Counting Number Fields by Discriminant

Evan Dummit | February 9, 2017 | QVNTS

- Among the most fundamental objects of study in number theory are algebraic number fields and extensions of number fields. The most basic such question is: how many number fields (with some particular set of properties) are there?
- A natural way to order number fields of a given degree is by absolute discriminant. For a relative extension L/K, there is an analogous object known as the relative discriminant $D_{L/K}$, which (although it is no longer an integer but rather an ideal in the ring of integers $\mathcal{O}_K$ of K) has essentially the same properties as the absolute discriminant D_L. One reason this might be a good idea is that the number of number fields of a given degree whose (absolute) discriminant is less than X is finite, as originally shown by Minkowski.
- The goal of this talk is to discuss upper bounds on the number of extensions of a fixed degree, bounded (relative) discriminant, and specified Galois closure. (The main theorems in this talk appeared in my Ph.D. thesis, but a lot of followup work is still ongoing.)

1 Notation and Background

- To introduce some notation, let K be a number field and L/K be an extension of degree n. We will let $\mathcal{O}_L$ and $\mathcal{O}_K$ be the rings of integers, and D_L and D_K be the absolute discriminants of L and K respectively, and $D_{L/K}$ be the relative discriminant ideal in $\mathcal{O}_K$. We also take $\mathrm{Nm}_{K/\mathbb{Q}}$ to be the absolute norm on ideals or elements (as appropriate).
  - Essentially none of the flavor is lost by assuming $K = \mathbb{Q}$, so feel free to make this assumption at any time.
- We will employ the standard notations $f(X) \sim g(X)$ to mean $\lim_{X\to\infty} \frac{g(X)}{f(X)} = 1$, and $f(X) \ll g(X)$ to mean that $f(x) < c\, g(X)$ for some constant $c > 0$ and X sufficiently large (where c may depend on other parameters such as n and ϵ that will be clear from the context). The group G will also always be a finite, transitive subgroup of S_n (and is to be interpreted as a Galois group equipped with an embedding).

1.1 Counting extensions of fixed degree

- Definition: For a fixed K and n, we define $N_{K,n}(X)$ to be the number of number fields L (up to K-isomorphism) with extension degree $[L : K] = n$ and absolute discriminant norm $\mathrm{Nm}_{K/\mathbb{Q}}(\mathcal{D}_{L/K}) < X$.
- A folk conjecture, sometimes attributed to Linnik, says that

$$N_{K,n}(X) \sim C_{K,n} X$$

  for fixed n and as $X \to \infty$, for some positive constant $C_{K,n}$ depending on K and n.
  - Even for $K = \mathbb{Q}$, the best known results for large n are far away from this conjectured result. Only in some low-degree cases ($n \leq 5$) is this conjecture proven. (I will give a list of known results in a moment.)
- The first upper bound for general n was proven by Schmidt in the 1980s:
- Theorem ([Schmidt]) For all n and all base fields K,

$$N_{K,n}(X) \ll X^{(n+2)/4}.$$

- The best upper bound for general n was established by Ellenberg and Venkatesh in 2006:
- Theorem ([Ellenberg, Venkatesh]) For all $n > 2$ and all base fields K,

$$N_{K,n}(X) \ll (X\, D_{K/\mathbb{Q}}^n\, A_n^{[K:\mathbb{Q}]})^{\exp(C\sqrt{\log n})},$$

  where A_n is a constant depending only on n and C is an absolute constant.
- For sufficiently large n (roughly on the order of $n = 20$), the result of Ellenberg-Venkatesh improves that of Schmidt.
  - The approach of Schmidt is as follows: first, construct a lattice attached to the ring of integers $\mathcal{O}_L$ of L, and use Minkowski's lattice theorem to obtain an element $\alpha \in \mathcal{O}_L$ whose archimedean norms are small (in terms of X). This gives bounds on the coefficients of the minimal polynomial of α; counting the number of possibilities for α yields the upper bound on the number of possible extensions L/K.
  - The approach of Ellenberg-Venkatesh, in brief, modifies this technique to instead count linearly-independent r-tuples of elements of $\mathcal{O}_L$, using properties of the invariant theory of products of symmetric groups and by rephrasing the problem into one about counting integral points on a scheme which is a generically-finite cover of affine space.

1.2 Counting extensions with fixed degree and particular Galois group

- We may refine the basic counting problem by restricting our attention to extensions L/K whose Galois closure $\hat{L}/K$ has Galois group isomorphic to a particular finite permutation group G.
- For fixed K and n, and a transitive permutation group $G \hookrightarrow S_n$ with a given embedding into S_n, we define $N_{K,n}(X; G)$ to be the number of number fields L (up to K-isomorphism) such that
  1. The degree $[L : K] = n$,
  2. The absolute norm of the relative discriminant $\mathrm{Nm}_{K/\mathbb{Q}}(\mathcal{D}_{L/K})$ is less than X, and
  3. The action of the Galois group of the Galois closure of L/K on the complex embeddings of L is permutation-isomorphic to G.
- For shorthand, we refer to extensions satisfying these conditions as G-extensions. We will also abuse terminology and refer to G as the "Galois group" of the extension L/K, despite the fact that this extension is not typically Galois.
- A series of conjectures of Malle give expected growth rates for $N_{K,n}(X; G)$ depending on the group G.
  - Explicitly, for G a transitive subgroup acting on $\Omega = \{1, 2, \ldots, n\}$, and for g in G, define the index of an element

$$\mathrm{ind}(g) = n - [\text{number of orbits of } g \text{ on } \Omega],$$

    which is also equal to the sum of the lengths of all the cycles, minus the number of cycles, in the cycle decomposition of g in S_n.
  - Next define the index of G to be

$$\mathrm{ind}(G) = \min\{\mathrm{ind}(g) \,:\, 1 \neq g \in G\}.$$

  - We also set

$$a(G) = 1/\mathrm{ind}(G).$$

    Note that the index of a transposition is equal to 1, and (since an element with index 1 has $n-1$ orbits) the transpositions are the only elements of index 1.
  - The absolute Galois group of K acts on the conjugacy classes of G via the action on $\bar{\mathbb{Q}}$-characters of G. We define the orbits (of that action) to be the "K-conjugacy classes" of G. Since all elements in a K-conjugacy class have the same index, we define the index of a conjugacy class to be the index of any element in that class.

- With the terminology defined above, the strong form of Malle's conjecture is as follows:
- Conjecture (Malle, strong form) There exists a constant $c(k,G) > 0$ such that

$$N_{K,n}(X;G) \sim c(K,G) \cdot X^{a(G)} \cdot \log(X)^{b(K,G)-1},$$

where $a(G) = \dfrac{1}{\mathrm{ind}(G)}$ and $b(K,G) = \#\{C : C \text{ a } K\text{-conjugacy class of minimal index } \mathrm{ind}(G)\}$.

  - We would expect by Linnik's conjecture that for any group G, the asymptotics should not exceed X^1, and indeed if $a(G) = 1$ then $b(K,G)$ is also 1.
  - The strong form of Malle's conjecture holds for all abelian groups; this is a result of Wright.
  - However, Klüners has constructed a counterexample to the $\log(X)$ part of the conjecture for the non-abelian group $G = C_3 \wr C_2$ of order 18 embedded in S_6. (Klüners also notes that this is not a unique example, and that all groups of the form $C_p \wr C_2$ yield counterexamples to Malle's conjecture as formulated above.)
  - The ultimate difficulty is the potential existence of an intermediate cyclotomic subfield inside the extension: in this case, $\mathbb{Q}(\zeta_3)$ (or $\mathbb{Q}(\zeta_p)$ in the general family).
  - There is a recent refinement of the exponent of the log-term in Malle's conjecture over function fields, due to Turkelli, which appears to avoid all of the known counterexamples. Turkelli's refinement is motivated by counting points on components of non-connected Hurwitz schemes.
  - The question of counting points on connected Hurwitz schemes was related to counting extensions of function fields in a paper of Ellenberg-Venkatesh, and their heuristics (subject to some assumptions) aligned with Malle's. Turkelli extended their arguments to cover non-connected Hurwitz schemes, and the difference in the results compared to those of Ellenberg-Venkatesh suggested a modification to Malle's conjecture.

- It is generally believed that the power of X in Malle's conjecture is essentially correct. Explicitly:
- Conjecture: (Malle, weak form) For any $\epsilon > 0$ and any number field K, $X^{a(G)} \ll N_{K,n}(X;G) \ll X^{a(G)+\epsilon}$, where $a(G) = \dfrac{1}{\mathrm{ind}(G)}$.

  - If true, Malle's conjecture, even when we restrict to the "weak form" that only considers the power of X, and only for extensions of $\mathbb{Q}$, would for example imply that every finite group is a Galois group over $\mathbb{Q}$. As such, even this weak version (let alone the full version) is naturally considered to be entirely out of reach of current methods.

- An upper bound at least as strong as the weak form of Malle's conjecture is known to hold in the following cases:

  1. For any abelian group (Wright), with the asymptotic constants (in principle).
  2. For any nilpotent group (Kluners-Malle). For a nilpotent group in its regular representation, the lower bound is also known.
  3. For S_3 (Davenport-Heilbronn, Datskovsky-Wright), and the asymptotic constants are also known. In fact, in this case there is a second main term, and its asymptotic constant is also known (Bhargava-Shankar-Tsimerman, Taniguchi-Thorne).
  4. For $D_{2\cdot 4}$ and S_4 (in principle over general K) (Baily, Bhargava, Cohen-Diaz y Diaz-Olivier). The asymptotic constants are also known.
  5. For S_5 (in principle over general K) (Kable-Yukie, Bhargava), as well as the asymptotic constant.
  6. For degree-6 S_3 extensions (Bhargava-Wood), as well as the asymptotic constant.
  7. Under mild restrictions, for wreath products of the form $C_2 \wr H$ where H is nilpotent (Kluners).

- Also of significant interest is the inverse question (though I won't talk much about it), about giving estimates on *lower* bounds.

- For the first question, for certain n one can use known results to get easy lower bounds of the correct exponent in X; e.g., if n is even one can count quadratic extensions of any field E with $[E:K] = n/2$, and this already gives the correct exponent of X.
- In order to avoid such trivial cases we would want to impose conditions on the Galois group, but then the lower bounds (for general G) are quite hard: after all, if one could merely show a positive lower bound for $N_{\mathbb{Q},n}(X;G)$ for all G, one would have solved the inverse Galois problem!

- We will remark here that it is very important that the fields are ordered by discriminant; using other orderings can produce very different results.
  - For example, consider the case of degree-4 extensions and suppose instead we wanted to count polynomials by the maximum height of their coefficients.
  - If we let a_i for $1 \leq i \leq n$ be indeterminates, then the polynomial $p(x) = x^n + a_{n-1}x^{n-1} + \cdots + a_0 \in K(a_1, \cdots, a_n)$ has Galois group S_n over $K(a_1, \ldots, a_n)$. Hilbert's Irreducibility Theorem then implies that almost all specializations (when ordered by the coefficient height) of this polynomial still have Galois group S_n.
  - However, the results of Cohen et al. collectively show that, when ordered by discriminant, a positive proportion (roughly 17%) of extensions of degree 4 have an associated Galois group isomorphic to the dihedral group $D_{2\cdot 4}$: the difference is entirely caused by ordering the fields by discriminant.
  - Malle's conjectures, moreover, indicate that the non-S_n extensions should have a positive density for any composite n, but should have zero density for prime n, though this is not known to be true for any $n > 5$. In general (again per Malle), the Galois groups which are expected to occur with positive density are precisely those Galois groups containing a transposition.

2 Counting G-Extensions by Discriminant

- My first goal is to discuss a sharpening of the upper-bound results of Ellenberg-Venkatesh for arbitrary G-extensions. Let me give the statement of my theorem now, and then talk about the ingredients:

- Theorem ([D.]): Let $n \geq 2$, let K be any number field, and let G be a proper transitive subgroup of S_n. Also, let t be such that if G' is the intersection of any point stabilizer in S_n with G, then any subgroup of G properly containing G' has index at least t. Then for any $\epsilon > 0$,

$$N_{K,n}(X;G) \ll X^{\frac{1}{2(n-t)}\left[\sum_{i=1}^{n-1}\deg(f_{i+1}) - \frac{1}{[K:\mathbb{Q}]}\right]+\epsilon},$$

  where the f_i for $1 \leq i \leq n$ are a set of primary invariants for G, whose degrees (in particular) satisfy $\deg(f_i) \leq i$.
  - The parameter t is necessary because there could potentially be proper subfields between L and K. In the event that there are no such subfields (i.e., if G is a primitive permutation group, e.g., if G is simple) then the result is strictly stronger than Schmidt's bound.
  - If $G = S_n$, then the result still holds without the $-1/[K:\mathbb{Q}] + \epsilon$ part, and the result becomes exactly Schmidt's bound – the exponent is $\frac{1}{2(n-1)} \cdot \sum_{i=1}^{n-1}\deg(f_{i+1}) = \frac{(2+3+\cdots+n)}{2(n-1)} = \frac{n(n+1)/2-1}{2(n-1)} = \frac{n+2}{4}$.

- Let me also give an example using my favorite group, the simple group of order 168:

- Corollary: Let $G = PSL_2(\mathbb{F}_7) \cong GL_3(\mathbb{F}_2)$, appearing in its 7-dimensional (primitive) permutation representation. (An explicit embedding is $G = \langle (1\,2\,3\,4\,5\,6\,7),\ (1\,2)(3\,6)\rangle$.) Then for any $\epsilon > 0$,

$$N_{\mathbb{Q},7}(X;G) \ll X^{11/6+\epsilon}.$$

  For comparison, Schmidt's bound (for general septic extensions) gives an upper bound of $X^{9/4}$, and the Ellenberg-Venkatesh bound is weaker.

2.1 Primary Invariants

- Let G be a finite group and $\rho : G \to GL_n(\mathbb{C})$ be a (faithful) complex representation, and let G act on $\mathbb{C}[x_1, \cdots, x_n]$ via ρ.
- If $f_1, \cdots, f_n$ are algebraically independent, homogeneous G-invariant elements of $\mathbb{C}[x_1, \cdots, x_n]$ with the property that $\mathbb{C}[x_1, \cdots, x_n]^G$, the ring of G-invariant polynomials, is a finitely-generated module over $\mathbb{C}[f_1, \cdots, f_n]$, we say these polynomials f_i are a set of primary invariants for G.
    - The Noether normalization lemma implies that such polynomials exist; that there are n of them follows from comparing transcendence degrees.
    - The primary invariants are not unique: one can (for example) take linear combinations or powers of the f_i and still retain the finite-generation property.
    - When we speak of primary invariants, we generally mean a set of primary invariants which are homogeneous and of minimal degree, and we will arrange them in nondecreasing order of degree.
- Denote $A = \mathbb{C}[f_1, \cdots, f_n]$, and $R = \mathbb{C}[x_1, \cdots, x_n]^G$.
- The theorem of Hochster-Roberts implies that R is a Cohen-Macaulay ring and, moreover, that there exist homogeneous G-invariant polynomials g_1, $g_2, \cdots, g_k$ with $g_1 = 1$ such that $R = A \cdot g_1 + \cdots + A \cdot g_k$.
    - These polynomials g_i are called secondary invariants of G and will depend intrinsically on the choice of primary invariants, and are not uniquely determined even for a fixed set of primary invariants.
- Example: Let $G = S_n$ and ρ be the natural representation of G acting by index permutation on $\mathbb{C}[x_1, \cdots, x_n]$. It is easy to see that the elementary symmetric polynomials are invariants under the action of G on $\mathbb{C}[x_1, \cdots, x_n]$, and that they are algebraically independent: thus, they form a set of primary invariants for G.
    - In fact, for any subgroup of S_n, the elementary symmetric polynomials form a set of (possibly non-minimal-degree) primary invariants.
    - Hence, by a simple replacement argument, for any permutation representation ρ of degree n, there exists a set of primary invariants of ρ such that $\deg(f_i) \leq i$ for each $1 \leq i \leq n$.
- Of course, the elementary symmetric polynomials do not necessarily have minimal degree.
- Example: Let G be the simple group of order 168, appearing in its 7-dimensional (primitive) permutation representation. A computation with MAGMA shows that primary invariants can be chosen as

$$
\begin{aligned}
f_1 &= x_1 + x_2 + x_3 + x_4 + x_5 + x_6 + x_7 \\
f_2 &= x_1^2 + x_2^2 + x_3^2 + x_4^2 + x_5^2 + x_6^2 + x_7^2 \\
f_3 &= x_1^3 + x_2^3 + x_3^3 + x_4^3 + x_5^3 + x_6^3 + x_7^3 \\
f_4 &= x_1x_2x_3 + x_1x_2x_5 + x_1x_2x_6 + x_1x_2x_7 + x_1x_3x_4 + x_1x_3x_6 + x_1x_3x_7 + x_1x_4x_5 \\
&\quad + x_1x_4x_6 + x_1x_4x_7 + x_1x_5x_6 + x_1x_5x_7 + x_2x_3x_4 + x_2x_3x_5 + x_2x_3x_7 + x_2x_4x_5 \\
&\quad + x_2x_4x_6 + x_2x_4x_7 + x_2x_5x_6 + x_2x_6x_7 + x_3x_4x_5 + x_3x_4x_6 + x_3x_5x_6 + x_3x_5x_7 \\
&\quad + x_3x_6x_7 + x_4x_5x_7 + x_4x_6x_7 + x_5x_6x_7 \\
f_5 &= x_1^4 + x_2^4 + x_3^4 + x_4^4 + x_5^4 + x_6^4 + x_7^4 \\
f_6 &= x_1^2x_2x_3 + x_1^2x_2x_5 + x_1^2x_2x_6 + x_1^2x_2x_7 + x_1^2x_3x_4 + x_1^2x_3x_6 + x_1^2x_3x_7 + x_1^2x_4x_5 \\
&\quad + x_1^2x_4x_6 + x_1^2x_4x_7 + x_1^2x_5x_6 + x_1^2x_5x_7 + x_1x_2^2x_3 + x_1x_2^2x_5 + x_1x_2^2x_6 + x_1x_2^2x_7 \\
&\quad + x_1x_2x_3^2 + x_1x_2x_5^2 + x_1x_2x_6^2 + x_1x_2x_7^2 + x_1x_3^2x_4 + x_1x_3^2x_6 + x_1x_3^2x_7 + x_1x_3x_4^2 \\
&\quad + x_1x_3x_6^2 + x_1x_3x_7^2 + x_1x_4^2x_5 + x_1x_4^2x_6 + x_1x_4^2x_7 + x_1x_4x_5^2 + x_1x_4x_6^2 + x_1x_4x_7^2 \\
&\quad + x_1x_5^2x_6 + x_1x_5^2x_7 + x_1x_5x_6^2 + x_1x_5x_7^2 + x_2^2x_3x_4 + x_2^2x_3x_5 + x_2^2x_3x_7 + x_2^2x_4x_5 \\
&\quad + x_2^2x_4x_6 + x_2^2x_4x_7 + x_2^2x_5x_6 + x_2^2x_6x_7 + x_2x_3^2x_4 + x_2x_3^2x_5 + x_2x_3^2x_7 + x_2x_3x_4^2 \\
&\quad + x_2x_3x_5^2 + x_2x_3x_7^2 + x_2x_4^2x_5 + x_2x_4^2x_6 + x_2x_4^2x_7 + x_2x_4x_5^2 + x_2x_4x_6^2 + x_2x_4x_7^2 \\
&\quad + x_2x_5^2x_6 + x_2x_5x_6^2 + x_2x_6^2x_7 + x_2x_6x_7^2 + x_3^2x_4x_5 + x_3^2x_4x_6 + x_3^2x_5x_6 + x_3^2x_5x_7 \\
&\quad + x_3^2x_6x_7 + x_3x_4^2x_5 + x_3x_4^2x_6 + x_3x_4x_5^2 + x_3x_4x_6^2 + x_3x_5^2x_6 + x_3x_5^2x_7 + x_3x_5x_6^2 \\
&\quad + x_3x_5x_7^2 + x_3x_6^2x_7 + x_3x_6x_7^2 + x_4^2x_5x_7 + x_4^2x_6x_7 + x_4x_5^2x_7 + x_4x_5x_7^2 + x_4x_6^2x_7 \\
&\quad + x_4x_6x_7^2 + x_5^2x_6x_7 + x_5x_6^2x_7 + x_5x_6x_7^2 \\
f_7 &= x_1^7 + x_2^7 + x_3^7 + x_4^7 + x_5^7 + x_6^7 + x_7^7
\end{aligned}
$$

of degrees 1, 2, 3, 3, 4, 4, 7 respectively.

2.2 Outline of Proof, Remarks

- The proof of my theorem is in roughly three steps. (I will describe it for $K = \mathbb{Q}$ since it is slightly easier, but the general version is the same):

  1. Given an extension $L/\mathbb{Q}$, use Minkowski's theorems and the geometry of numbers to construct an algebraic integer α generating the extension whose archimedean norms are small (in terms of the discriminant bound X).
     - This is a fairly standard technique: construct the Minkowski lattice by embedding the ring of integers $\mathcal{O}_L$ in $\mathbb{R}^{\#}$ by sending

       $$\alpha \mapsto \Big(\rho_1(\alpha), \ldots, \rho_r(\alpha), \sqrt{2}\operatorname{Re}\sigma_1(\alpha), \sqrt{2}\operatorname{Im}\sigma_1(\alpha), \ldots, \sqrt{2}\operatorname{Re}\sigma_s(\alpha), \sqrt{2}\operatorname{Im}\sigma_s(\alpha)\Big)$$

       where the ρ_i are the r real embeddings and the σ_i are the s nonconjugate complex embeddings.
     - Use Minkowski's theorems to show the existence of an element $\alpha \in \mathcal{O}_L$ which generates $L/\mathbb{Q}$ and all of whose archimedean embeddings satisfy $|\alpha|_i \ll X^{1/2(n-t)}$, where t is the maximal degree of an intermediate subfield.
  2. Use the invariant theory of the group G to construct a finite scheme map to affine space, under which the image of a point corresponding to α has integral image.
     - The map is only a slightly more complicated version of the ring map that evaluates the primary invariant polynomials on the archimedean embeddings of α.
  3. Count the number of possible images of α ("points in a box"). Since α generates $L/\mathbb{Q}$, this will give an upper bound on the number of possible L.
     - By the construction of the primary invariant polynomials, the scheme map from step 2 is finite.

- Since α is integral, and we know that each archimedean embedding satisfies $|\alpha|_i \ll X^{1/2l(n-t)}$ from step 1, we are then essentially reduced to counting the number of points in a box whose dimensions are $\ll X^{\deg(f_i)/2(n-t)}$.
- The laziest possible method would just be to count the number of points in this box, and apply finiteness. The resulting bound is

$$N_{\mathbb{Q},n}(X;G) \ll X^{\frac{1}{2(n-t)}\left[\sum_{i=1}^{n-1}\deg(f_{i+1})\right]},$$

- We can do slightly better by using some sieving: this is the final ingredient in getting the original statement I gave above.

- We would naturally expect the actual number of integral points to be lower than the bound from the theorem, per Malle's heuristics. There are (at least) three ways in which we lose accuracy:

1. The map associating an element $\mathbf{x}$ to an extension L/K is not injective: any extension has many different generators. Worse still, there is no uniform way to account for this non-injectivity: an extension of small discriminant will have many generators of small archimedean norm, and thus it will show up in the count much more frequently than an extension of larger discriminant.
2. The simple techniques employed above for counting integral points on the scheme Z give weaker bounds than could be hoped for. Most points in affine space are not actually the image of an integral point on Z, but we do not expect that the sieving we performed is actually sharp: it is likely only extracting a small amount of the potential savings that should be realizable.
3. If L/K has any intermediate extensions, the bound given in the theorem is weaker than for a primitive extension. The worst losses occur when L/K has a subfield of small index (e.g., index 2), in which case the exponent obtained is nearly doubled. In principle, a more careful analysis of towers of fields could deal with this issue (perhaps not completely, but at least partly).

- We will also note that as n grows, finding a set of primary invariants becomes very computationally intensive, and it becomes infeasible (at present) to compute them for nontrivial transitive subgroups of S_n when n is larger than 9. I have computed the results for all transitive subgroups of S_n for $n \leq 8$, and in general they tend to beat Schmidt's bound by quite a bit when there are no big intermediate subfields.
  - Here is a table for the nontrivial transitive subgroups of S_7 (which are all necessarily primitive):

#	Ord	Isom to	Generators	Invariant Degrees	Result	Malle	Schmidt
T_2	14	$D_{2\cdot 7}$	(1 2 3 4 5 6 7), (1 6)(2 5)(3 4)	1,2,2,2,3,4,7	$X^{19/12}$	$X^{1/3}$	$X^{27/12}$
T_3	21	F_{21}	(1 2 3 4 5 6 7), (1 2 4)(3 6 5)	1,2,3,3,3,4,7	$X^{21/12}$	$X^{1/4}$	$X^{27/12}$
T_4	42	F_{42}	(1 2 3 4 5 6 7), (1 3 2 6 4 5)	1,2,3,3,4,6,7	$X^{24/12}$	$X^{1/3}$	$X^{27/12}$
T_5	168	$PSL_2(\mathbb{F}_7)$	(1 2 3 4 5 6 7), (1 2)(3 6)	1,2,3,3,4,4,7	$X^{22/12}$	$X^{1/2}$	$X^{27/12}$
T_6	2520	A_7	(3 4 5 6 7), (1 2 3)	1,2,3,4,5,6,7	$X^{26/12}$	$X^{1/2}$	$X^{27/12}$

3 The ρ-Discriminant

- I will now briefly discuss some ongoing work in generalizing the counting theorem presented above.
- One way of reinterpreting my theorem is to view it as a result about reduced permutation representations of groups. The invariant theory involved in the proof carries over to general representations ρ, and so one could ask: is there a way to generalize the result to arbitrary permutations ρ?
- The answer turns out to be yes, however, it is necessary to introduce a new counting metric attached to the representation ρ, which I have termed the "ρ-discriminant". (It takes the place of the square root of the relative discriminant, for counting purposes.)

3.0.1 Motivation for the ρ-Discriminant

- Let L/K be a degree-n Galois extension of number fields with Galois group G, and let the respective rings of integers be $\mathcal{O}_L$ and $\mathcal{O}_K$. Additionally, let $\rho : G \to GL_d(\mathcal{O}_K)$ be a faithful representation of G.
  - Note, for emphasis: the extension L/K is now assumed to be Galois! (This is not really a restriction, as I will later explain.)
- The proof of the Theorem can be interpreted as follows:
  - First, we construct a generator α of the extension L/K that has small archimedean valuations relative to the discriminant of the extension.
  - Then we compute the primary invariant polynomials $f_1, \cdots, f_n$ for the permutation representation $\rho : G \to S_n \hookrightarrow GL_n(\mathbb{Z})$, and we observe that $f_i(\mathbf{x})$ lies in K, where $\mathbf{x}$ is the vector of archimedean embeddings of α (on which G acts through ρ and through the Galois action).
  - Next, we use the finiteness of a scheme map originating from invariant theory to conclude that if we fix the values $f_1(\mathbf{x}), \ldots, f_n(\mathbf{x})$, then there are only a bounded number of possibilities for $\mathbf{x}$.
  - Finally, we count the number of possibilities for these invariant values $f_1(\mathbf{x}), \ldots, f_n(\mathbf{x})$, yielding an upper bound for the number of possible $\mathbf{x}$ and in turn the number of possible α, hence (at last) bounding the number of possible L.
- We would like to adapt this technique to a setting with a general representation: so suppose, now, that ρ is an arbitrary degree-d representation.
  - The scheme map originating from invariant theory is still finite, and everything following that point in the argument still holds, provided we can construct some vector $\mathbf{x} \in \mathcal{O}_L^{\oplus d}$ with the property that $f_i(\mathbf{x}) \in K$ for all of the primary invariants f_i.
  - By Galois theory, $f_i(\mathbf{x}) \in K$ if and only if $g \cdot f_i(\mathbf{x}) = f_i(g \cdot \mathbf{x})$ is in K, where $g \in G$ is acting on $\mathbf{x}$ via the Galois action.
  - If we demand that $g \cdot \mathbf{x} = \rho(g)\mathbf{x}$, where we view $\rho(g)$ as acting on $\mathbf{x}$ via the representation action, then since f_i is an invariant polynomial, we would have $g \cdot f_i(\mathbf{x}) = f_i(g \cdot \mathbf{x}) = f_i(\rho(g)\mathbf{x}) = f_i(\mathbf{x})$, which is precisely the outcome we are seeking.
- The correct object to work with is, therefore, the set of tuples of elements of $\mathcal{O}_L$ which the two natural actions of G (via the Galois action or via the representation ρ) agree.

3.0.2 The ρ-Discriminant

- We now carry through the details of the construction we just motivated.
- Observe that there are two natural actions of G on the space

$$\mathcal{O}_L \otimes_{\mathcal{O}_K} \mathcal{O}_K^{\oplus d} \cong \mathcal{O}_L^{\oplus d}.$$

  - First, there is the action δ stemming from the Galois action of G on $\mathcal{O}_L$ (which acts on the left side in the tensor product and diagonally on each copy of $\mathcal{O}_L$ in the direct sum): thus,

$$\delta : G \to \mathrm{Aut}_{\mathcal{O}_K}(\mathcal{O}_L)^{\oplus d}. \tag{1}$$

  - There is also the action τ obtained by acting in the right component of the tensor product by ρ (which in the direct sum is equivalent to extending the representation ρ from its action on $\mathcal{O}_K$ to an action on $\mathcal{O}_L$): thus,

$$\tau : G \to GL_d(\mathcal{O}_K) \hookrightarrow GL_d(\mathcal{O}_L). \tag{2}$$

- The object we are interested in, per the argument above, is the subset of elements where these actions agree:

- Definition: For a given Galois extension L/K with Galois group G and a faithful representation $\rho : G \to GL_d(\mathcal{O}_K)$, we define the tuning submodule Ξ_ρ to be the subset of elements of the space $\mathcal{O}_L^{\oplus d}$ on which the two actions δ and τ from 1 and 2 coincide; namely,

$$\Xi_\rho = \left\{ x \in \mathcal{O}_L^{\oplus d} : \forall g \in G,\, \delta(g)(x) = \tau(g)(x) \right\}. \tag{3}$$

  - The tuning submodule Ξ_ρ is a torsion-free $\mathcal{O}_K$-module of rank d.
  - The torsion-free part is obvious, but the rank is not quite as immediate. Here is an easy way to see the rank is at least d: tensor with K and choose a basis for L/K. Then we obtain a system of nd linear equations in nd variables, where $n = [L : K]$, and the d equations corresponding to $g = 1$ are all trivial. (Hence, the rank of the solution space is at least d.)

- To construct a discriminant-like object using the tuning submodule Ξ_ρ, we follow the analogy with the construction of the classical relative discriminant $D_{L/K}$ and use determinants of elements in Ξ_ρ.

- Definition: Let L/K be a Galois extension with Galois group G, $\rho : G \to GL_d(\mathcal{O}_K)$ is a faithful representation of G and let Ξ_ρ be the tuning submodule attached to (ρ, L, K). We define the ρ-discriminant ideal $D_{L/K}^{(\rho)}$ to be the ideal of $\mathcal{O}_L$ generated by all $d \times d$ determinants of the form $\det(\xi_1, \xi_2, \ldots, \xi_d)$, where each ξ_i, for $1 \le i \le d$, is an element of Ξ_ρ (thought of as a length-d column vector).

  - As defined above, the ρ-discriminant ideal is only an ideal of $\mathcal{O}_L$, since the entries in each determinant lie in $\mathcal{O}_L$.
  - However, in some cases, the ρ-discriminant descends naturally to an ideal of $\mathcal{O}_K$: for any $g \in G$, $\delta(g) \cdot \det(\xi_1, \xi_2, \ldots, \xi_d) = \det(\rho(g)) \cdot \det(\xi_1, \xi_2, \ldots, \xi_d)$, so in the event that $\rho(g) = 1$ for all $g \in G$, we see that $\det(\xi_1, \ldots, \xi_d)$ is Galois-invariant and therefore lies in $\mathcal{O}_K$.
  - Also, if ρ is a permutation representation, then $\det(\rho(g)) = \pm 1$. Then the square of the determinant $\det(\xi_1, \xi_2, \ldots, \xi_d)^2$ will be Galois-invariant; ensuring this Galois-invariance is precisely why the definition of the classical discriminant $D_{L/K}$ uses squares of determinants. In general, the ρ-discriminant will behave analogously to the square root of the classical discriminant.

- Here are a few concrete examples of tuning submodules and ρ-discriminants:

- Example 1: Let $K = \mathbb{Q}$ and $L = \mathbb{Q}(\sqrt{D})$ for a squarefree D. Then $G = \mathbb{Z}/2\mathbb{Z}$, so let ρ be the nontrivial 1-dimensional representation of G.

  - For g the nonidentity element of G, we see that for $x = a + b\sqrt{D} \in \mathcal{O}_L$, we have $\delta(g)(x) = a - b\sqrt{D}$, and $\tau(g)(x) = -x = -a - b\sqrt{D}$, which are equal precisely when $a = 0$.
  - Hence $\Xi_\rho = \mathbb{Z}\sqrt{D}$, and then we can readily see that $D_{L/\mathbb{Q}}^{(\rho)} = \sqrt{D}$.
  - Since the (classical) discriminant D_L of this extension is D or $4D$, depending on whether D is or is not congruent to 1 modulo 4, respectively, we see that the ρ-discriminant and classical discriminant behave similarly, but not identically, for these extensions.

- Example 2: Let L/K be an arbitrary Galois extension of degree n with Galois group G, and take ρ to be the regular representation of G as a subgroup of S_n.

  - Then Ξ_ρ is the set of n-tuples of elements of $\mathcal{O}_L$ such that the permutation action of G agrees with the Galois action of G.
  - Since G is transitive on the n coordinates, we see that the tuning submodule Ξ_ρ is precisely the set of n-tuples $(\alpha^{(1)}, \cdots, \alpha^{(n)})$ of Galois orbits of elements $\alpha \in \mathcal{O}_L$, where $x^{(j)} = \sigma_j x$ represents the image of x under the jth element σ_j of G (for some fixed labeling of the elements in G).
  - The ρ-discriminant $D_{L/K}^{(\rho)}$ ideal is then generated by the determinants of all possible matrices $\{\beta_i^{(j)}\}_{1 \le i,j \le n}$ for $\beta_i \in \mathcal{O}_L$, which is precisely the same set of determinants used to define the (classical) relative discriminant ideal $D_{L/K}$.
  - In this case, due to our slightly different normalization, we have $D_{L/K}^{(\rho)} = \sqrt{D_{L/K}}$.

- If E/K is a (non-Galois) field extension with Galois closure L/K, then I believe (but have not actually done out all of the necessary calculations to verify) that there is a fairly natural representation ρ attached to L/K whose ρ-discriminant recovers the classical discriminant of E/K.
    - Thus, we lose no generality by only working with Galois extensions: by choosing the appropriate representation ρ, we can effectively study the non-Galois extensions of K having a particular Galois group of their Galois closure (in exactly the same way we did earlier).
    - Wood and Varma have a recent result on counting dihedral quartic extensions using a modification of the discriminant that (they have argued) seems more natural for that setting, and which appears to be a special case of the ρ-discriminant construction.

3.0.3 Counting by ρ-Discriminant

- We now obtain a new class of counting problems: counting extensions L/K whose ρ-discriminant is bounded. Explicitly, if we define $N_{K,n}(X;\rho)$ to be the number of number fields L (up to K-isomorphism) such that
    1. The degree $[L:K] = n$,
    2. The Galois group $\mathrm{Gal}(L/K) = G$, and $\rho : G \to GL_d(\mathcal{O}_K)$ is a faithful representation of G, and
    3. The ρ-discriminant norm $\mathrm{Nm}_{L/\mathbb{Q}} D_{L/K}^{(\rho)}$ is less than X.

    we then have the following generalization of the previous theorem:
- Theorem ([D.]): Let K be any number field, G be a finite group of order n, and $\rho : G \to GL_d(\mathcal{O}_K)$ be a faithful d-dimensional representation of G on $\mathcal{O}_K$. Also define $t(\rho)$ to be the smallest positive integer such that for any nontrivial subgroup H of G, $(\mathcal{O}_K^d)^{\rho(H)}$ has rank $\leq t(\rho)$ as an $\mathcal{O}_K$-module. Then

$$N_{K,n}(X;\rho) \ll X^{\frac{1}{d-t(\rho)}\left[\sum_{i=1}^{d} \deg(f_i)\right]},$$

  where the f_i for $1 \leq i \leq d$ are a set of primary invariants for ρ. Furthermore, if ρ has a nontrivial secondary invariant, then we can replace the upper bound by $X^{\frac{1}{d-t(\rho)}\left[\sum_{i=1}^{d} \deg(f_i) - \frac{\deg(f_1)}{2[K:\mathbb{Q}]}\right]+\epsilon}$.
- We have reached the end of the results I wanted to talk about, but I'd like to close with a few other things.
    - Much of this is still very much work in progress: there are a number of properties that the classical discriminant has, and which any sensible object should be known to possess before we should call it a discriminant, but which I have not yet established for the ρ-discriminant.
    - There are many different approaches to "alternate discriminants" (e.g., Varma and Wood's use of an Artin conductor to count dihedral quartics, Yasuda's "V-discriminant", and some work of Silas Johnson on "weighted discriminants"), and it is unclear at the moment how all of these ideas are related. It seems very likely that there are relations between all of these notions, and some of them may even be the same.
    - The pipe dream would be to establish relations between all of these discriminants and study what kinds of arithmetic information they provide. In addition to trying to tackle the kinds of counting questions I discussed, there are many other different directions this could go, but for example: does knowing all of the ρ-discriminants for a given extension L/K necessarily characterize L up to isomorphism? Are there analytic formulas involving the ρ-discriminant? What kind of local data does the ρ-discriminant provide? (And so forth.)