

Optimal Class-Group Tomography for Hecke Eigensystems over Number Fields

Luca Bianchi

June 2026

Abstract

Recent work on Hecke operators over number fields shows that, in important Bianchi and Hilbert settings, principal Hecke operators recover complete eigensystems only up to unramified quadratic twist. This note reformulates that ambiguity as an observable fibre and gives an optimal presentation-theoretic resolver for it. The residual fibre is controlled by the group

$$G_2(K) = \text{Hom}(\text{Cl}(K), \{\pm 1\})$$

of unramified quadratic characters. Non-principal Hecke operators evaluate these characters on ideal classes whenever the relevant Hecke values are nonzero. Therefore resolving a principal-Hecke fibre becomes a finite tomography problem on the class group.

The main theorem includes the self-twist and zero-value phenomena that occur in computations. For an eigensystem λ , let $\mathcal{P}_\lambda$ be the set of ideal classes admitting nonzero non-principal probe values, let $W_\lambda \subseteq \text{Cl}(K)/2\text{Cl}(K)$ be their span, and let $H_{\text{eff}}(\lambda)$ be the image of the residual twist group on W_λ . A set of nondegenerate probe classes resolves the principal fibre exactly when its evaluation map is injective on $H_{\text{eff}}(\lambda)$. The minimum number of independent nondegenerate probes is therefore

$$\dim_{\mathbb{F}_2} H_{\text{eff}}(\lambda),$$

and such a minimum resolver always exists. In restricted or degenerate computational windows, the optimal remaining resolver is a weighted set cover on pairs of residual twists.

The result gives a concrete example where Presentation Theory does not merely rename a known ambiguity. It identifies the fibre, quotients out invisible inner-twist data, chooses observables adapted to the visible fibre, proves a sharp lower bound, gives an algorithm, and produces finite verification data.

1 Introduction

Let K be a number field. In the computational theory of modular and automorphic forms over K , Hecke operators are naturally indexed by ideals of K . Principal ideals and non-principal ideals do not always carry the same amount of information. Work of Cremona on Hecke operators and formal modular symbols over number fields shows that, in the relevant setting, principal Hecke data determine the complete eigensystem only up to unramified twist; after fixing the appropriate character, the remaining ambiguity is quadratic [1]. In Bianchi computations, the role of non-principal Hecke operators and self-twists is also visible in the class-group structure of the algorithm [2].

This note turns that residual ambiguity into a finite tomography theorem. Principal Hecke operators define an observable map. The fibre of this map is controlled by unramified quadratic characters, hence by

$$G_2(K) = \text{Hom}(\text{Cl}(K), \{\pm 1\}).$$

Non-principal Hecke operators evaluate these characters on ideal classes. The only subtlety is that some Hecke values vanish, and some twists are invisible on the nonzero support actually available in a finite computation. The correct resolver is therefore not the full class group, but the effective quotient seen by the nondegenerate probe classes.

The main result is optimal. For a fixed eigensystem λ , let $\mathcal{P}_\lambda$ be the set of ideal classes on which at least one available Hecke value is nonzero. Let W_λ be the span of these classes in $\text{Cl}(K)/2\text{Cl}(K)$, and let $H_{\text{eff}}(\lambda)$ be the image of the residual twist group on W_λ . Then the minimum number of independent nondegenerate class probes needed to resolve the principal fibre, modulo twists invisible on all available nonzero probes, is exactly

$$\dim_{\mathbb{F}_2} H_{\text{eff}}(\lambda).$$

2 Principal data and residual twists

Let $\text{Cl}(K)$ be the ideal class group, and define

$$G_2(K) = \text{Hom}(\text{Cl}(K), \{\pm 1\}).$$

Elements of $G_2(K)$ are the unramified quadratic characters of the ideal class group. We restrict throughout to ideals prime to the declared level.

Throughout this note, an unramified class character means a quadratic character of the finite class quotient acting on the prime-to-level Hecke operators in the normalization of the declared computation. In the simplest case this quotient is $\text{Cl}(K)$. In variants involving ray class characters, nebentypus, or additional local level structure, $\text{Cl}(K)$ should be replaced everywhere by the corresponding finite class quotient. The linear-algebra statements are unchanged after that replacement.

Let $\mathcal{E}$ be a finite computational window of Hecke eigensystems. For $\lambda \in \mathcal{E}$, write

$$\lambda(T_{\mathfrak{a}})$$

for the value of the Hecke operator indexed by the ideal $\mathfrak{a}$. Let

$$\mathcal{O}_{\text{prin}}(\lambda) = (\lambda(T_{\mathfrak{a}}))_{[\mathfrak{a}]=1}$$

be the principal Hecke data.

Definition 2.1 (Residual twist ledger). *The principal residual ledger of λ is*

$$\text{TwLed}(\lambda) = \{\mu \in \mathcal{E} : \mathcal{O}_{\text{prin}}(\mu) = \mathcal{O}_{\text{prin}}(\lambda)\}.$$

In a twist-governed window, we choose a residual twist group

$$H \subseteq G_2(K)$$

such that every element of the ledger is represented by a twist

$$\lambda^\chi, \quad \chi \in H.$$

The group H is part of the finite record. It may be smaller than $G_2(K)$, and it depends on the fixed character, the finite window, the normalization conventions, and possible self-twist phenomena. The theorem below does not assume that the principal fibre is the full quadratic character group.

The twist rule has the form

$$\lambda^\chi(T_{\mathfrak{a}}) = \chi([\mathfrak{a}])\lambda(T_{\mathfrak{a}})$$

for the ideals and normalizations in the declared computation. Principal ideals have class 1, so unramified class characters do not change principal Hecke values. Non-principal ideals can detect the residual character.

3 Observable support

A non-principal class is useful only when it is represented by a Hecke operator whose baseline value is nonzero. This is the finite support on which twist ratios can be read.

Definition 3.1 (Nondegenerate probe class). *A class $c \in \text{Cl}(K)$ is nondegenerate for λ if there is an ideal $\mathfrak{a}$, prime to the level, such that*

$$[\mathfrak{a}] = c \quad \text{and} \quad \lambda(T_{\mathfrak{a}}) \neq 0.$$

Let $\mathcal{P}_{\lambda} \subseteq \text{Cl}(K)$ be the set of nondegenerate probe classes.

Since the characters in $G_2(K)$ are quadratic, they factor through $\text{Cl}(K)/2\text{Cl}(K)$. Let

$$\overline{\mathcal{P}}_{\lambda} \subseteq \text{Cl}(K)/2\text{Cl}(K)$$

be the image of $\mathcal{P}_{\lambda}$, and set

$$W_{\lambda} = \text{span}_{\mathbb{F}_2}(\overline{\mathcal{P}}_{\lambda}) \subseteq \text{Cl}(K)/2\text{Cl}(K).$$

Definition 3.2 (Effective twist group). *The effective twist group of λ is*

$$H_{\text{eff}}(\lambda) = \text{im}(H \longrightarrow \text{Hom}(W_{\lambda}, \{\pm 1\})).$$

The invisible twist ledger is

$$I_{\lambda} = \ker(H \longrightarrow \text{Hom}(W_{\lambda}, \{\pm 1\})).$$

Thus $H_{\text{eff}}(\lambda) \cong H/I_{\lambda}$. The subgroup I_{λ} consists of the residual twists that cannot be distinguished by any nonzero probe in the declared support. It includes the part of the self-twist or zero-value phenomenon that is invisible to the chosen computational window.

4 Optimal tomography theorem

For $c \in \mathcal{P}_{\lambda}$, choose an ideal $\mathfrak{a}$ with class c and $\lambda(T_{\mathfrak{a}}) \neq 0$. For every residual twist $\chi \in H$, the ratio

$$\frac{\lambda^{\chi}(T_{\mathfrak{a}})}{\lambda(T_{\mathfrak{a}})}$$

is equal to $\chi(c)$. A class probe is therefore an evaluation functional on the effective twist group.

Theorem 4.1 (Optimal class-group tomography). *Let $C = (c_1, \dots, c_s)$ be a list of nondegenerate probe classes for λ , and choose ideals $\mathfrak{a}_i$ representing them with*

$$\lambda(T_{\mathfrak{a}_i}) \neq 0.$$

Then the Hecke values

$$\lambda^{\chi}(T_{\mathfrak{a}_1}), \dots, \lambda^{\chi}(T_{\mathfrak{a}_s})$$

separate the principal residual fibre modulo I_{λ} if and only if the evaluation map

$$\text{ev}_C : H_{\text{eff}}(\lambda) \longrightarrow \{\pm 1\}^s, \quad \eta \longmapsto (\eta(c_1), \dots, \eta(c_s))$$

is injective. Equivalently, the classes $c_1, \dots, c_s$ span the dual of $H_{\text{eff}}(\lambda)$ through evaluation.

Proof. Let $\chi, \psi \in H$. The selected Hecke values of λ^χ and λ^ψ agree for all i exactly when

$$\chi(c_i)\lambda(T_{\mathbf{a}_i}) = \psi(c_i)\lambda(T_{\mathbf{a}_i})$$

for all i . Since the chosen baseline values are nonzero, this is equivalent to

$$\chi(c_i) = \psi(c_i)$$

for all i . Equivalently, the image of $\chi\psi^{-1}$ in $H_{\text{eff}}(\lambda)$ lies in the kernel of ev_C .

Therefore the selected Hecke operators distinguish two residual twists precisely when their effective classes in H/I_λ have different images under ev_C . They separate the fibre modulo I_λ precisely when no nontrivial element of $H_{\text{eff}}(\lambda)$ is trivial on all the selected classes. This is exactly injectivity of ev_C . $\square$

Corollary 4.2 (Sharp minimum). *Let*

$$d_\lambda = \dim_{\mathbb{F}_2} H_{\text{eff}}(\lambda).$$

Every independent nondegenerate class-probe resolver of the principal fibre modulo I_λ has length at least d_λ , and there is one of length exactly d_λ . Hence

$$R_{\text{CGT}}(\lambda) = d_\lambda.$$

Proof. Each selected class contributes one binary evaluation on $H_{\text{eff}}(\lambda)$. If s probes separate $H_{\text{eff}}(\lambda)$, then $H_{\text{eff}}(\lambda)$ injects into $\{\pm 1\}^s$. As an elementary abelian 2-group, it has $\mathbb{F}_2$ -dimension d_λ , so $s \geq d_\lambda$.

For existence, use the definition of W_λ . The images of the classes in $\mathcal{P}_\lambda$ span W_λ . Evaluation by elements of W_λ gives the full dual space of $H_{\text{eff}}(\lambda)$, modulo the annihilator of $H_{\text{eff}}(\lambda)$. Since the columns coming from $\mathcal{P}_\lambda$ span that dual, one can choose d_λ pivot columns. The corresponding d_λ nondegenerate probe classes give an injective evaluation map. $\square$

Corollary 4.3 (Uniform families). *Suppose a family has a common residual twist group $H_{\mathcal{F}}$ and a common nondegenerate support space $W_{\mathcal{F}}$. Define*

$$H_{\text{eff}}(\mathcal{F}) = \text{im}(H_{\mathcal{F}} \longrightarrow \text{Hom}(W_{\mathcal{F}}, \{\pm 1\})).$$

Then the minimum number of independent class probes resolving the family, modulo the common invisible ledger, is

$$\dim_{\mathbb{F}_2} H_{\text{eff}}(\mathcal{F}).$$

Proof. The preceding proof is linear algebra over the common support space $W_{\mathcal{F}}$. Applying it to the shared effective image gives the stated minimum. $\square$

In practice a family may not have a single natural support space. One can instead use an intersection support, a union support, or a fibrewise support, producing different universal-resolver notions. The corollary applies only after that support convention is declared.

Example 4.4 (Four-twist fibre). *Suppose, in a declared finite window, that*

$$\text{Cl}(K)/2\text{Cl}(K) \cong \mathbb{F}_2 a \oplus \mathbb{F}_2 b$$

and that the residual twist group is the full dual $H = G_2(K)$. Then there are four residual twists. Assume that the classes a and b are nondegenerate for λ , while every available ideal in the class $a + b$ has baseline Hecke value zero in the window. Then

$$W_\lambda = \mathbb{F}_2 a \oplus \mathbb{F}_2 b, \quad H_{\text{eff}}(\lambda) = H,$$

and the two probes a, b have evaluation matrix

$$\begin{pmatrix} 1 & 1 \\ -1 & 1 \\ 1 & -1 \\ -1 & -1 \end{pmatrix}$$

on the four twists, after choosing an ordering of H . Equivalently, after writing a negative sign as 1 and a positive sign as 0, the two columns have rank 2 over $\mathbb{F}_2$. Thus two nondegenerate probes are necessary and sufficient. The zero values in the class $a + b$ are recorded in the zero-value ledger; they do not change the optimum because a and b already span the effective support.

5 Restricted windows and set cover

The sharp theorem applies to the clean linear setting where all chosen class probes are nondegenerate and where the support space is known. A finite computation may impose additional restrictions: only some ideals have been computed, some values vanish, and costs may vary by norm, ramification, or implementation.

Let F be the finite residual candidate set left by the principal data after identifying candidates that are already indistinguishable in the intended quotient. For an available ideal $\mathfrak{a}$, define

$$C_{\mathfrak{a}} = \{\{\mu, \nu\} \subseteq F : \mu(T_{\mathfrak{a}}) \neq \nu(T_{\mathfrak{a}})\}.$$

Thus $C_{\mathfrak{a}}$ is the set of unordered residual pairs separated by $T_{\mathfrak{a}}$.

Proposition 5.1 (Restricted resolver). *In a finite restricted window, choosing a minimum-cost list of non-principal Hecke operators that separates all remaining residual candidates is exactly weighted set cover on unordered pairs:*

$$\bigcup_{\mathfrak{a} \in J} C_{\mathfrak{a}} = \{\{\mu, \nu\} \subseteq F : \mu \neq \nu\}.$$

Proof. An operator contributes exactly the pairs it separates. A list separates the finite residual set if and only if every unordered pair is separated by at least one chosen operator. Adding an operator cost function therefore gives the standard weighted set-cover problem on the universe of residual pairs. $\square$

6 The CGT resolver

The class-group tomography resolver is the following finite procedure.

- (i) Compute the class group $\text{Cl}(K)$ and the quadratic class-character group $G_2(K)$.
- (ii) Compute the principal Hecke data and identify the residual twist group H in the declared finite window.
- (iii) List the nondegenerate probe classes $\mathcal{P}_{\lambda}$, form their images in $\text{Cl}(K)/2\text{Cl}(K)$, and compute

$$W_{\lambda} = \text{span}_{\mathbb{F}_2}(\overline{\mathcal{P}}_{\lambda}).$$

- (iv) Restrict H to W_{λ} , producing $H_{\text{eff}}(\lambda)$ and I_{λ} .

- (v) Build the evaluation matrix with rows indexed by effective twists and columns indexed by available nondegenerate probe classes:

$$M_\lambda(\eta, c) = \eta(c).$$

- (vi) In the clean case, select pivot columns of rank $\dim_{\mathbb{F}_2} H_{\text{eff}}(\lambda)$. In a restricted window, run weighted set cover on residual pairs.
- (vii) Output the selected non-principal operators and the remaining invisible ledger I_λ .

7 Verification data

A class-group tomography record contains:

- (i) the field K , the level, and the computed class group $\text{Cl}(K)$;
- (ii) the convention identifying the relevant finite class quotient, if it is not the ordinary ideal class group;
- (iii) the principal Hecke data and the residual twist group H ;
- (iv) the nondegenerate support $\mathcal{P}_\lambda$, the space W_λ , and the invisible ledger I_λ ;
- (v) the selected ideals $\mathfrak{a}_1, \dots, \mathfrak{a}_s$ and their classes;
- (vi) the evaluation matrix on $H_{\text{eff}}(\lambda)$;
- (vii) a rank witness for the selected columns;
- (viii) the zero-value ledger recording unavailable classes;
- (ix) the prime-to-level restriction and the Hecke-operator normalization ledger;
- (x) the residual ledger after adding the selected non-principal observables.

The audit is finite. It checks the class group computation, the twist ledger, the nonzero Hecke values, the evaluation matrix, and the rank condition proving optimality in the clean case.

8 Relation with Presentation Theory

This application has the presentation-theoretic form

$$\text{principal Hecke data} \longrightarrow \text{effective class-group fibre} \longrightarrow \text{minimal non-principal resolver}.$$

The first arrow is a quotient: principal operators forget unramified quadratic twist data. The second arrow is a resolver: non-principal ideal classes evaluate the remaining characters. The self-twist and zero-value phenomena are not exceptions to the method; they determine the support space W_λ and the invisible ledger I_λ .

The useful output is not just that extra Hecke operators may distinguish twists. The theorem says exactly which quotient is distinguishable, how many independent probes are necessary, how to choose them, and how to audit the choice in a finite computation.

References

- [1] J. E. Cremona, *Hecke operators, Hecke eigensystems, and formal modular symbols over number fields*, arXiv:2601.17524, 2026.
- [2] K. Thalagoda and D. Yasaki, *Bianchi Modular Forms over Class Number 4 Fields*, arXiv:2502.00141, 2025.
- [3] J. Cremona, J. Balakrishnan, N. Dunfield, M. Harrison, D. Roberts, W. Stein, and others, *The L-functions and modular forms database project*, arXiv:1511.04289, 2015.
- [4] J. Neukirch, *Algebraic Number Theory*, Grundlehren der mathematischen Wissenschaften 322, Springer, 1999.
- [5] T. Miyake, *Modular Forms*, Springer Monographs in Mathematics, Springer, 2006.