n)) = (\sigma h \sigma^{-1} \sigma) * f(r_1, \dots, r_n) = \sigma * f(r_1, \dots, r_n).$$

Given the correspondence established in lemma 2.9, this is equivalent to $\gamma(\lambda) = \lambda$. Since N/K is a Galois extension, the fixed field of $\text{Gal}(N/K)$ is K , and $\gamma(\lambda) = \lambda$ implies $\lambda \in K$.

Conversely, suppose $\lambda = \sigma * f(r_1, \dots, r_n)$ lies in K . Then, for all $\gamma \in \text{Gal}(N/K)$, we have $\gamma(\lambda) = \lambda$, which implies $\phi(\gamma) * (\sigma * f(r_1, \dots, r_n)) = \sigma * f(r_1, \dots, r_n)$ by lemma 2.9. Letting σ^{-1} act on both sides of the equation, we have

$$\sigma^{-1} * (\phi(\gamma) * (\sigma * f(r_1, \dots, r_n))) = (\sigma^{-1} \phi(\gamma) \sigma) * f(r_1, \dots, r_n) = f(r_1, \dots, r_n).$$

By assumption (1), this implies $\sigma^{-1} \phi(\gamma) \sigma \in H$. Therefore, $\phi(\gamma) \in \sigma H \sigma^{-1}$, and $\text{Gal}(p)$ is a subgroup of $\sigma H \sigma^{-1}$ as claimed. \square

Hypothesis (1) appears quite strong, but in fact, it can be achieved by applying a Tschirnhausen transformation to the polynomial.

Lemma 3.2. *If K is an infinite field, then in the situation of proposition 3.1 there exist c_0, \dots, c_{n-1} in K such that, for $\beta_i = \sum_{j=0}^{n-1} c_j r_i^j$ we have $\beta_i \neq \beta_j$ for $i \neq j$, and $\sigma * f(\beta_1, \dots, \beta_n) \neq f(\beta_1, \dots, \beta_n)$ for all $\sigma \in G \setminus H$. In fact, the c_i may be chosen from any subset of K with cardinality larger than $\deg(f) \cdot \binom{[G:H]}{2} + \binom{n}{2}$.*

Proof. See Lemma 5.2.2 in Derksen and Kemper's "Computational Invariant Theory" [3]. \square

The above result guarantees that we can always construct a separable polynomial with the same Galois group as p whose roots do satisfy (1). In particular, with β_1, \dots, β_n as in lemma 3.2, we can form the product

$$\hat{p} = \prod_{i=1}^n (x - \beta_i) \in K[x].$$

This amounts to a transformation of p which does not change the Galois group (see [3, Section 5.2] for details). In light of these results, an algorithm for computing the Galois group of a separable polynomial is now clear.

Algorithm 3.3 (Stauduhar's Method). First, compute the subgroup lattice of S_n . Then, beginning with $G = S_n$, select a maximal proper subgroup $H \leq G$ and construct a G -relative H -invariant. Apply proposition 3.1 to traverse the subgroup lattice until a group $G \leq S_n$ is reached such that $\text{Gal}(p) \leq G$ but $\text{Gal}(p)$ is not contained in any maximal subgroup of G , then $\text{Gal}(p) = G$.

It is clear that G -relative H -invariants are central to Stauduhar's method, but good invariants are not always easy to construct. For practical applications, we would like to find invariants which minimize the computational cost – determined by the degree of the invariant and the number of multiplications required for evaluation.

We will investigate the question of constructing invariants in the next two sections, with a particular focus on pairs of groups $H \leq G$ where H is maximal (proper) in G , for these are the pairs which naturally arise in Stauduhar's method.

4. GENERIC INVARIANTS

We will begin this section with a constructive proof for the existence of relative invariants for any pair of permutation groups. The analysis which follows motivates the study of alternative methods.

Lemma 4.1. *Let $H \leq G$ be permutation groups of degree n . There exists a G -relative H -invariant.*

Proof. For any pair of permutation groups, we can construct a standard example of a relative invariant, namely

$$f = \sum_{\sigma \in H} \sigma * \left(\prod_{i=1}^{n-1} x_i^i \right).$$

The action of permutation in H simply permutes the order of the sum, leaving f fixed. The action of $\pi \in G \setminus H$ moves the sum onto the left coset πH , i.e.

$$\pi * f = \sum_{\sigma \in \pi H} \sigma * \left(\prod_{i=1}^{n-1} x_i^i \right).$$

The element $\pi \in \pi H$ maps the monomial $x_1^1 \cdots x_{n-1}^{n-1}$ to $\pi * x_1^1 \cdots x_{n-1}^{n-1}$ which does not appear in the original sum. Hence, f is a G -relative H -invariant. \square

Lemma 4.1 demonstrates that relative invariants exist for every pair of permutation groups, and that we know how to construct them. The invariant which appears in the proof however, is expensive; assuming the powers of each x_i are stored, evaluation of the invariant requires $|H|(n-2)$ multiplications. Being a proper subgroup of S_n , the order of H is bounded above by $n!/2$, meaning the complexity of this invariant is, at worst, factorial in the degree of the polynomial.

This example motivates the development of generic invariants, or generic methods, which produce a relative invariant for any pair of permutation groups, while trying to minimize the computational cost. One commonly used generic invariant is the method of Fieker and Klüners [7], which produces invariants of minimal degree. In this section, we will introduce two new methods for constructing relative invariants of minimal degree. Our methods have the advantage that they are formulated in terms of matrix operations.

4.1. The Reynolds Operator.

Definition 4.2. Let G be a finite group, and K a field whose characteristic does not divide the order of G . Given a linear representation $\rho : G \rightarrow \text{GL}_n(K)$, the Reynolds operator is the map

$$\mathcal{R}_G : K[X] \rightarrow K[X]^G, \quad f \mapsto \frac{1}{|G|} \sum_{\sigma \in G} \sigma * f.$$

The primary application of the Reynolds operator is the construction of invariants (not relative invariants); it is easy to see that the action an element in G simply permutes the sum, leaving $\mathcal{R}_G(f)$ fixed. In this section, we will repurpose the Reynolds operator for the computation of relative invariants, but first we will examine \mathcal{R}_G more closely. In particular, we observe the following.

Lemma 4.3.

- (1) *The Reynolds operator is graded and K -linear.*
- (2) *The Reynolds operator is a projection.*

Proof. The fact that \mathcal{R}_G is graded, meaning it does not change the degree of its argument, is clear from the fact that the group action does not change the degree of polynomials it acts on. The linearity of \mathcal{R}_G follows immediately from the fact that the group action is linear.

If f belongs to $K[X]^H$, then $g = \mathcal{R}_G(f)$ is G -invariant. Applying \mathcal{R}_G to g , we see that the action of $\sigma \in G$ leaves g fixed, so we sum over $|G|$ copies of g . Therefore, after dividing by $|G|$, we get $\mathcal{R}_G(g) = g$. Hence, the Reynolds operator is idempotent. \square

Since $K[X]^G$ is contained in $K[X]^H$ (polynomials invariant under G must also be invariant under subgroup H), the fact that \mathcal{R}_G is idempotent implies that \mathcal{R}_G is surjective. This is an extremely useful fact, and suggests an algorithm for computing $K[X]^G$ which we will develop in subsection 4.4. For now, we want to extend the Reynolds operator by defining the relative Reynolds operator for pairs of groups $H \leq G$.

Definition 4.4. Let G be a finite group with subgroup H , and let K be a field whose characteristic does not divide the index of H in G . Given a linear representation $\rho : G \rightarrow \text{GL}_n(K)$, the relative Reynolds operator is the map

$$\mathcal{R}_{G/H} : K[X]^H \rightarrow K[X]^G, \quad f \mapsto \frac{1}{[G:H]} \sum_{\sigma H \in G/H} \sigma * f.$$

Lemma 4.5.

- (1) *The relative Reynolds operator is independent of the choice of coset representatives.*
- (2) *The Reynolds operator is graded and K -linear.*
- (3) *The Reynolds operator is a projection.*

Proof. The proof of properties 2 and 3 is identical to the proof given for the Reynolds operator in lemma 4.3. For this reason, we will prove only the first claim. Let π belong to the left coset σH . We can write $\pi = \sigma h$ for some $h \in H$, and therefore $\pi * f = \sigma * (h * f)$. Since f is stabilized by $h \in H$, we have $\pi * f = \sigma * f$. \square

Again, since $\mathcal{R}_{G/H}$ is a projection and $K[X]^G \subseteq K[X]^H$, the relative Reynolds operator is surjective. Furthermore, the image and kernel of $\mathcal{R}_{G/H}$ have trivial intersection – suppose f in $K[X]^H$ belongs to $K[X]^G$, then $\mathcal{R}_{G/H}(f) = f$. Unless f is zero, then f is not in the kernel of $\mathcal{R}_{G/H}$. As a consequence, we obtain the vector space decomposition

$$K[X]^H = \ker \mathcal{R}_{G/H} \oplus K[X]^G. \tag{2}$$

This observation leads to the following new result.

Proposition 4.6. *Let $\rho : G \rightarrow \mathrm{GL}_n(K)$ be a linear representation of finite group G , and let H be a maximal (proper) subgroup of G . If f is a non-zero element in the kernel of $\mathcal{R}_{G/H}$, then f is a G -relative H -invariant.*

Proof. The domain of $\mathcal{R}_{G/H}$ is $K[X]^H$, so H is certainly a subgroup of $\mathrm{Stab}_G(f)$. Since the kernel and image of $\mathcal{R}_{G/H,d}$ have trivial intersection and $\mathcal{R}_{G/H}$ is surjective, f cannot be G invariant, thus the stabilizer of f must be a proper subgroup of G . Since H is maximal in G , we conclude $\mathrm{Stab}_G(f) = H$. \square

This result gives us a way to compute every relative invariant for the pair $H \leq G$. To be specific, by finding a basis for the kernel of $\mathcal{R}_{G/H,d}$, we can obtain every relative invariant as a linear combination of the basis vectors and elements from $K[X]^G$. However, $K[X]^H$ is infinite dimensional; to compute the kernel of $\mathcal{R}_{G/H}$, we will restrict $\mathcal{R}_{G/H}$ to the finite dimensional homogeneous piece $K[X]_d$:

$$\mathcal{R}_{G/H,d} : K[X]_d^H \rightarrow K[X]_d^G, \quad f \mapsto \mathcal{R}_{G/H}(f).$$

Then, given a basis for $K[X]_d^H$, we can form the matrix of $\mathcal{R}_{G/H,d}$ and compute the kernel of $\mathcal{R}_{G/H,d}$ by solving a system of linear equations over K . We will discuss an algorithm for computing such a basis in subsection 4.4, but for now we will proceed, taking this for granted.

Before turning these observations into an algorithm, there is one more question that needs to be addressed, namely, how do we choose a degree d for this calculation? Certainly we want to choose d so that the kernel is non-trivial, and we would like to minimize d to make our relative invariant as efficient as possible. In subsection 4.3 we will introduce the Hilbert series, and Molien's formula for calculating the Hilbert series of an invariant algebra. These tools will allow us to choose the smallest d possible. For now, we will take these for granted and describe our first new algorithm for computing relative invariants.

Algorithm 4.7 (Method 1). To compute a relative invariant from a basis for $K[X]_d^H$, construct the matrix of $\mathcal{R}_{G/H,d}$ and compute the kernel by solving a system of linear equations over K . Any non-zero element in the kernel is a relative invariant.

4.2. The Linear Algebra Method.

Definition 4.8. Let $G//H$ denote a set of left coset representatives of G/H , and let $S = (G//H) \setminus \{e\}$ so that S , when combined with H , generates G . We define a map from $K[X]^H$ to the $|S|$ -fold direct sum of $K[X]$ by

$$\psi : K[X]^H \rightarrow \bigoplus_{i=1}^{|S|} K[X], \quad f \mapsto (\sigma * f - f)_{\sigma \in S}. \quad (3)$$

Lemma 4.9.

- (1) *The map ψ is K -linear.*
- (2) *The kernel of ψ is $K[X]^G$.*

Proof. A direct computation shows that for $f, g \in K[X]^G$ and $\lambda \in K$,

$$\psi(f + g) = (\sigma * f + \sigma * g - f - g)_{\sigma \in S} = (\sigma * f - f)_{\sigma \in S} + (\sigma * g - g)_{\sigma \in S} = \psi(f) + \psi(g),$$

$$\psi(\lambda f) = (\lambda(\sigma * f) - \lambda f)_{\sigma \in S} = \lambda(\sigma * f - f)_{\sigma \in S} = \lambda\psi(f).$$

If $\psi(f) = 0$ then $\sigma * f = f$ for all $\sigma \in S$. Since f is also invariant under H , and S together with H generates G , f must be G -invariant. The converse is clear – if f is G -invariant then $\sigma * f - f = 0$ for all $\sigma \in S \subseteq G$ – so we conclude $\psi(f) = 0$ if and only if f is G -invariant. \square

This result implies another algorithm for computing a basis for $K[X]_d^G$; for details, we refer readers to section 3.1 in Derksen and Kemper’s “Computational Invariant Theory” [3]. The next new result demonstrates the utility of ψ as a tool for computing relative invariants.

Proposition 4.10. *Let $\rho : G \rightarrow \text{GL}_n(K)$ be a linear representation of finite group G , and let H be a maximal proper subgroup of G . Then $f \in K[X]^H$ is a relative invariant if and only if $\psi(f) \neq 0$, where ψ is the map defined in equation (3).*

Proof. If f is a relative invariant, then by definition, its G -stabilizer is H . Therefore, the action of $\sigma \in S$ will not leave f fixed, meaning $\sigma * f - f \neq 0$. In essence, $\psi(f)$ is not zero in any component.

If $\psi(f)$ is non-zero in any component, then f is not G -invariant. Of course, H is contained in the stabilizer, but there are no groups between H and G , so we must have $\text{Stab}_G(f) = H$. That is, f is a relative invariant. \square

Lemma 4.11. *If \mathcal{B} is a basis for $K[X]_d^H$, then $\psi(\mathcal{B})$ is a spanning set for $\psi(K[X]_d^H)$.*

Proof. This is an immediate consequence of the fact that ψ is a linear transformation. \square

Corollary 4.12. *If there is a relative invariant of degree d for the pair $H \leq G$, then there will be some element in the basis \mathcal{B} which are not mapped to zero by ψ .*

Proof. If there is a relative invariant in $K[X]_d^H$, then the image of $K[X]_d^H$ will not equal $\{0\}$. Since $\psi(\mathcal{B})$ spans the image, this implies there is a non-zero element in $\psi(\mathcal{B})$, and therefore, a relative invariant in \mathcal{B} . \square

These insights give rise to a new algorithm for computing relative invariants which we describe now.

Algorithm 4.13 (Method 2). To compute a relative invariant from a basis for $K[X]_d^H$, apply ψ , the map defined by (3), to each basis vector. Basis vectors which are not mapped to zero are relative invariants.

Corollary 4.12 guarantees that, if there exists a relative invariant of degree d for the pair $H \leq G$, then the above algorithm will produce a relative invariant.

4.3. The Molien Series. Both of our methods rely on knowledge of the minimal degree of a relative invariant. In this section we will develop the tools to extract this information from our knowledge of the groups. We know

$$K[X]_d^H = \ker \mathcal{R}_{G/H,d} \oplus K[X]_d^G,$$

which implies

$$\dim(\ker \mathcal{R}_{G/H,d}) = \dim K[X]_d^H - \dim K[X]_d^G.$$

Therefore, we need to determine the smallest degree d for which there are more linearly independent H -invariants than linearly independent G -invariants. Fortunately, we can answer this question by computing the Hilbert series of $K[X]^G$ and $K[X]^H$ with the help of Molien's formula.

Definition 4.14. For a naturally graded algebra of finite type M , the Hilbert series of M is the formal power series

$$H(M, t) = \sum_{d=0}^{\infty} \dim(M_d) t^d.$$

The Hilbert series encodes all the information about the dimensions of each homogeneous component of M . It is not obvious from the definition, but as Theodor Molien showed, there is a nice formula for computing the Hilbert series of an invariant algebra directly from our knowledge of the group and its representation.

Proposition 4.15 (Molien's Formula). *For a linear representation $\rho : G \rightarrow \mathrm{GL}_n(K)$ of a finite group G , the Hilbert series of the invariant algebra is given by*

$$H(K[X]^G, t) = \frac{1}{|G|} \sum_{\sigma \in G} \frac{1}{\det(I - \rho(\sigma)t)}.$$

Proof. See Theorem 11.16 in Neusel's "Invariant Theory" [10]. □

Using this formula, we can easily obtain the Hilbert series of $K[X]^G$ and $K[X]^H$ and take the Taylor series about zero to compute as many coefficients as needed. Comparing these coefficients will allow us to determine the minimal degree of a relative invariant. In particular, we can compute

$$H(\ker \mathcal{R}_{G/H}, t) = H(K[X]^H, t) - H(K[X]^G, t).$$

By (2), the first degree with a non-zero coefficient in the Hilbert series of the kernel is the minimal degree of a relative invariant for the pair $H \leq G$.

4.4. Computing a Basis. Both of our methods rely on knowledge of a basis for a homogeneous piece of $K[X]^H$. In this section we will discuss a practical and widely known [3] approach to computing $K[X]_d^H$. First, we observe the following.

Proposition 4.16. *If \mathcal{B} is a basis for $K[X]_d^H$, then $\mathcal{R}_{G/H,d}(\mathcal{B})$ is a spanning set for $K[X]_d^G$.*

Proof. This is an immediate consequence of the fact that $\mathcal{R}_{G/H}$ is linear. □

By taking H to be the trivial group and recognizing that $\mathcal{R}_{G/\{e\}} = \mathcal{R}_G$, we see that applying $\mathcal{R}_{G,d}$ to a basis for $K[X]$ will produce a spanning set for $K[X]_d^G$. However, the dimension of $K[X]_d$ could be very large – as we discussed in the introduction, the dimension of $K[X]_d$ is $\binom{n+d-1}{d}$. Moreover, the Reynolds operator requires us to compute the action of $|G|$ group elements on every basis vector. This can become quite expensive when n (the degree of the polynomial whose Galois group we are interested in) is large.

To address this issue, we have an alternative approach to computing the relative Reynolds operator for a pair of finite groups $H \leq G$. The key observation is that, by forming a chain of groups

$H = W_1 \leq \cdots \leq W_r = G$, we can compute the relative Reynolds operator as the composition $\mathcal{R}_{G/H} = \mathcal{R}_{W_r/W_{r-1}} \circ \cdots \circ \mathcal{R}_{W_2/W_1}$. This has the advantage that each function

$$\mathcal{R}_{W_{i+1}/W_i} : K[X]^{W_i} \rightarrow K[X]^{W_{i+1}}, \quad f \mapsto \frac{1}{[W_{i+1} : W_i]} \sum_{\sigma H \in W_{i+1}/W_i} \sigma * f$$

requires the action of only $[W_{i+1} : W_i]$ group elements. Therefore, the total number of actions required is

$$\sum_{i=1}^{r-1} [W_{i+1} : W_i].$$

In most cases, this is much smaller than the number of actions required to compute in the standard way:

$$\prod_{i=1}^{r-1} [W_{i+1} : W_i] = [G : H].$$

Moreover, if at each step we take the spanning set produced by applying $\mathcal{R}_{W_i/W_{i-1}}$ to a basis of $K[X]^{W_{i-1}}$ and refine it to a basis for $K[X]^{W_i}$, then we need to apply $\mathcal{R}_{W_{i+1}/W_i}$ to only $\dim K[X]^{W_i}$ elements, as opposed to all $\dim K[X]^H$ elements which form a basis for $K[X]^H$. We summarize this well-known [3] procedure with the next algorithm.

Algorithm 4.17. To compute a basis for $K[X]_d^G$ from a basis of $K[X]_d^H$, first form a chain of groups $H = W_1 \leq \cdots \leq W_r = G$. Then, apply \mathcal{R}_{W_2/W_1} to each element in the basis and refine the resulting spanning set to a basis for $K[X]_d^{W_2}$. Repeat $r - 1$ times until a basis for $K[X]_d^G$ is obtained.

5. SPECIAL INVARIANTS

In this section, we will focus our attention on a specific pair of groups which is particularly difficult to treat with existing methods, namely, the projective semi-linear group on the projective line over \mathbb{F}_{25} , and its index two subgroup, the projective linear group. The methods outlined in [5] are not applicable to this particular pair, and one must use a generic invariant requiring 14735 multiplications – the most of any pair tested in [5]. Our goal is to study these groups and construct a special invariant.

5.1. Projective Geometry & Transformations.

Definition 5.1. Given a vector space V over a field K , the projective space $\mathbb{P}(V)$ is the set of one dimensional subspaces of V . If $\dim V = n$, we say $\mathbb{P}(V)$ has projective dimension $\text{pdim } \mathbb{P}(V) = n - 1$.

As in the Euclidean case, we call a projective space of projective dimension 0 a projective point, and a projective space of projective dimension 1 a projective line. If we take $V = K^2$, then the projective line $\mathbb{P}(V)$ is the set of lines through the origin, each of which is a projective point in this space. All but one of these subspaces intersects the horizontal line $x_2 = 1$ at exactly 1 point. In particular, a projective point spanned by $(z, 1)$ intersects $x_2 = 1$ at z on the horizontal axis. In this way, we can uniquely identify (almost) all the projective points in $\mathbb{P}(V)$ with an element in K .

But what about the line spanned by $(1, 0)$? This subspace does not intersect $x_2 = 1$ at any point, but as a projective point approaches the line spanned by $(1, 0)$, the point of intersection with $x_2 = 1$

becomes arbitrarily large. In this way, the line spanned by $(1, 0)$ is playing the role of infinity in the projective line. We will make this more precise with the next definition.

Definition 5.2. For a field K , we define the projectively extended line \widehat{K} to be the set $K \cup \{\infty\}$. The projectively extended line is in 1-1 correspondence with the projective line $\mathbb{P}(K^2)$ with the bijection given by $\text{span}(z, 1) \mapsto z$, and $\text{span}(0, 1) \mapsto \infty$.

This correspondence allows us to identify projective points with elements in the projectively extended line which we will do frequently, particularly in the next two subsections.

Definition 5.3. An automorphism $f : V \rightarrow V$ induces a bijection of $\mathbb{P}(V)$ called a homography by mapping subspaces to their image. The group of homographies induced by the linear group $\text{GL}(V)$ is called the projective linear group, denoted $\text{PGL}(V)$. When $V = \mathbb{F}_q^n$, the notation $\text{PGL}(n, q)$ is used.

Lemma 5.4. *Given two automorphisms $f : V \rightarrow V$ and $g : V \rightarrow V$, the homographies induced by f and g are equivalent if and only if there is some non-zero scalar $\lambda \in K$ such that $f = \lambda g$.*

Proof. See proposition 3.2.1 in Gruenberg and Weir's "Linear Geometry" [9]. □

As a result, we see that $\text{PGL}(V)$ is isomorphic to the linear group $\text{GL}(V)$ modulo its center. Unlike their linear counterparts, which are determined by their action on n distinct vectors, the homographies in $\text{PGL}(V)$ are uniquely determined by their action on $n+1$ distinct projective points, or to be more specific, their action on a *projective frame*.

Definition 5.5. Let $\mathbb{P}(V)$ have projective dimension $n - 1$. A projective frame in $\mathbb{P}(V)$ is an $(n + 1)$ -tuple of distinct projective points.

Lemma 5.6. *Let $\mathbb{P}(V)$ have projective dimension $n - 1$. Given any two projective frames $(A_i)_{1 \leq i \leq n+1}$ and $(B_i)_{1 \leq i \leq n+1}$ in $\mathbb{P}(V)$, there is a unique homography $h : \mathbb{P}(V) \rightarrow \mathbb{P}(V)$ so that $h(A_i) = B_i$ for $1 \leq i \leq n + 1$.*

Proof. See lemma 5.5.4 in Gallier's "Geometric Methods and Applications" [8]. □

We can extend the projective linear group by considering the permutations induced by *semi-linear transformations*.

Definition 5.7. We say a bijection $f : V \rightarrow V$ is a semi-linear transformation with respect to $\alpha \in \text{Aut}(K)$ if

- (1) $f(u + v) = f(u) + f(v)$ for all $u, v \in V$,
- (2) $f(\lambda v) = \alpha(\lambda)f(v)$ for all $v \in V$ and $\lambda \in K$.

Definition 5.8. The group of semi-linear transformations of V is called the semi-linear group, denoted $\Gamma\text{L}(V)$.

An important example can be constructed from an automorphism of the ground field K by applying the automorphism to the coordinates of a vector. More generally, we have the following.

Lemma 5.9. *If we fix a basis \mathcal{B} for V , then we can define an injective group homomorphism $\psi : \text{Aut}(K) \rightarrow \Gamma\text{L}(V)$ by*

$$\psi(\alpha) \left(\sum_{b \in \mathcal{B}} c_b b \right) = \sum_{b \in \mathcal{B}} \alpha(c_b) b.$$

Proof. First we will verify that the stated co-domain is correct. Let $u, v \in V$ and $\lambda \in K$; a direct computation shows that $\psi(\alpha)$ is indeed semi-linear with respect to α :

$$\begin{aligned} \psi(\alpha)(u + v) &= \left(\sum_{b \in \mathcal{B}} (c_b + d_b) b \right) = \sum_{b \in \mathcal{B}} \alpha(c_b + d_b) b = \sum_{b \in \mathcal{B}} \alpha(c_b) b + \alpha(d_b) b = \psi(\alpha)(u) + \psi(\alpha)(v), \\ \psi(\alpha)(\lambda v) &= \psi(\alpha) \left(\sum_{b \in \mathcal{B}} (\lambda c_b) b \right) = \sum_{b \in \mathcal{B}} \alpha(\lambda) \alpha(c_b) b = \alpha(\lambda) \psi(\alpha) \left(\sum_{b \in \mathcal{B}} c_b b \right). \end{aligned}$$

Now we will verify that ψ is a group homomorphism:

$$\psi(\alpha \circ \gamma) \left(\sum_{b \in \mathcal{B}} c_b b \right) = \sum_{b \in \mathcal{B}} \alpha(\gamma(c_b)) b = \psi(\alpha) \left(\sum_{b \in \mathcal{B}} \gamma(c_b) b \right) = (\psi(\alpha) \circ \psi(\gamma)) \left(\sum_{b \in \mathcal{B}} c_b b \right).$$

Finally, to show ψ is injective, suppose $\psi(\alpha) = \psi(\gamma)$. Then,

$$\psi(\alpha) \left(\sum_{b \in \mathcal{B}} c_b b \right) = \psi(\gamma) \left(\sum_{b \in \mathcal{B}} c_b b \right) \implies \sum_{b \in \mathcal{B}} (\alpha(c_b) - \gamma(c_b)) b = 0.$$

The elements of \mathcal{B} are linearly independent, so this implies each coefficient is zero, meaning $\alpha(c_b) = \gamma(c_b)$ for all $b \in \mathcal{B}$. Of course, the coordinates c_b could be any elements in K , so $\alpha = \gamma$. \square

This result demonstrates that $\text{Aut}(K)$ is isomorphic to a subgroup of $\Gamma\text{L}(V)$. If we take \mathcal{B} to be the standard basis for V , then $\psi(\alpha)$ is equivalent to applying α component wise.

Lemma 5.10. *If $f : V \rightarrow V$ is a semi-linear map, then $W \mapsto f(W)$ defines a homography of $\mathbb{P}(V)$.*

Proof. See section 3.5 of Gruenberg and Weir's "Linear Geometry" [9]. \square

Combining lemmas 5.9 and 5.10, we see that the automorphisms of K permute the projective space $\mathbb{P}(V)$ by $W \mapsto \psi(\alpha)(W)$. The homographies induced by $\text{Aut}(K)$ form a subgroup of the symmetric group on $\mathbb{P}(V)$ which we will identify with $\text{Aut}(K)$ for the rest of the report.

Definition 5.11. The projective semi-linear group is the product

$$\text{PFL}(V) = \text{PGL}(V) \text{Aut}(K) = \{g\alpha : g \in \text{PGL}(V), \alpha \in \text{Aut}(K)\}.$$

When $V = \mathbb{F}_q^n$, the notation $\text{PFL}(n, q)$ is used.

Proposition 5.12. *The projective semi-linear group is the semi-direct product of normal subgroup $\text{PGL}(V)$ by $\text{Aut}(K)$. That is $\text{PFL}(V) = \text{PGL}(V) \rtimes \text{Aut}(K)$.*

Proof. We will begin by showing that $N := \text{PGL}(V)$ is normal in $\text{PFL}(V)$. Consider the conjugate $\sigma N \sigma^{-1}$ for a permutation σ induced by the semi-linear map obtained by applying automorphism $\alpha \in \text{Aut}(K)$ component wise. An element $\sigma n \sigma^{-1} \in \sigma N \sigma^{-1}$ permutes the projective space by

$$W \mapsto \{\sigma n \sigma^{-1}(v) : v \in W\}$$

where in this context, $\sigma n \sigma^{-1}$ is a composition of semi-linear maps. Therefore, it suffices to show that $\sigma n \sigma^{-1}$ is linear, and therefore, the permutations in $\sigma N \sigma^{-1}$ belong to N . Let $\lambda \in K$ and $v \in V$ and observe that

$$(\sigma n \sigma^{-1})(\lambda v) = \sigma n(\alpha^{-1}(\lambda) \sigma^{-1}(v)) = \sigma(\alpha^{-1}(\lambda) n \sigma^{-1}(v)) = \lambda \sigma n \sigma^{-1}(v).$$

Hence $\sigma N \sigma^{-1} \subseteq N$. Of course, the conjugation is an automorphism of $\text{PGL}(V)$, so $|\sigma N \sigma^{-1}| = |N|$ and therefore, $\sigma N \sigma^{-1} = N$.

Now we will show that $N \cap \text{Aut}(K)$ is trivial. If $\sigma \in \text{PGL}(V)$ lies in the intersection, then the map which induces σ is simultaneously semi-linear with respect to $\alpha \in \text{Aut}(K)$, and linear. This implies that α is the trivial automorphism, and therefore, σ is the identity permutation.

Finally, we will show that every element $\sigma \in \text{PGL}(V)$ can be written uniquely as the product $\sigma = n\alpha$ for some $n \in N$ and $\alpha \in \text{Aut}(K)$. By definition every element can be constructed this way, so it only remains to show uniqueness. If $\sigma = n\alpha = n'\alpha'$ for some $n, n' \in N$ and $\alpha, \alpha' \in \text{Aut}(K)$, then $n = n'\alpha'\alpha^{-1}$. For this to be true, $n'\alpha'\alpha^{-1}$ must be linear, and therefore, $\alpha\alpha^{-1}$ must be the trivial automorphism. In other words $\alpha' = \alpha$. Then, returning to $n\alpha = n'\alpha$, we cancel α on both sides and obtain $n = n'$. \square

For more information about the semi-direct product, see glossary entry [6.4](#).

5.2. The Pair $\text{PGL}(2, 25) \triangleleft \text{PGL}(2, 25)$. In this subsection, we want to study the particular pair $\text{PGL}(2, 25) \triangleleft \text{PGL}(2, 25)$. We regard these group as permutations of the projective space $\mathbb{P}(\mathbb{F}_{25}^2)$, which has 26 elements.

Lemma 5.6 tells us that the action of $\text{PGL}(2, 25)$ on $\mathbb{P}(\mathbb{F}_{25}^2)$ is sharply 3-transitive, meaning for all 3-tuples (A, B, C) , (D, E, F) of projective points, there is a unique permutation σ in $\text{PGL}(2, 25)$ so that $(\sigma(A), \sigma(B), \sigma(C)) = (D, E, F)$.

With this information we can determine the order of $\text{PGL}(2, 25)$ by counting the number of distinct projective frames in $\mathbb{P}(\mathbb{F}_{25}^2)$. There are 26 elements in $\mathbb{P}(\mathbb{F}_{25}^2)$, so we have $15600 = 26 \cdot 25 \cdot 24$ projective frames, meaning $|\text{PGL}(2, 25)| = 15600$.

Now we will turn our attention to $\text{Aut}(\mathbb{F}_{25})$.

Lemma 5.13. *The automorphisms of \mathbb{F}_{p^k} fix the prime subfield \mathbb{F}_p .*

Proof. The prime subfield \mathbb{F}_p is generated by 1, and every field automorphism must send 1 to 1. \square

Hence, we see that $\text{Aut}(\mathbb{F}_{25}) = \text{Aut}(\mathbb{F}_{25}/\mathbb{F}_5)$.

Lemma 5.14. *\mathbb{F}_{p^k} is the splitting field of $f(x) = x^{p^k} - x$ over \mathbb{F}_p .*

Proof. The multiplicative group of \mathbb{F}_{p^k} has order $p^k - 1$, and therefore by Lagrange, $x^{p^k-1} = 1$ for all $x \neq 0 \in \mathbb{F}_{p^k}$. Hence, every element of \mathbb{F}_{p^k} is a root of the polynomial f . The degree of the polynomial is the order of \mathbb{F}_{p^k} , so the elements of the field are exactly the roots of f . \square

The proof also demonstrates that f has no repeated roots, and therefore is separable. Taking these facts together, we conclude that $\mathbb{F}_{p^k}/\mathbb{F}_p$ is a Galois extension.

Definition 5.15. For a finite field \mathbb{F}_{p^k} , the Frobenius automorphism is the map $\varphi : \mathbb{F}_{p^k} \rightarrow \mathbb{F}_{p^k}$ defined by $\varphi(x) = x^p$.

Lemma 5.16. *The degree of the extension $\mathbb{F}_{p^k}/\mathbb{F}_p$ is k .*

Proof. There are only p coefficients in \mathbb{F}_p to choose from, so \mathbb{F}_{p^k} must have k basis vectors in order to generate p^k distinct elements. \square

Lemma 5.17. *The multiplicative group of a finite field is cyclic.*

Proof. See theorem 6.5.10 in Beachy and Blair's "Abstract Algebra" [2]. \square

Lemma 5.18. *The automorphism group $\text{Aut}(\mathbb{F}_{p^k}/\mathbb{F}_p)$ is cyclic of order k and generated by the Frobenius automorphism.*

Proof. Since \mathbb{F}_{p^k} is a Galois extension of \mathbb{F}_p , we know $|\text{Aut}(\mathbb{F}_{p^k}/\mathbb{F}_p)| = [\mathbb{F}_{p^k} : \mathbb{F}_p] = k$. Hence, it suffices to show that the Frobenius automorphism has order k .

The multiplicative group of \mathbb{F}_{p^k} is cyclic of order $p^k - 1$, so there is some $a \in \mathbb{F}_{p^k}^*$ with order $p^n - 1$. Thus, $a^p, a^{p^2}, \dots, a^{p^k}$ are distinct element, which is to say $\varphi(a), \varphi^2(a), \dots, \varphi^k(a)$ are distinct, which implies the Frobenius automorphism has order k . \square

Given these results, we know $\text{Aut}(\mathbb{F}_{25}/\mathbb{F}_5)$ is cyclic of order two and generated by the map $x \mapsto x^5$. Hence, $\text{PTL}(2, 25)$ consists of $\text{PGL}(2, 25)$ and the left coset $\varphi \text{PGL}(2, 25)$. Therefore, to produce a relative invariant, it suffices to find $f \in K[X]$ which is invariant under $\text{PGL}(2, 25)$, and not invariant under the Frobenius automorphism.

5.3. Constructing an Invariant.

Lemma 5.19. *If x is an element in the prime subfield of \mathbb{F}_{p^k} , then $\varphi(x) = x$. If x is an element in $\mathbb{F}_{p^k} \setminus \mathbb{F}_p$, then $\varphi(x) \neq x$.*

Proof. The first claim is exactly lemma 5.13. We saw in the proof of lemma 5.18 that $x^p - x \in \mathbb{F}_p[x]$ has p distinct roots in \mathbb{F}_{p^k} . If there were some $x \in \mathbb{F}_{p^k} \setminus \mathbb{F}_p$ for which $\varphi(x) = x^p = x$, then x would also be a root of the polynomial, contradicting that a polynomial of degree p has at most p distinct roots. \square

To understand the action of φ on the projective line over \mathbb{F}_{25} , recall that the permutation induced by automorphisms of \mathbb{F}_{25} is obtained by applying the automorphism to each component of a spanning vector.

With this in mind, we see that the subspace spanned by (x, y) is left fixed by φ if and only if x and y belong to the prime subfield \mathbb{F}_5 . If $y = 0$, then y is in the prime subfield, and we can make $x = 1$ belong to the subfield as well. Therefore, φ leaves the infinity point fixed.

For points spanned by (x, y) , we can make $y = 1$ by multiplying the spanning vector by $1/y$. Hence, φ leaves projective points spanned by $(\lambda, 1)$ fixed if and only if λ lies in \mathbb{F}_5 .

Definition 5.20. Given four projective points (A, B, C, D) lying in the same projective line, with A, B, C distinct, the cross ratio is defined as the element $h(D) \in \mathbb{P}(V)$ where $h : \mathbb{P}(V) \rightarrow \mathbb{P}(V)$ is the unique homography mapping (A, B, C) to $(\infty, 0, 1)$. The cross ratio is denoted $[A, B; C, D]$.

Lemma 5.21. *The cross ratio is invariant under the projective linear group, meaning for all σ in $\text{PGL}(V)$*

$$[A, B; C, D] = [\sigma(A), \sigma(B); \sigma(C), \sigma(D)].$$

Proof. See lemma 5.8.2 in Gallier's "Geometric Methods and Applications" [8]. \square

Lemma 5.22. *Let φ denote the Frobenius automorphism, and let $[A, B; C, D] = \lambda$. Then*

$$[\varphi(A), \varphi(B); \varphi(C), \varphi(D)] = \varphi(\lambda).$$

Proof. We begin by applying φ to each side of $[A, B; C, D] = \lambda$ to obtain $\varphi\sigma(D) = \varphi(\lambda)$ where σ is the homography representing the cross ratio. Since $\text{PGL}(2, 25)$ is normal in $\text{PTL}(2, 25)$, we have $\sigma = \varphi^{-1}\pi\varphi$ for some $\pi \in \text{PGL}(2, 25)$, so

$$\varphi\sigma(D) = \pi\varphi(D) = \varphi(\lambda).$$

Now we will show that π is the map which sends $\varphi(A)$ to ∞ , $\varphi(B)$ to 0, and $\varphi(C)$ to 1. We know $\sigma(A) = \infty$ and that φ leaves ∞ fixed; this implies $\varphi\sigma(A) = \pi\varphi(A) = \infty$. Repeating this argument for B and C , we see that π is the unique homography taking $(\varphi(A), \varphi(B), \varphi(C))$ to $(\infty, 0, 1)$. \square

For the remainder of the section, let us identify projective points with their corresponding elements of the projectively extended line $\widehat{\mathbb{F}}_{25}$. In light of the result above, if we choose projective points (A, B, C, D) so that their cross ratio is $\lambda \in \mathbb{F}_{25} \setminus \mathbb{F}_5$, then the action of the Frobenius automorphism will send (A, B, C, D) to some 4-tuple having cross ratio not equal to λ .

Hence, we can construct a relative invariant for the pair by associating each projective point A with an indeterminate x_A , and forming the sum

$$\sum_{[A, B; C, D] = \lambda} x_A x_B^2 x_C^3 x_D^4 \in K[X]. \quad (4)$$

That is, we sum over all four tuples such that $[A, B; C, D] = \lambda$. The action of a permutation in $\text{PGL}(2, 25)$ will send each four tuple to another with the same cross ratio, permuting the sum. On the other hand, φ will send each 4-tuple to a 4-tuple having cross ratio not equal to λ , meaning the associated monomial will not be in the sum.

Now we want to show that the set of tuples having cross ratio equal to λ is equal to the orbit of a tuple (A, B, C, D) such that $[A, B; C, D] = \lambda$ under the action of $\text{PGL}(2, 25)$. To begin, we remark that, by virtue of the cross ratio being a projective invariant, the orbit is clearly a subset the set of tuples with cross ratio λ . To show the reverse inclusion, we will count the number of elements in each set and see that they have the same cardinality.

We know that given any projective frame (A, B, C) , and any fourth point λ , there is exactly 1 point D such that $[A, B; C, D] = \lambda$. Hence, there are $26 \cdot 25 \cdot 24 = 15600$ 4-tuples (A, B, C, D) with cross ratio λ . On the other hand, $\text{PGL}(2, 25)$ is sharply 3-transitive, so for a given (A, B, C) there is a unique element which maps (A, B, C) to every other 3-tuple. Again, there are $26 \cdot 25 \cdot 24 = 15600$ such tuples, so the orbit of (A, B, C) under $\text{PGL}(2, 25)$ has cardinality 15600. Finally, we remark that D is uniquely determined by (A, B, C) , so the orbit of (A, B, C, D) has the same size as the orbit of (A, B, C) .

Therefore, we can rewrite the special invariant (4) as a sum over the $\mathrm{PGL}(2, 25)$ -orbit of the monomial $x_A x_B^2 x_C^3 x_D^4$:

$$\sum_{\sigma \in \mathrm{PGL}(2, 25)} \sigma * (x_A x_B^2 x_C^3 x_D^4). \quad (5)$$

We remark that the cross ratio, though it does depend on the order, is actually invariant under the action of the Klein four group $K_4 = \{e, (AB)(CD), (AC)(BD), (AD)(BC)\}$. Therefore, the action of φ on $\{A, B, C, D\}$ must correspond to some permutation of these elements which is not an element of K_4 .

If we can construct a monomial $m(x_A, x_B, x_C, x_D)$ which is stabilized by K_4 , but not by φ , then m will have a non-trivial stabilizer in $\mathrm{PGL}(2, 25)$, and we can simplify (5) by summing over the smaller orbit of m . Furthermore, if we identify the projective point $[A, B; C, D] \neq \infty$, with its corresponding element in \mathbb{F}_{25} , the orbit of the cross ratio under the action of the symmetric group on $\{A, B, C, D\}$ is described by

$$\begin{aligned} [A, B; C, D] &= \lambda, & [A, B; D, C] &= \frac{1}{\lambda}, \\ [A, C; D, B] &= \frac{1}{1 - \lambda}, & [A, C; B, D] &= 1 - \lambda, \\ [A, D; B, C] &= \frac{\lambda - 1}{\lambda}, & [A, D; C, B] &= \frac{\lambda}{\lambda - 1}. \end{aligned}$$

Therefore, if we choose λ so that $\lambda = (1 - \lambda)^{-1} = \lambda^{-1}(\lambda - 1)$, then the cross ratio will be invariant under the even permutations of $\{A, B, C, D\}$. It turns out that two such elements exists in $\mathbb{F}_{25} \setminus \mathbb{F}_5$; if we regard \mathbb{F}_{25} as the quotient ring $\mathbb{F}_5[x]/\langle x^2 - 2 \rangle$ so that $x^2 \equiv 2$, these two elements are $2x + 3$ and $3x + 3$, and they are related by the fact that $\varphi(2x + 3) = 3x + 3$.

Given these symmetries, we see that, for carefully chosen λ , the stabilizer of $[A, B; C, D] = \lambda$ will have order 12. By choosing a monomial (or polynomial) with the same stabilizer in $\mathrm{PGL}(2, 25)$, we can shrink the orbit to $15600/12 = 1300$ elements, greatly simplifying the orbit sum (5).

We will discuss our efforts in this direction in the next section.

6. RESULTS & FUTURE WORK

6.1. Results. To simplify the notation, let $G = \mathrm{PTL}(2, 25)$ and $H = \mathrm{PGL}(2, 25)$. To begin applying our algorithms, we computed the Hilbert series of the invariant algebras $K[X]^G$ and $K[X]^H$ by applying Molien's formula and found that

$$\begin{aligned} H(K[X]^G, t) &= 1 + t + 2t^2 + 3t^3 + 8t^4 + \dots, \\ H(K[X]^H, t) &= 1 + t + 2t^2 + 3t^3 + 9t^4 + \dots. \end{aligned}$$

From these results, we concluded that the minimal degree of a relative invariant is four (which agrees with the method of Fieker and Klüners), and that the kernel of $\mathcal{R}_{G/H, 4}$ is one-dimensional.

With our knowledge of the minimal degree of a relative invariant for the pair, we computed a basis for $K[X]^H$ by applying algorithm 4.17. We then applied algorithm 4.7 (method 1) to the basis to produce a relative invariant for our pair. The resulting invariant requires 23400 multiplications – a few thousand more than the invariant produced by the benchmark generic method.

With a basis of $K[X]_4^H$ computed, we applied algorithm 4.13 (method 2) to determine which of the basis elements is a relative invariant. This procedure yielded an invariant requiring 11700 multiplications, improving upon the benchmark.

Finally, we consider the special invariant described in section 5. The order of $\mathrm{PGL}(2, 25)$ is 15600, and assuming the powers of each x_i are stored, we can evaluate each monomial with 3 products, giving a total cost of $15600 \cdot 3 = 46800$ multiplications.

We summarize the results obtained from our new methods in the table below. Here “benchmark” refers to the generic method of Fieker and Klüners [7].

Method	Degree	Products
Benchmark	4	14675
Method 1	4	23400
Method 2	4	11700
Special Method	10	46800

TABLE 1. Computational Cost of Invariants.

6.2. Future Work. We introduced two new algorithmic approaches (algorithms 4.7 and 4.13) for computing relative invariants in pairs of permutation groups. In order to evaluate them, we would like to test our generic methods on other pairs of groups to obtain a more comprehensive comparison to the benchmark. An analysis of the computational complexity of the algorithms (not the invariants produced) is also desired.

In addition to this work, the new special invariant (5) we constructed in section 5 may be simplified by choosing points A, B, C, D so that their cross ratio is either $2x + 3$ or $3x + 3$, and choosing a polynomial $p(x_A, x_B, x_C, x_D)$ which is invariant under the even permutations of $\{A, B, C, D\}$.

A canonical polynomial with stabilizer A_4 in S_4 is the Vandermonde determinant, defined by $V(x_1, x_2, x_3, x_4) = \prod_{i < j} (x_i - x_j)$. We might suspect that we could form the orbit sum below to obtain a relative invariant:

$$\sum_{\sigma H \in \mathrm{PGL}(2, 25) / \mathrm{Stab}(V)} \sigma * V(x_A, x_B, x_C, x_D).$$

However, calculations performed in SageMath [11] suggest that this sum works out to be zero. Since scalars are invariant under the action of the full symmetric group, the effect of the Frobenius automorphism will leave this sum fixed, meaning it is unfortunately not a relative invariant for this pair. We will continue to investigate this sum, and alternatives, to try to reduce the complexity of our new special invariant.

Finally, A.-S. Elsenhans recently [4] identified a new class of polynomials which are unusually difficult to treat with existing methods, namely, polynomials with primitive affine Galois groups. We would like to test our algorithms on these polynomials to see if they fare any better than existing methods.

REFERENCES

- [1] Michael Artin. *Algebra*. Pearson Education, 2011, pp. xv+543. ISBN: 978-0-132-41377-0.
- [2] John A. Beachy and William D. Blair. *Abstract algebra*. Waveland Press, Inc., Long Grove, IL, 2004, pp. vii+531. ISBN: 978-1-4786-3869-8.
- [3] Harm Derksen and Gregor Kemper. *Computational invariant theory*. Vol. 130. Encyclopedia of Mathematical Sciences. Springer, Heidelberg, 2015, pp. xxii+366. ISBN: 978-3-642-07796-8.
- [4] Andreas-Stephan Elsenhans. “Computing Galois groups”. In: *International Journal of Group Theory* 13.3 (2024), pp. 241–250.
- [5] Andreas-Stephan Elsenhans. “Improved methods for the construction of relative invariants for permutation groups”. In: *Journal of Symbolic Computation* 79 (2017), pp. 211–231.
- [6] Andreas-Stephan Elsenhans. “Invariants for the computation of intransitive and transitive Galois groups”. In: *Journal of Symbolic Computation* 47.3 (2012), pp. 315–326.
- [7] Claus Fieker and Jürgen Klüners. “Computation of Galois groups of rational polynomials”. In: *LMS Journal of Computation and Mathematics* 17.1 (2014), pp. 141–158.
- [8] Jean Gallier. *Geometric methods and applications*. Vol. 38. Texts in Applied Mathematics. Springer, New York, NY, 2011, pp. xxviii+680. ISBN: 978-1-4419-9960-3.
- [9] K. W. Gruenberg and A. J. Weir. *Linear geometry*. Vol. 49. Graduate Texts in Mathematics. Springer, New York, NY, 1977, pp. x+199. ISBN: 978-0-387-90227-2.
- [10] Mara D. Neusel. *Invariant Theory*. Vol. 36. Student Mathematical Library. American Mathematical Society, Providence, RI, 2015, pp. viii+314. ISBN: 978-0-8218-4132-7.
- [11] The Sage Developers. *SageMath, the Sage Mathematics Software System (Version 2.3.0)*. <https://www.sagemath.org>. 2024.
- [12] Richard P. Stauduhar. “The determination of Galois groups”. In: *Mathematics of Computation* 27 (1973), pp. 981–996.

GLOSSARY

Definition 6.1 (Homogeneous Polynomial). A polynomial is called homogeneous if it is the sum of monomials of the same degree.

Definition 6.2 (Linear representation). A linear representation of a group G is a group homomorphism from G to the linear group $\mathrm{GL}_n(K)$, where K is a field. The number n is called the degree of the representation.

Definition 6.3 (Dual Space). Given a vector space V over K , the dual space V^* is the vector space of all linear maps $V \rightarrow K$.

Definition 6.4 (Semi-Direct Product). If G is a group with subgroups N and K such that N is normal in G , the intersection $N \cap K$ is trivial, and $G = NK$, then we say G is the semi-direct product of N by K , denoted $G = N \rtimes K$.

Definition 6.5 (Field Automorphism). Let K be a field. A map $\alpha : K \rightarrow K$ is an automorphism of K if it is bijective and satisfies the following for all a, b in K :

- (1) $\alpha(a + b) = \alpha(a) + \alpha(b)$ and,
- (2) $\alpha(ab) = \alpha(a)\alpha(b)$.

Definition 6.6 (Stabilizer). Let G be a group acting a set X . The stabilizer of an element $x \in X$ is the set of elements in G which leave x fixed. The stabilizer is a subgroup of G , and it is denoted $\text{Stab}_G(x)$.