Lattice methods for algebraic modular forms on orthogonal groups

John Voight
Dartmouth College

joint work with
Matthew Greenberg
and Jeffery Hein and Gonzalo Tornaría

Computational Challenges in the Theory of Lattices
ICERM, Providence
24 April 2018

Lattices in a quadratic space

Slight shift in perspective: consider lattices in a (fixed) quadratic space.

Let V be a $\mathbb{Q}$-vector space with $\dim_{\mathbb{Q}} V = n$.
Let $Q\colon V \to \mathbb{Q}$ be a positive definite quadratic form. Write $T\colon V \times V \to \mathbb{Q}$ for the associated bilinear form, defined by $T(x, y) := Q(x + y) - Q(x) - Q(y)$ for $x, y \in V$.

Let $\Lambda < V$ be a (full) lattice, the $\mathbb{Z}$-span of a basis for V. Suppose that Λ is **integral**, i.e., $Q(\Lambda) \subseteq \mathbb{Z}$.

We represent a lattice in bits by a basis $\Lambda = \mathbb{Z}e_1 + \cdots + \mathbb{Z}e_n$; we obtain a quadratic form

$$Q_\Lambda(x) = Q(x_1 e_1 + \cdots + x_n e_n) \in \mathbb{Z}[x_1, \ldots, x_n]$$

and vice versa. Define

$$\operatorname{disc}(\Lambda) := \frac{\det(T(e_i, e_j))_{i,j}}{2^{n \bmod 2}} \in \mathbb{Z}_{>0}.$$

Isometry classes

We define the orthogonal group

$$\begin{aligned} \mathrm{O}(V) &:= \{g \in \mathrm{GL}(V)\colon Q(gx) = Q(x) \text{ for all } x \in V\} \\ \mathrm{O}(\Lambda) &:= \{g \in \mathrm{O}(V) : g\Lambda = \Lambda\}. \end{aligned}$$

We have $\# \mathrm{O}(\Lambda) < \infty$.

Lattices $\Lambda, \Pi \subset V$ are **isometric**, written $\Lambda \simeq \Pi$, if there exists $g \in \mathrm{O}(V)$ such that $g\Lambda = \Pi$.

Same with isometric over $\mathbb{Q}_p$, with $g_p \in \mathrm{O}(V \otimes \mathbb{Q}_p)$ for a prime p.

The **genus** of Λ is

$$\mathrm{Gen}(\Lambda) := \{\Pi < V : \Lambda_p \simeq \Pi_p \text{ for all } p\}.$$

The **class set** $\mathrm{Cl}(\Lambda) := \mathrm{Gen}(\Lambda)/\simeq$ is the set of (global) isometry classes in $\mathrm{Gen}(\Lambda)$. By the geometry of numbers, we have $\# \mathrm{Cl}(\Lambda) < \infty$.

Kneser's p-neighbors

Kneser's theory of p-neighbors (1957) gives an effective method to compute the class set. It will also give us a Hecke action!

Let $p \nmid \operatorname{disc}(\Lambda)$ be prime; $p = 2$ is OK.

We say that a lattice $\Pi < V$ is a p-**neighbor** of Λ, and write $\Pi \sim_p \Lambda$, if Π is integral and

$$[\Lambda : \Lambda \cap \Pi] = [\Pi : \Lambda \cap \Pi] = p.$$

If $\Lambda \sim_p \Pi$, then:

- $\operatorname{disc}(\Lambda) = \operatorname{disc}(\Pi)$,
- $\Pi \in \operatorname{Gen}(\Lambda)$.

Explicit neighbors

Let $H\colon \mathbb{Q}^2 \to \mathbb{Q}$ denote the **hyperbolic plane** defined by $H(x,y) = xy$.

- $\Pi \sim_p \Lambda$ if and only if $\Lambda_q = \Pi_q$ for all $q \neq p$, and there exists a splitting

  $$\Lambda_p = (\mathbb{Z}_p e_1 \oplus \mathbb{Z}_p e_2) \boxplus \Lambda'_p \simeq H_p \boxplus \Lambda'_p$$

  such that

  $$\Pi_p = \mathbb{Z}_p(\tfrac{1}{p} e_1) + \mathbb{Z}_p(p e_2) \boxplus \Lambda'_p.$$

- $\Pi \sim_p \Lambda$ if and only if there exists $v \in \Lambda$ such that $Q(v) \equiv 0 \pmod{2p^2}$ and

  $$\Pi = (p^{-1}v)\mathbb{Z} + \{w \in \Lambda : T(v,w) \in p\mathbb{Z}\}.$$

  The line spanned by v uniquely determines Π, so there are as many p-neighbors as their are isotropic $\mathbb{Z}_p$-lines in Λ_p.

Computing p-neighbors

The number of p-neighbors is $O(p^m)$ where m is the **Witt index** of $\Lambda/p\Lambda$ over $\mathbb{F}_p$, i.e., the maximum dimension of a totally isotropic $\mathbb{F}_p$-subspace. We have $n = 2m, 2m+1, 2m+2$. If $n = 3$, then the number of p-neighbors is $p+1$.

The set of p-neighbors can be computed in time $O(p^{m+\epsilon} H_n(\|\Lambda\|))$ where $\|\Lambda\|$ is the bit size and H_n is a polynomial depending on n (the bit operations in computing a Hermite normal form).

All classes are neighbors

There is an effectively computable finite set S of primes such that every $[\Lambda'] \in \mathrm{Cl}(\Lambda)$ is an **iterated S-neighbor**

$$\Lambda \sim_{p_1} \Lambda_1 \sim_{p_2} \cdots \sim_{p_r} \Lambda_r \simeq \Lambda'$$

with $p_i \in S$.

Typically (when there is only one spinor genus in the genus) we may take $S = \{p\}$ for any $p \nmid \mathrm{disc}(\Lambda)$.

Example

Let $\Lambda = \mathbb{Z}^3 = \mathbb{Z}e_1 + \mathbb{Z}e_2 + \mathbb{Z}e_3 \subset \mathbb{Q}^3$ have the quadratic form

$$Q_\Lambda(x, y, z) = x^2 + y^2 + 3z^2 + xz$$

and bilinear form given by

$$(T(e_i, e_j))_{i,j} = \begin{pmatrix} 2 & 0 & 1 \\ 0 & 2 & 0 \\ 1 & 0 & 6 \end{pmatrix}.$$

Thus $\operatorname{disc}(Q_\Lambda) = 11$. We have $\#\operatorname{Cl}(\Lambda) = 2$, with the nontrivial class represented by the 3-neighbor

$$\Lambda' = \mathbb{Z}e_1 + 3\mathbb{Z}e_2 + \tfrac{1}{3}\mathbb{Z}(e_1 + 2e_2 + e_3)$$

with corresponding quadratic form

$$Q_{\Lambda'}(x, y, z) = x^2 + 9y^2 + z^2 + 4yz + xz.$$

Hecke action

The space of **orthogonal modular forms for** Λ (with trivial weight) is

$$M(\mathrm{O}(\Lambda)) := \mathrm{Map}(\mathrm{Cl}(\Lambda), \mathbb{C}).$$

In the basis of characteristic functions for Λ we have $M(\mathrm{O}(\Lambda)) \simeq \mathbb{C}^h$ where $h = \#\,\mathrm{Cl}(\Lambda)$.

For $p \nmid \mathrm{disc}(\Lambda)$, define the **Hecke operator**

$$\begin{aligned} T_p \colon M(\mathrm{O}(\Lambda)) &\to M(\mathrm{O}(\Lambda)) \\ f &\mapsto T_p(f) \\ T_p(f)([\Lambda']) &:= \sum_{\Pi' \sim_p \Lambda'} f([\Pi']). \end{aligned}$$

The operators T_p commute and are self-adjoint with respect to a natural inner product. So there is a basis of simultaneous eigenvectors, called **eigenforms**.

This is case of the orthogonal group for the theory of *algebraic modular forms* (Gross 1999).

Example

In the running example with discriminant 11, we compute

$$[T_2] = \begin{pmatrix} 1 & 2 \\ 3 & 0 \end{pmatrix}, \quad [T_3] = \begin{pmatrix} 2 & 2 \\ 3 & 1 \end{pmatrix}, \quad [T_5] = \begin{pmatrix} 4 & 2 \\ 3 & 3 \end{pmatrix}, \quad \ldots.$$

These can also be thought of as *adjacency matrices* for the p-neighbor graph.

We find eigenvectors $e = \begin{pmatrix} 1 \\ 1 \end{pmatrix}, f = \begin{pmatrix} 2 \\ -3 \end{pmatrix} \in M(\mathrm{O}(\Lambda))$. The eigenvector e is an *Eisenstein series* with $T_p(e) = (p+1)e$. We have $T_p(f) = a_p f$ with

$$a_2 = -2,\ a_3 = -1,\ a_5 = 1, \ldots$$

We match it with the classical modular form

$$\sum_{n=1}^{\infty} a_n q^n = \prod_{n=1}^{\infty} (1-q^n)^2 (1-q^{11n})^2 = q - 2q^2 - q^3 + \ldots \in S_2(\Gamma_0(11)).$$

The *Atkin–Lehner* involution $z \mapsto \dfrac{-1}{11z}$ acts on $f(z)\,dz$ with eigenvalue $w_{11} = -a_{11} = -1$.

Quick review of modular forms

Let $N \in \mathbb{Z}_{>0}$. A **modular form** of weight $k \in 2\mathbb{Z}_{\geq 0}$ and level N is a holomorphic function $f \colon \mathcal{H} \to \mathbb{C}$ such that

$$f\left(\frac{az+b}{cz+d}\right) = (cz+d)^k f(z)$$

for all $\begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{SL}_2(\mathbb{Z})$ with $N \mid c$ and such that f is holomorphic at the cusps. We accordingly write $f \in M_k(\Gamma_0(N))$.

We further say f is a **cusp form** if f vanishes at the cusps. There are several maps $M_k(\Gamma_0(M)) \to M_k(\Gamma_0(N))$ for $M \mid N$, and we say f is **new** if it is not in the span of the images of these maps.

Computing the Hecke action

To compute the matrix representing the Hecke operator, we need to identify the isometry classes of the p-neighbors of a lattice.

This can be accomplished on lattices using an algorithm of Plesken–Souveignier: match up short vectors and use lots of tricks to compute an isometry or rule it out as early as possible. This is very fast in practice and in fixed dimension theoretically efficient.

Theorem (Haviv–Regev 2014)

There exists a deterministic algorithm that,
given as input lattices $\Lambda, \Lambda' < V$,
computes as output all $g \in \mathrm{O}(V)$ such that $g\Lambda = \Lambda'$
using $n^{O(n)} s^{O(1)}$ bit operations and space $s^{O(1)}$, where s is the input size.

In our setting, we need only *one* isometry $g \in \mathrm{O}(V)$.

Hashing isometry classes

How efficiently (in theory, in practice) can the isometry class of a lattice be identified? Is there an efficiently computable "hash function" on $\mathrm{Gen}(\Lambda)$ that is well-defined on $\mathrm{Cl}(\Lambda)$?

If so, we do not need to do $O(h)$ isometry tests!

Example answer over $\mathbb{Q}$ in small dimension: use *reduction theory*. For example, if $n = 3$ there is explicit reduction theory of integral ternary quadratic forms due to Eisenstein. The result is a *unique* reduced form, so that isometry testing is replaced by table lookup.

More generally, use Minkowski reduction or a Voronoi region?

Or use number of short vectors? With action of small automorphisms? Theta series? Use the duals? Can you leverage the fact that $\Lambda \sim_p \Pi$ agree on a large sublattice? ...

(There is an idea of using *almost autometries* in the Ph.D. thesis of Daniel Kim Murphy to minimize isometry tests.)

We would like this also for lattices over rings of integers.

Classical modular forms

Let $S(\mathrm{O}(\Lambda)) \subset M(\mathrm{O}(\Lambda))$ be the orthogonal complement of the constant functions.

Theorem (Birch 1991, Hein 2016)

Suppose $n = 3$ and $N = \operatorname{disc}(\Lambda)$ is squarefree. Let $\epsilon_p \in \{\pm 1\}$ be the p-Witt invariant for $p \mid N$ and let $D = \prod_{p : \epsilon_p = -1} p$. Then there is a Hecke-equivariant inclusion

$$S(\mathrm{O}(\Lambda)) \hookrightarrow S_2(\Gamma_0(N))$$

whose image is

$S_2(\Gamma_0(N); D\text{-new}; w = \epsilon) :=$
$\{f \in S_2(\Gamma_0(N)) : f \text{ is new at all } p \mid D \text{ and } W_p f = \epsilon_p f \text{ for all } p \mid N\}.$

Computing classical modular forms

Theorem (Hein–Tornaría–V)

There exists an explicit, deterministic algorithm that, given input

$$\begin{gathered}\text{a weight } k \in 2\mathbb{Z}_{>0},\\ \text{a factored } \textit{nonsquare} \text{ level } N = \textstyle\prod_i p_i^{e_i},\\ D \mid \textstyle\prod_{2\nmid e_i} p_i \;\; \text{with an } \textit{odd} \text{ number of factors,}\\ \text{and } \epsilon \in \{\pm 1\}^r,\end{gathered}$$

computes as output the space $S_k(\Gamma_0(N); D\text{-new}; w = \epsilon)$ as a Hecke module.

After precomputation steps (hard to analyze, instantaneous in practice), the running time of the algorithm to compute T_p is $\widetilde{O}(pd)$, where

$$d = \dim S(\mathrm{O}(\Lambda), \rho) = \dim S_k(\Gamma_0(N); D\text{-new}; w = \epsilon) = O(2^{-r}kN).$$

Computational results

For level $N = 1062347 = 11 \cdot 13 \cdot 17 \cdot 19 \cdot 23$ and $D = N$ (so all forms are new), we take

$$Q(x, y, z) = x^2 + 187y^2 + 1467z^2 - 187xz$$

and have $\# \operatorname{Cl}(\Lambda) = 2016$.

Given Q, we can compute $[T_2], [T_3], [T_5], [T_7]$ for *all* signs (giving all newforms) in 4 seconds on a standard desktop machine. Then 1 minute of linear algebra computing kernels with sparse matrices in Magma gives that there are exactly 5 elliptic curves with conductor N.

This isn't a “generic” level! That being said, the same computation with modular symbols in Magma crashed after consuming all 24 GB of available memory.

Examples in higher rank, applications

The space $S(\mathrm{O}(\Lambda))$ with $n = 4$ computes certain Hilbert modular forms over real quadratic fields.

For $n = 5$, we find Siegel paramodular forms. For example, with $N = 61$ we compute Euler factors for $p < 100$ giving enough terms of the degree 4 L-series to verify the functional equation. Conjecturally, this is attached to a Calabi–Yau threefold!

For large n but discriminant 1, Chenevier–Lannes have many results. For example, $\#\,\mathrm{Cl}(E_8 \boxplus E_8) = 2$ and

$$T_p = c_{16}(p) + (1 + p + p^2 + p^3)\frac{1 + p^{11} - \tau(p)}{691}\begin{pmatrix} -405 & 286 \\ 405 & -286 \end{pmatrix}$$

where

$$\Delta(q) = q\prod_{n=1}^{\infty}(1 - q^n)^{24} = \sum_{n=1}^{\infty}\tau(n)q^n.$$

Conclusion

- Isometry classes of lattices with the p-neighbor relation compute spaces of orthogonal modular forms.
- The main workhorse is efficiently identifying isometry classes of p-neighbors. Can these classes be hashed efficiently?
- For $n = 3$ (ternary quadratic forms), orthogonal modular forms compute classical modular forms; with the right weight module, we can carve out desirable subspaces (new, Atkin–Lehner signs) and this implementation is *very fast*!
- More generally, we can consider other definite forms of classical groups (*algebraic modular forms*), including unitary and symplectic groups, over totally real fields.
- For large n, this approach allows us to peer deeply into the world of automorphic forms.

Thank you for your attention!