Exercise 1 (Reducing Approximate SVP to Hermite-SVP)
For $\gamma := \gamma(n)$, recall that the γ-SVP and γ-HermiteSVP problems on an n-dimensional lattice $\mathcal{L}$ are to find a non-zero vector $\mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}$ such that $\|\mathbf{y}\| \leq \gamma \cdot \lambda_1(\mathcal{L})$ and $\|\mathbf{y}\| \leq \gamma \cdot \det(\mathcal{L})^{1/n}$ respectively. Note that γ-HermiteSVP is only solvable when $\gamma(n) \geq \gamma_n^{1/2}$, the Hermite constant in dimension n. The goal of this exercise is to show that γ^2-SVP reduces to γ-HermiteSVP. The reduction will proceed via the following algorithm:

Algorithm 1: γ^2-SVP to γ-HermiteSVP reduction

Input : n-dimensional lattice $\mathcal{L} \subseteq \mathbb{R}^n$.
Output: $\mathbf{y} \in \mathcal{L} \setminus \{0\}$ satisfying $\|\mathbf{y}\| \leq \gamma^2(n)\lambda_1(\mathcal{L})$.

$\mathcal{L}_n := \mathcal{L}$
for $i = n$ **downto** 1 **do**
  $\mathbf{y}_i := \gamma$-HermiteSVP$(\mathcal{L}_i)$
  $\mathbf{y}_i^* := \gamma$-HermiteSVP$(\mathcal{L}_i^*)$
  $\mathcal{L}_{i-1} := \mathcal{L}_i \cap \text{span}(\mathbf{y}_i^*)^\perp$
end
return shortest vector in $\mathbf{y}_1, \ldots, \mathbf{y}_n$.

Prove that the above algorithm is correct. That is, assuming that all the calls to γ-HermiteSVP return correctly, that it outputs at γ^2 approximation to the shortest vector of $\mathcal{L}$.
(Hint: Show that $\|\mathbf{y}_i\|\|\mathbf{y}_i^*\| \leq \gamma(n)^2$. How short can a vector in $\mathcal{L}_i$ but not in $\text{span}(\mathbf{y}_i)^\perp$ be?)

Exercise 2 (Similar Distributions)
Assume that $\eta_\varepsilon(\mathcal{L}) \leq 1$ for $\varepsilon \in (0,1)$. Examine the following distributions:

1. $\mathbf{X} \sim D_{\mathbb{R}^n}$.
2. Sample $\mathbf{U}$ uniformly from $\mathbb{R}^n/\mathcal{L}$. Then sample $\mathbf{Y}$ as $D_{\mathcal{L}+\mathbf{U}}$.

Prove that for any measurable $A \subseteq\in \mathbb{R}^n$, that

$$\frac{1-\varepsilon}{1+\varepsilon}\Pr[\mathbf{Y} \in A] \leq \Pr[\mathbf{X} \in A] \leq \frac{1+\varepsilon}{1-\varepsilon}\Pr[\mathbf{Y} \in A].$$

Exercise 3 (Covariance of Discrete Gaussians)
Let $\mathcal{L} \subset \mathbb{R}^n$ be a full-rank lattice. Let $\mathbf{X} \sim D_{\mathcal{L}}$ and $\mathbf{Y} \sim D_{\mathcal{L}^*}$. Let $\mathbf{v} \in \mathbb{R}^n$, $\|\mathbf{v}\| = 1$, and $\varepsilon \in (0,1)$.

1. Show that $\mathbb{E}[\langle \mathbf{X}, \mathbf{v}\rangle^2] + \mathbb{E}[\langle \mathbf{Y}, \mathbf{v}\rangle^2] = \frac{1}{2\pi}$.
   (Hint: Combine the derivatives $f(\mathbf{t}) := \rho(\mathcal{L} + \mathbf{t})$ in the direction of $\mathbf{v}$, together with the Poisson summation formula, to deduce the identity).
2. Show that if $\rho(\mathcal{L}^*) = 1 + \varepsilon$, then

$$\mathbb{E}[\langle \mathbf{Y}, \mathbf{v}\rangle^2] \leq \frac{1}{2\pi}\frac{\varepsilon}{1+\varepsilon}\left(1 + \ln\left(\frac{2(1+\varepsilon)}{\varepsilon}\right)\right).$$

   (Hint: First show that $\mathbb{E}[\langle \mathbf{Y}, \mathbf{v}\rangle^2] = \int_0^\infty 2s \Pr[|\langle \mathbf{Y}, \mathbf{v}\rangle| \geq s]ds$. What bound do you get on $\Pr[|\langle \mathbf{Y}, \mathbf{v}\rangle| \geq s]$ from $\rho(\mathcal{L}^*) = 1 + \varepsilon$? From the 1D discrete Gaussian tailbound? Combine the information from both to derive the above bound.)

3. Conclude that if $\eta_\varepsilon(\mathcal{L}) \leq 1$, then

$$\frac{1}{2\pi}(1 - O(\varepsilon \ln(1/\varepsilon))) \leq \mathbb{E}[\langle \mathbf{X}, \mathbf{v} \rangle^2] \leq \frac{1}{2\pi}.$$

Exercise 4 (Upper Bounding the Empirical Covariance)
Let $\mathcal{L} \subset \mathbb{R}^n$ ben an n-dimensional lattice. Prove the following:

1. For $\mathbf{v} \in \mathbb{R}^n$, $\|\mathbf{v}\| = 1$, $\lambda < 1$, $\mathbb{E}_{\mathbf{X} \sim D_\mathcal{L}}[e^{\pi\lambda\langle \mathbf{X}, \mathbf{v}\rangle^2}] \leq \frac{1}{\sqrt{1-\lambda}}$.
   (Hint: Adapt the proof that $\rho_s(\mathcal{L}) \leq s^n \rho(\mathcal{L})$ to the "1D dimensional" setting.)

2. Let $\mathbf{v} \in \mathbb{R}^n$, $\|\mathbf{v}\| = 1$, $\varepsilon > 0$, $\mathbf{X}_1, \ldots, \mathbf{X}_N \sim D_\mathcal{L}$ i.i.d.

   (a) Show that

$$\Pr[\sum_{i=1}^N \langle \mathbf{X}_i, \mathbf{v} \rangle^2 \geq N\frac{1+\varepsilon}{2\pi}] \leq \min_{\lambda \in (0,1)} \sqrt{1-\lambda}^{-N} e^{N(1+\varepsilon)/2}.$$

   (Hint: Use the previous part together with Markov's inequality.)

   (b) Optimize over λ to deduce that

$$\Pr[\sum_{i=1}^N \langle \mathbf{X}_i, \mathbf{v} \rangle^2 \geq N\frac{1+\varepsilon}{2\pi}] \leq e^{-\frac{N}{2}(\varepsilon - \ln(1+\varepsilon))} = e^{-\frac{N}{2}\Omega(\min\{\varepsilon, \varepsilon^2\})}.$$

3. Use a packing bound to show that there exists $N_\varepsilon \subseteq \partial\mathcal{B}_2^n$ (unit sphere), $|N_\varepsilon| = (1 + \frac{1}{\varepsilon})^n$, such that $\forall \mathbf{v} \in \partial\mathcal{B}_2^n$, $\exists \mathbf{w} \in N_\varepsilon$ such that $\|\mathbf{w} - \mathbf{v}\| \leq \varepsilon$.

4. Let $\mathbf{x}_1, \ldots, \mathbf{x}_N \in \mathbb{R}^n$ be vectors and assume that $\forall \mathbf{w} \in N_\varepsilon$, that $(\frac{1}{N}\sum_{i=1}^N \langle \mathbf{x}_i, \mathbf{w} \rangle^2)^{1/2} \leq 1 + \varepsilon$. Show that this implies that $\forall \mathbf{v} \in \mathbb{R}^n, \|\mathbf{v}\| = 1$, that $(\frac{1}{N}\sum_{i=1}^N \langle \mathbf{x}_i, \mathbf{v} \rangle^2)^{1/2} \leq \frac{1+\varepsilon}{1-\varepsilon}$.
   (Hint: Use the triangle inequality)

5. Deduce that for $\varepsilon \in (0,1)$, $N \geq \Omega(\frac{n}{\varepsilon^2} \ln \frac{1+\varepsilon}{\varepsilon})$, $\mathbf{X}_1, \ldots, \mathbf{X}_N \sim D_\mathcal{L}$ i.i.d., that

$$\Pr[\forall\ \mathbf{v} \in \mathbb{R}^n, \|\mathbf{v}\| = 1, \frac{1}{N}\sum_{i=1}^N \langle \mathbf{X}_i, \mathbf{v} \rangle^2 \leq \frac{1+\varepsilon}{2\pi}] \geq 1 - 2^{-n}.$$

   (Hint: First show this holds for say $\varepsilon/10$ over the net $N_{\varepsilon/10}$ via the union bound, and then extend this to the hold sphere using the previous part.)