

CSE206A: Lattice Algorithms and Applications

Harmonic Analysis

UCSD – Fall 2025

Daniele Micciancio



- The problem:
 - Use a lattice basis B to generate a probability distribution D on lattice $\Lambda = \mathcal{L}(B)$
 - Make D independent of B and Λ
- Example Motivation 1: Worst-case/average-case connection
 - Goal: find short vectors in a given (fixed, worst-case) lattice Λ
 - Want to map Λ to a random distribution R on an (average-case) problem R which does not depend on Λ
- Example Motivation 2: Digital Signatures
 - Generate lattice points (sigs) close to a target (msg) using a secret lattice basis B
 - Do not want signatures to “leak” information about secret basis B

- Continuous gaussian distribution on $\mathbb{R}^n$
 - Pick each coordinate independently with probability $\rho(x_i) = e^{-\pi x_i^2}$
 - Point $\vec{x}$ has probability $\rho(\vec{x}) = \prod_i \rho(x_i) = \prod_i e^{-\pi x_i^2} = e^{-\pi \sum_i x_i^2} = e^{-\pi \|\vec{x}\|^2}$
 - Distribution D is spherically symmetric! $\rho(\vec{x}) = e^{-\pi \|\vec{x}\|^2}$ depends only on $\|\vec{x}\|$
- Why π ? So that $\int_{\vec{x}} \rho(\vec{x}) d\vec{x} = 1$. More generally, $c + s D$ has p.d.f. $\frac{1}{s^n} \rho\left(\frac{\vec{x} - \vec{c}}{s}\right)$
- Discrete Gaussian: pick $\vec{x}$ from Λ with probability proportional to $\approx \rho(\vec{x})$
 - How? Possibly with the help of a basis B
 - Still, probability $\rho(\vec{x})$ depends only on $\vec{x} \in \Lambda$, and not B
 - Add (gaussian) noise to $\vec{x}$ to make distribution $\approx$ uniform in $\mathbb{R}^n(\text{mod } \Lambda)$

- Let X be the gaussian distribution in $\mathbb{R}^n$ with $\Pr\{\vec{x}\} = e^{-\pi \|\vec{x}\|^2}$
- Apply a linear transformation $T: \mathbb{R}^n \rightarrow \mathbb{R}^k$. Is $T(X)$ still gaussian?
 - Yes, if $T = Q$ is an orthogonal transformation ($Q^t Q = I$)
 - No, in general. E.g., consider $T(x_1, x_2, \dots, x_n) = (x_1, 2x_2, \dots, nx_n)$
 - Still, $T(X)$ is an “ellipsoidal gaussian”
- If T is an arbitrary linear transformation
 - $A = T^t T$ is a symmetric, positive semidefinite matrix: $x^t A x = \|A x\|^2 \geq 0$
 - Diagonalize $A = Q^t S Q$ where $S = \text{diag}(s_1, \dots, s_n)$, $s_i \geq 0$, $Q^t Q = Q Q^t = I$
 - $\|Tx\|^2 = x^t A x = x^t Q^t S Q x = \|\sqrt{S} Q x\|^2$: $T(X)$ is an ellipsoidal Gaussian

Definition 1. A symmetric matrix $A \in \mathbb{R}^{n \times n}$ is positive semidefinite if $x^t A x \geq 0$ for all $x \in \mathbb{R}^n$. It is positive definite if $x^t A x = 0$ if and only if $x = 0$. (Notation: $A \succcurlyeq 0$, $A \succ 0$)

- Partial ordering on symmetric matrices: $A \succcurlyeq B$ if $(A - B) \succcurlyeq 0$
- Any positive definite matrix can be written as $A = Q^t S Q$ where $S = \text{diag}(s_1, \dots, s_n)$ with $s_i \geq 0$ and Q is orthonormal
 - Change of coordinates $x \mapsto y = Q x$. Then A maps x to $\sum_i s_i y_i^2$
 - A is positive definite if $s_i > 0$ for all i
- If $s_i = 0$ for $i > k$, then use only top k rows Q' to map $x \in \mathbb{R}^n \mapsto Q' x \in \mathbb{R}^k$
 - $x^t A x$ depends only on the projection $y = Q' x$

- $T \in \mathbb{R}^{k \times n}: \mathbb{R}^n \rightarrow \mathbb{R}^k$ linear transformation, $D_n = D_1^n$ gaussian distribution on $\mathbb{R}^n$
- $\Sigma = T^t T = Q^t S Q$ where $Q \in \mathbb{R}^{k \times n}$, $S = \text{diag}(s_1, \dots, s_k)$, $\sqrt{S} = \text{diag}(\sqrt{s_1}, \dots, \sqrt{s_k})$
- $T(D) = \sqrt{S} Q(D)$
 - Q maps the gaussian distribution D_n to D_k
 - $\sqrt{S}$ scales each dimension by $\sqrt{s_i}$
 - Output distribution depends only on Σ : $\rho(Tx) = e^{-\pi \|Tx\|^2} = e^{-\pi x^t \Sigma x}$
 - Write $\sqrt{\Sigma}$ for any T such that $T^t T = \Sigma$
- Ellipsoidal gaussian with parameter $\sqrt{\Sigma}$, $\Sigma \succcurlyeq 0$: $\Pr\{D_{\sqrt{\Sigma}} = x\} = \Pr\{\sqrt{\Sigma} D = x\}$
- If $\Sigma = s^2 I$ is diagonal, this is the same as the scaled gaussian $s D$

- Lattice $\Lambda \subset \mathbb{R}^n$, discrete gaussian D_Λ :
 - $\forall x \in \Lambda. \Pr\{x \leftarrow D_\Lambda\} = \rho(x) / \rho(\Lambda)$
 - $\rho(\Lambda) = \sum_{x \in \Lambda} \rho(x)$ so that $\sum_{x \in \Lambda} \Pr\{x \leftarrow D_\Lambda\} = 1$
 - Informally: continuous gaussian distribution $\Pr\{x\} = \rho(x)$ conditioned on $x \in \Lambda$
- Discrete Gaussian on cosets $\forall x \in \Lambda + c. \Pr\{x \leftarrow D_{\Lambda+c}\} = \rho(x) / \rho(\Lambda + c)$
- Shifted Discrete Gaussian $D_{\Lambda,c}$: $\forall x \in \Lambda. \Pr\{x \leftarrow D_{\Lambda,c}\} = \rho(x - c) / \rho(\Lambda - c)$
- Scaled Shifted Discrete Gaussian $D_{\Lambda,c,\sqrt{\Sigma}}$:

$$\Pr\{x\} = \rho\left(\sqrt{\Sigma^{-1}}(x - c)\right) / \rho\left(\sqrt{\Sigma^{-1}}(\Lambda - c)\right)$$

- $T: \mathbb{R}^n \rightarrow \mathbb{R}^k$ linear transformation, $D_{\Lambda, \sqrt{\Sigma}}$ discrete gaussian
- Is $T(D_{\Lambda})$ a discrete gaussian?
 - Yes, if T is injective: $T(D_{\Lambda}) = D_{T(\Lambda), \sqrt{T^t \Sigma T}}$
 - No, in general: $T(\Lambda)$ may not even be a lattice!
- Assume $T(\Lambda)$ is a lattice. Then, $T(D_{\Lambda, \sqrt{\Sigma}})$ is well defined:
 - $\Pr\{y \leftarrow T(D_{\Lambda, \sqrt{\Sigma}})\} = \Pr\{T^{-1}(y) \leftarrow D_{\Lambda, \sqrt{\Sigma}}\}$ where $T^{-1}(y) = \{x \in \Lambda: T(x) = y\}$
 - But, is $T(D_{\Lambda, \sqrt{\Sigma}})$ a discrete gaussian on $T(\Lambda)$?

- Let $\Gamma \subseteq \Lambda$ be two lattices. Then $G = \Lambda/\Gamma = \{x + \Gamma : x \in \Lambda\}$ is a finitely generated group

Exercise 1. Any finitely generated group can be represented as Λ/Γ

- Assume Γ is a full rank sublattice, then $|G| = \det(\Gamma)/\det(\Lambda)$ is finite
- The dual lattices satisfy $\hat{\Lambda} \subseteq \hat{\Gamma}$. Define the dual group $\hat{G} = \hat{\Gamma}/\hat{\Lambda}$
- The definition of $\hat{G}$ does not depend on the choice of Λ, Γ
- The dual group has size $|\hat{G}| = |G|$. Why?

- For any finite abelian group $G = (\Lambda / \Gamma)$, define $V = G \rightarrow \mathbb{C} = \mathbb{C}^G$
- Hermitean (scalar) product: $\langle f, g \rangle = \frac{1}{|G|} \sum_{x \in G} f(x) \overline{g(x)}$
- Characters: $\chi_v(x) = e^{2\pi i \langle x, v \rangle}$ where $v \in \hat{G} = (\hat{\Gamma} / \hat{\Lambda})$ and $\chi_v: G \rightarrow \mathbb{C}$

Remark 2. $\chi_v(x)$ does not depend on representatives for $x + \Gamma \in G$ and $v + \hat{\Lambda} \in \hat{G}$:

$$\langle x + \Gamma, v + \hat{\Lambda} \rangle \in \langle x, v \rangle + \langle x, \hat{\Lambda} \rangle + \langle \Gamma, v \rangle + \langle \Gamma, \hat{\Lambda} \rangle \subseteq \langle x, v \rangle + \langle \Lambda, \hat{\Lambda} \rangle \equiv_1 \langle x, v \rangle$$

So, $\chi_v(x) = e^{2\pi i (\langle x, v \rangle + \mathbb{Z})} = e^{2\pi i \langle x, v \rangle} e^{2\pi i \mathbb{Z}} = e^{2\pi i \langle x, v \rangle}$

- Finite groups: $G = (\Lambda / \Gamma)$, $\hat{G} = (\hat{\Gamma} / \hat{\Lambda})$. Vector space: $V = G \rightarrow \mathbb{C} = \mathbb{C}^G$
- Characters $\{\chi_v : v \in \hat{G}\} \subset V$ where $\chi_v(x) = e^{2\pi i \langle x, v \rangle}$ for all $x \in G$

Theorem 3. $\langle \chi_v, \chi_w \rangle = \delta_{v,w}$

Proof. $\langle \chi_v, \chi_v \rangle = \mathbb{E}_x[\chi_v(x) \overline{\chi_v(x)}] = \mathbb{E}_x[e^{2\pi i \langle v, x \rangle} e^{-2\pi i \langle v, x \rangle}] = \mathbb{E}_x[1] = 1.$

Now assume $v \neq w$, so $(v - w) \notin \hat{\Lambda}$ and $\exists \mathbf{y} \in \Lambda. \langle v - w, \mathbf{y} \rangle \notin \mathbb{Z}$ and $G + \mathbf{y} = G$

$$\langle \chi_v, \chi_w \rangle = \mathbb{E}_x[\chi_v(x) \overline{\chi_w(x)}] = \mathbb{E}_x[e^{2\pi i \langle v - w, x \rangle}] = \mathbb{E}_x[e^{2\pi i \langle v - w, x + \mathbf{y} \rangle}] = \langle \chi_v, \chi_w \rangle e^{2\pi i \langle v - w, \mathbf{y} \rangle}$$

Since $e^{2\pi i \langle v - w, \mathbf{y} \rangle} \neq 1$ it must be $\langle \chi_v, \chi_w \rangle = 0.$



- $|G|$ -dimensional vector space: $V = \mathbb{C}^G$
- $|\hat{G}| = |G|$ mutually orthogonal vectors $\langle \chi_v, \chi_w \rangle = \delta_{v,w}$
- So, $\{\chi_v: v \in \hat{G}\}$ is an orthonormal basis for V
- For any vector $f: G \rightarrow \mathbb{C}$ in V
 - $f(x) = \sum_v c_v \chi_v(x)$
 - $c_v = \langle f, \chi_v \rangle$

- **Fourier coefficients** $c_v = \langle f, \chi_v \rangle$

- **Fourier Transform** $\hat{f}: \hat{G} \rightarrow \mathbb{C}$, $\hat{f}(v) = \langle f, \chi_v \rangle$:

$$\hat{f}(v) = \frac{1}{|G|} \sum_x f(x) e^{-2\pi i \langle x, v \rangle}$$

- **Inverse Fourier Transform** $f(x) = \sum_v \hat{f}(v) \chi_v(x)$:

$$f(x) = \sum_v \hat{f}(v) e^{2\pi i \langle x, v \rangle}$$

- A function $\varphi: \mathbb{R}^n \rightarrow \mathbb{C}$ is periodic modulo $\Lambda \subset \mathbb{R}^n$ if $\varphi(x) = \varphi(x + v)$ for all $v \in \Lambda$
 - $\varphi(x) = \text{dist}(x, \Lambda)$
 - $\varphi(x) = f(x + \Lambda) = \sum_{v \in \Lambda} f(x + v)$
- Assume $\varphi: \mathbb{R}^n / \Lambda \rightarrow \mathbb{C}$ is a “nice” function: $\varphi(x) \approx \varphi\left(\frac{\lfloor qx \rfloor_\Lambda}{q}\right)$ where $\lfloor \cdot \rfloor_\Lambda: \mathbb{R}^n \rightarrow \Lambda$
 - Restricting φ to $\frac{1}{q}\Lambda$ gives a function $\varphi_q: G_q \rightarrow \mathbb{C}$ where $G_q = \left(\frac{1}{q}\Lambda\right) / \Lambda \equiv \mathbb{Z}_q^n$
 - Fourier transform: $\hat{\varphi}_q: \hat{G}_q \rightarrow \mathbb{C}$ where $\hat{G}_q = \hat{\Lambda} / (q\hat{\Lambda})$ and $\hat{\varphi}_q(v) = \mathbb{E}_{x \in G_q}[\varphi_q(x) \overline{\chi_v(x)}]$
- Fourier Transform for periodic function $\varphi: \mathbb{R}^n / \Lambda \rightarrow \mathbb{C}$:

$$\hat{\varphi}(v) = \lim_{q \rightarrow \infty} \hat{\varphi}_q(v) = \mathbb{E}_{x \in \mathbb{R}^n / \Lambda}[\varphi(x) \overline{\chi_v(x)}]$$

- Fourier Transform: $\hat{\varphi}(v) = \lim_{q \rightarrow \infty} \hat{\varphi}_q(v) = \mathbb{E}_{x \in \mathbb{R}^n / \Lambda} [\varphi(x) \overline{\chi_v(x)}]$
 - $\hat{\varphi}_q: (\hat{\Lambda} / q \hat{\Lambda}) \rightarrow \mathbb{C}$ is periodic modulo $q \hat{\Lambda}$
 - As $q \rightarrow \infty$, $\hat{\varphi}: \hat{\Lambda} \rightarrow \mathbb{C}$
- Fourier Inversion: $\varphi(x) = \sum_{v \in \hat{\Lambda}} \hat{\varphi}(v) \chi_v(x)$
 - Characters $\chi_v: \mathbb{R}^n \rightarrow \mathbb{C}$ are periodic modulo Λ
 - Dual vectors $v \in \hat{\Lambda}$ describes the high dimensional “fundamental frequencies” of φ
 - Each character is used with Fourier coefficient $\hat{\varphi}(v)$

- For any (non-periodic) function $f: \mathbb{R}^n \rightarrow \mathbb{C}$ such that $\int_{\mathbb{R}^n} |f(x)| \, dx < \infty$, define

$$\hat{f}: \mathbb{R}^n \rightarrow \mathbb{C} \qquad \hat{f}(v) = \int_{\mathbb{R}^n} f(x) \overline{\chi_v(x)} \, dx$$

- There is a corresponding inversion formula $f(x) = \int_{\mathbb{R}^n} \hat{f}(v) \chi_v(x) \, dv$

..... but we won't use it!

Theorem 4. *The gaussian function $\rho(x) = e^{-\pi \|x\|^2}$ has fourier transform $\hat{\rho}(x) = \rho(x)$*

Proof. Note that $\hat{\rho}(x) = \prod_i \hat{\rho}(x_i)$. Then, use calculus.



Theorem 5. For any function $f: \mathbb{R}^n \rightarrow \mathbb{C}$ and lattice $\Lambda \subset \mathbb{R}^n$

$$f(\Lambda) = \det(\hat{\Lambda}) \hat{f}(\hat{\Lambda})$$

and, more generally, $f(x + \Lambda) = \det(\hat{\Lambda}) \sum_{v \in \hat{\Lambda}} \hat{f}(v) \chi_v(v, x)$

Proof. The Fourier transform of the periodic function $\varphi(x) = f(x + \Lambda)$ is

$$\det(\hat{\Lambda}) \int_{x \in \mathbb{R}^n / \Lambda} f(x + \Lambda) \overline{\chi_v(x + \Lambda)} dx = \det(\hat{\Lambda}) \int_{x \in \mathbb{R}^n} f(x) \overline{\chi_v(x)} dx$$

This proves that $\hat{\varphi}(v) = \det(\hat{\Lambda}) \hat{f}(v)$. Then, use the Fourier inversion formula, and set $x = 0$. $\square$

- Informal: amount of (gaussian) noise s that needs to be added to Λ to get $\approx \mathbb{R}^n$
- Notation $\eta_\varepsilon(\Lambda)$, parametrized by $\varepsilon > 0$, usually $\varepsilon \approx 0$.

Definition 6. (Smoothing parameter) $\eta_\varepsilon(\Lambda)$ is the smallest $s > 0$ such that $\rho(s \hat{\Lambda}) \leq 1 + \varepsilon$.

- $\rho(\hat{\Lambda})$ may be large, but if we scale up the lattice $\lim_{s \rightarrow \infty} \rho(s \hat{\Lambda}) = \rho(\{0\}) = 1$
- Note $\eta_\varepsilon(c \Lambda) = c \eta_\varepsilon(\Lambda)$ for all $c > 0$
- We say that Λ is ε -smooth if $\eta_\varepsilon(\Lambda) \leq 1$.
- $s \geq \eta_\varepsilon(\Lambda)$ if the scaled lattice Λ/s is smooth

Theorem 7. If Λ is ε -smooth, then $\rho(\Lambda + u) = (1 \pm \varepsilon) \det(\hat{\Lambda})$

Theorem. *If Λ is ε -smooth, then $\rho(\Lambda + u) = (1 \pm \varepsilon) \det(\hat{\Lambda})$*

- Assume $\rho(\hat{\Lambda}) \leq 1 + \varepsilon$. Goal: $\rho(\Lambda + u) \approx \det(\hat{\Lambda})$ up to relative error ε
- Note that $\widehat{f(\cdot + u)} = \hat{f}(\cdot) \chi_u(\cdot)$ for any f
- Using Poisson summation formula and triangle inequality

$$\frac{|\rho(\Lambda + u) - \det(\hat{\Lambda})|}{\det(\hat{\Lambda})} = \left| \sum_{v \in \hat{\Lambda}} \rho(v) \chi_u(v) - 1 \right| \leq \sum_{v \in \hat{\Lambda} \setminus \{0\}} \rho(v) |\chi_u(v)| = \rho(\hat{\Lambda}) - 1 \leq \varepsilon$$

- $\eta_\varepsilon(\mathbb{Z}) \approx \sqrt{\frac{\ln(2/\varepsilon)}{\pi}} = \omega(\sqrt{\log n})$ for negligible $\varepsilon = 1/n^{\omega(1)}$
- For orthogonal sums $\Lambda_1 \oplus \Lambda_2$, $\widehat{(\Lambda_1 \oplus \Lambda_2)} = \widehat{\Lambda_1} \oplus \widehat{\Lambda_2}$ and $\rho(\widehat{\Lambda_1 \oplus \Lambda_2}) = \rho(\widehat{\Lambda_1}) \rho(\widehat{\Lambda_2})$
 - $\eta(\Lambda_1 \oplus \Lambda_2) \approx \max(\eta(\Lambda_1), \eta(\Lambda_2))$
- Gaussian sums on cosets are maximized at $\rho(\Lambda) \geq \rho(\Lambda + u)$
 - For any lattice basis B , $\eta_\varepsilon(B) \leq \eta_\varepsilon(B^*)$
- Tail bounds and more (see lecture notes)
 - For any halfspace $0 \in \mathcal{H}$, $\rho(\Lambda \setminus \mathcal{H}) \leq \rho(h) \rho(\Lambda - h)$ where $h = \operatorname{argmin}\{\|h\| : h \in \mathcal{H}\}$
 - $\rho(\Lambda \setminus \mathcal{B}(\sqrt{n})) \leq 2^{-n} \rho(\Lambda)$

Theorem 8. For any $\Lambda \subset \mathbb{R}^n$, $\eta_\varepsilon(\Lambda) \leq \omega(\sqrt{\log n}) \lambda_n^*(\Lambda)$ [Gram-Schmidt $\lambda_n^* \leq \lambda_n$]

Theorem 9. For any $\Lambda \subset \mathbb{R}^n$, $\eta_\varepsilon(\Lambda) \leq \frac{\sqrt{n}}{\lambda_1(\hat{\Lambda})}$ [Dual distance $\hat{\lambda}_1$]

Theorem 10. $\eta_\varepsilon(\Lambda) \geq \mu(\Lambda) / \sqrt{n}$ [Covering radius μ]

Combining the last two bounds we get

Theorem 11. (Banaszczyk's transference theorem) For any lattice $\Lambda \subset \mathbb{R}^n$

$$1 \leq 2 \mu(\Lambda) \lambda_1(\hat{\Lambda}) \leq n$$

- Bounded Distance Decoding: $(B, \vec{t} = B^* \vec{y})$ such that $\delta = \text{dist}(\vec{t}, \mathcal{L}(B))$ is small
 - Assume we have a relatively good basis, with large $\beta = \min_i \|b_i^*\|$
 - Using (Babai) nearest plane algorithm, we can solve BDD up to distance $\delta < \beta/2$
- Consider a randomized version of the nearest plane algorithm (Klein/GPV):
 - Nearest plane: round y_n to closest integer c_n and subtract $c_n b_n$ from target
 - Instead: $c_n \leftarrow D_{\mathbb{Z}, y_n, s_n}$ where $s_n = \sigma / \|b_n^*\| \geq \eta_\epsilon(\mathbb{Z})$ for $\sigma = \eta_\epsilon(\mathbb{Z}) \|B^*\|$
 - Squared distance: $\delta = \sum_i (c_i - y_i)^2 \|b_i^*\|^2$ with probability

$$\prod_i \frac{\rho((c_i - y_i) / s_i)}{\rho((\mathbb{Z} - y_i) / s_i)} \approx \frac{\prod_i \rho((c_i - y_i) / s_i)}{\prod_i \rho(\mathbb{Z} / s_i)} = \frac{\rho(\delta / \sigma)}{C}$$