Continuous Learning with Errors

Oded Regev (NYU)

Joint work with Joan Bruna, Min Jae Song (NYU)
and Yi Tang (Umich)
June 14th 2021, Simons Institute

Failed attempts to solve lattice problems

- Learning with Errors (LWE) arose as a failed attempt to design a quantum algorithm for solving lattice problems
- The quantum algorithm had a gap
- That “gap” is LWE
- So instead of an algorithm, we got a hardness reduction

- In this work, we (slightly) modify the reduction to obtain a different “gap”, which we call Continuous LWE (CLWE).

Motivation: Gaussian pancakes

Is this the standard Gaussian distribution?

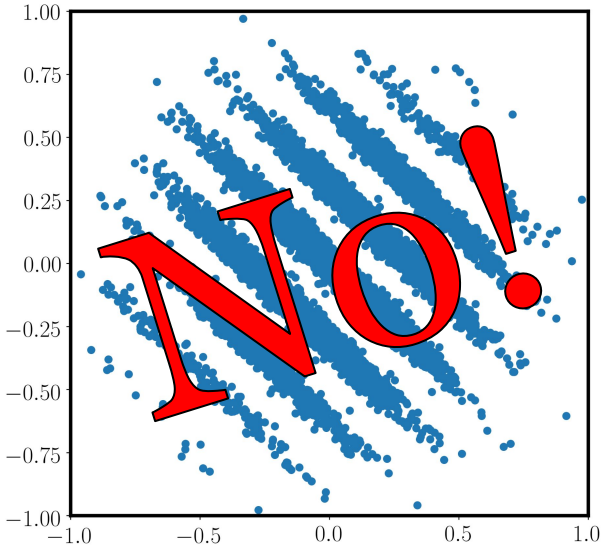

Gaussian pancakes

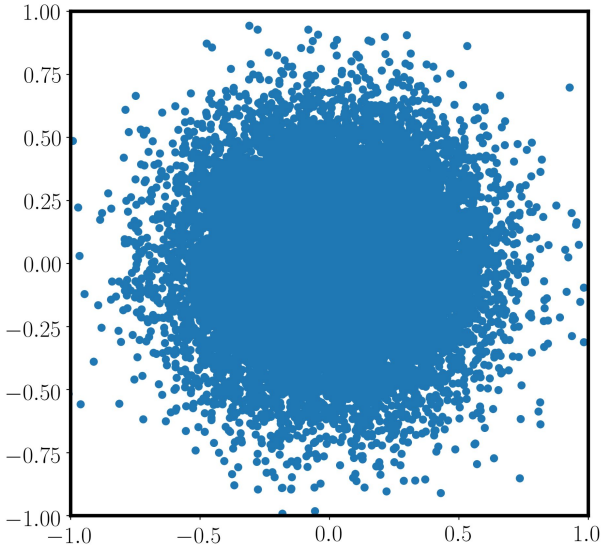

Standard Gaussian

Motivation: Gaussian pancakes

But what if you hide this discrete direction in higher dimensions?

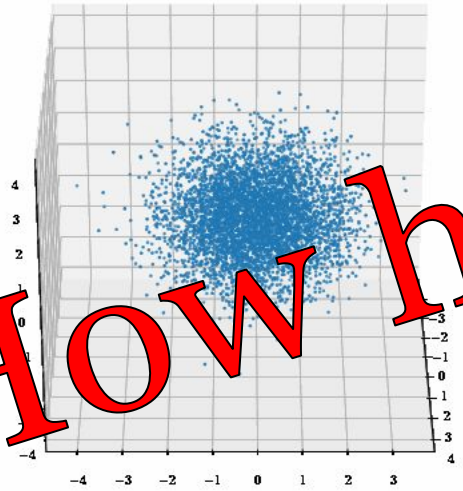

Gaussian pancakes!

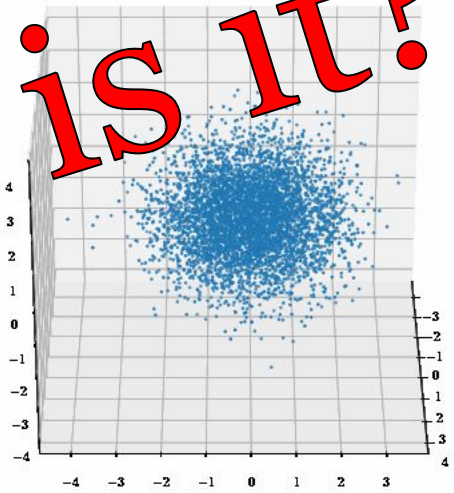

Standard Gaussian

How hard is it?

Gaussian pancakes

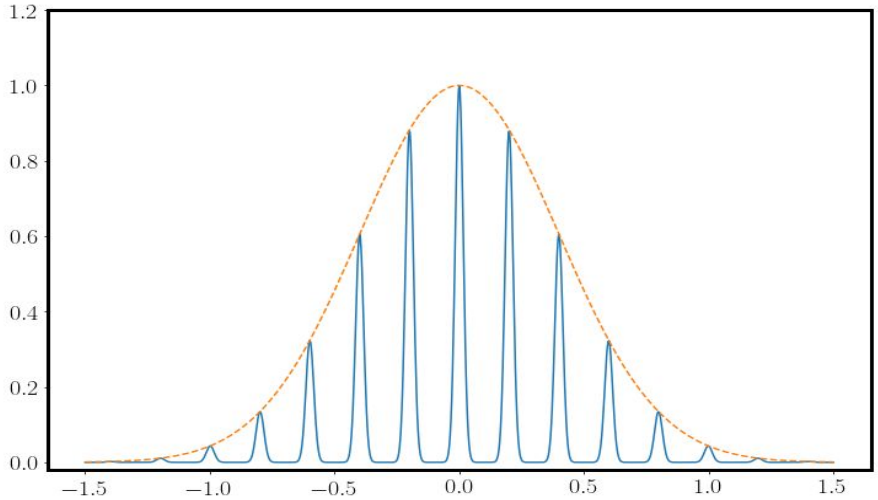

The Gaussian pancakes distribution is a noisy discrete Gaussian (blue) in one hidden direction. In other $n-1$ directions, the distribution is Gaussian (orange).

Spoiler: this is the **homogeneous Continuous Learning with Errors (hCLWE)** distribution.

SQ-hardness of distinguishing Gaussian pancakes [DiakonikolasKaneStewart17]

Thm [DKS17]: Distinguishing Gaussian pancakes from the standard Gaussian is hard for **statistical query (SQ) algorithms**.

Corollary [DKS17]**.** Improperly learning (=density estimation) mixtures of Gaussians is hard for SQ algorithms, even when the components are nearly non-overlapping and even when parameter recovery is info-theoretically possible with poly(n) samples.

Open question [BubeckLeePriceRazenshteyn19]**:**

Is detecting the pancake structure computationally hard for **any** algorithm?

We resolve this here in the **affirmative**.

Extension to multiple discrete directions: Gaussian baguettes

Thm [BLPR19]: SQ-hard also with **multiple** (up to $O(n)$) discrete directions (Gaussian "baguettes").

Proof: By hybrid argument.

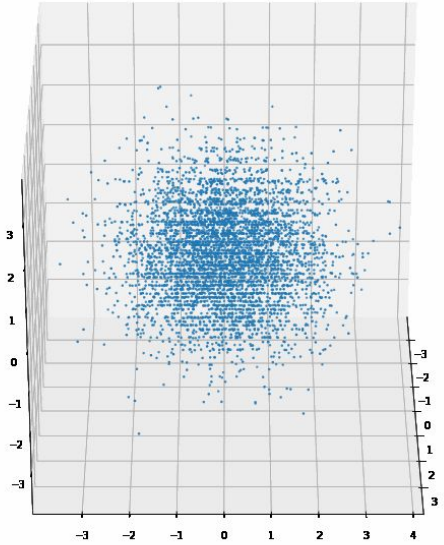

Application: adversarial examples and robust classifiers

[SzegedyZarembaSutskeverBrunaErhanGoodfellowFergus13,...]

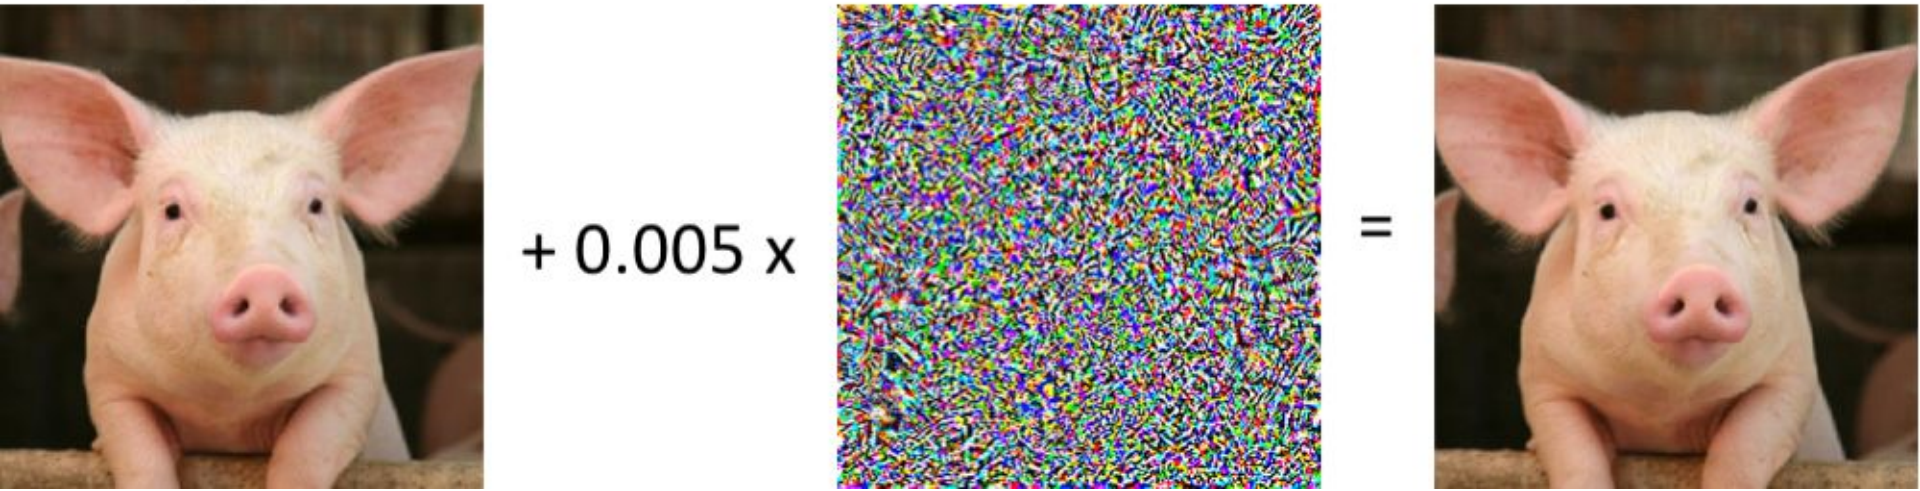

Corollary: Learning a classifier can be hard, even when **robust classifiers** exist and are learnable info-theoretically (and even when non-robust classifiers can be learned efficiently)

([BLPR19] show SQ hardness; we show computational).

Our goal

Prove that distinguishing the following two distributions (in high dimension) is computationally hard.

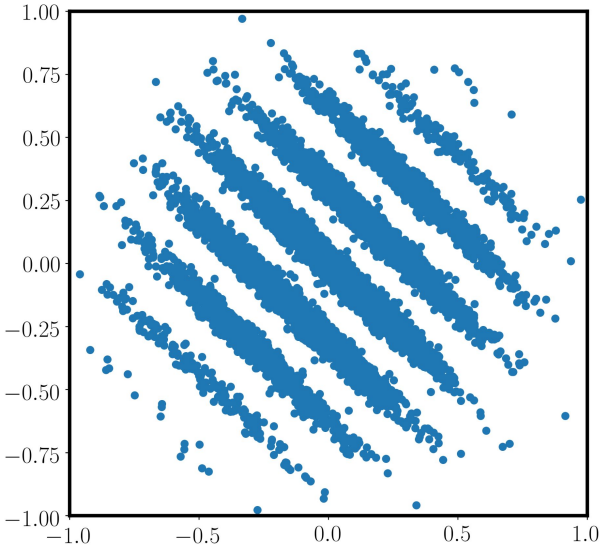

Gaussian pancakes

vs

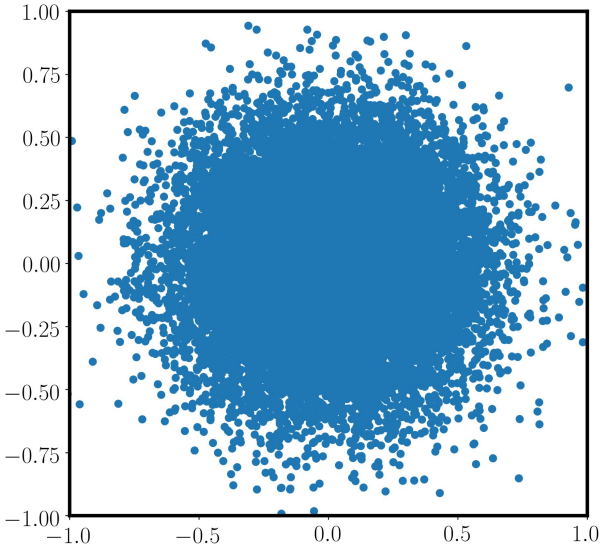

Standard Gaussian

Gaussian pancakes

γ : ~number of pancakes (or inverse of pancake spacing)

β : pancake (relative) thickness

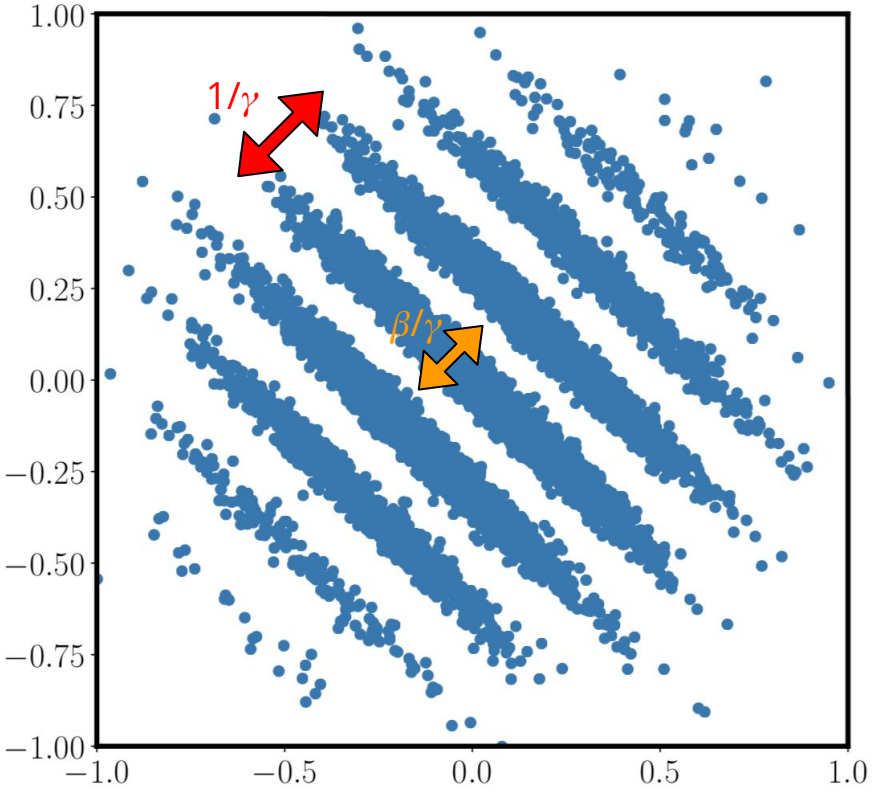

Our result: hardness of Gaussian pancakes

- Thm [BRST21]: Distinguishing Gaussian pancakes with spacing $1/\gamma < n^{-1/2}$ from the standard Gaussian is hard

(even with any non-negligible advantage, and assuming there are no poly-time quantum algorithms for worst-case lattice problems; β can be any inverse polynomial.)

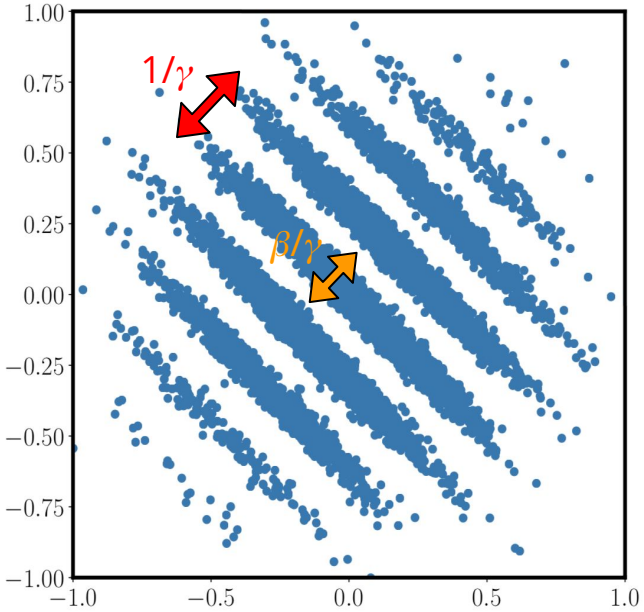

Implications of our hardness result

Assuming some worst-case lattice problems cannot be solved by polynomial-time quantum algorithms …

- Distinguishing Gaussian pancakes/baguettes from the standard Gaussian is hard for **any** polynomial-time algorithm.
- Improperly learning mixtures of Gaussians can be computationally hard even when the mixture components are nearly non-overlapping.

Remark: Hardness of improper learning results are generally rare (because there is no restriction on what hypothesis the learning algorithm can output). One notable example is [KlivansSherstov06], also based on lattice problems.

Hardness of (h)CLWE: proof overview

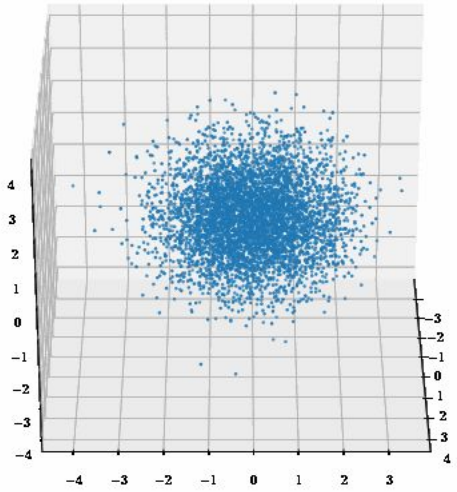

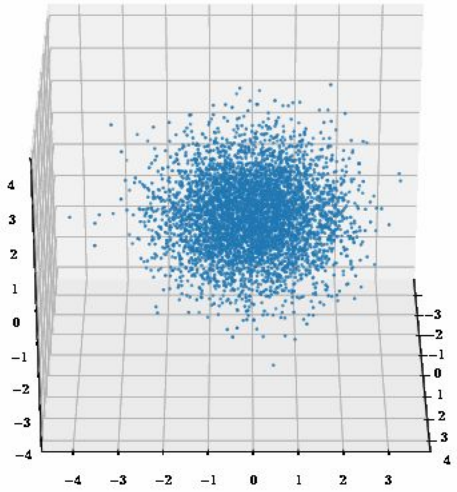

Hardness of (h)CLWE: proof overview

We actually prove a stronger hardness result, for a relaxed problem: (*inhomogeneous*) CLWE.

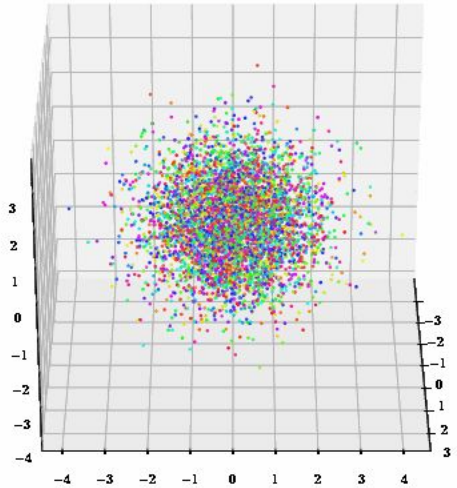

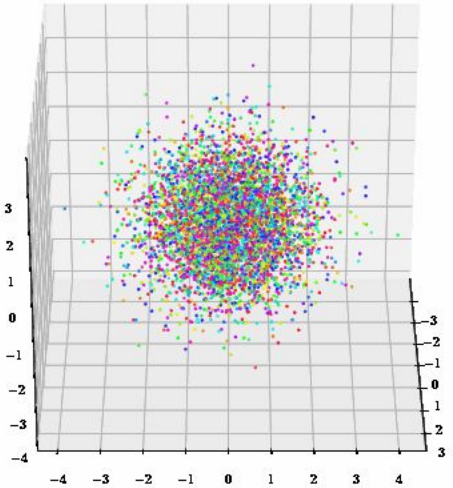

Continuous LWE (CLWE)

Def: (β,γ)-CLWE: Decide whether given samples of the form $(\mathbf{y},z)$ with $\mathbf{y} \sim \mathcal{N}(0,I_n)$ have either:

(1) periodic "colors" z along some secret direction $\mathbf{w} \in \mathbb{R}^n$, i.e., $z = (\gamma\langle\mathbf{y},\mathbf{w}\rangle + e) \bmod 1$ where $e \sim \mathcal{N}(0,\beta)$, or

(2) uniformly random "colors" $z \in [0,1)$.

hCLWE samples are essentially CLWE samples conditioned on $z \approx 0$. It therefore suffices to prove hardness of CLWE.

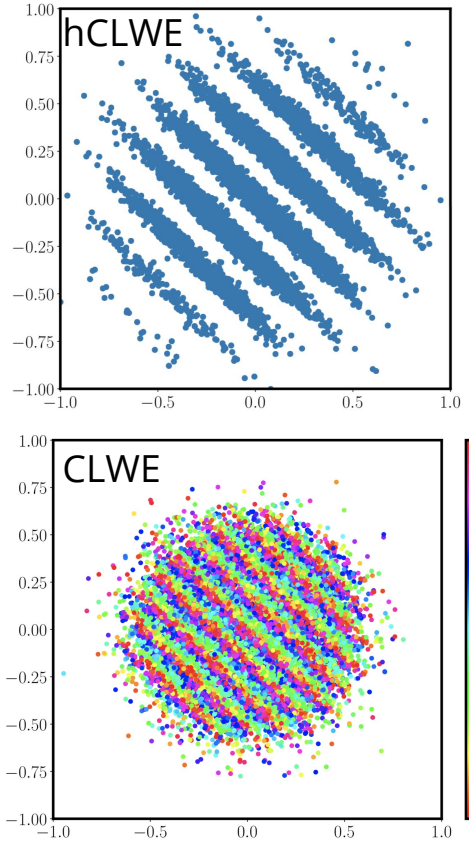

Hardness proof

We follow [R05,PeikertRStephens-Davidowitz17].

The core of the argument is showing how to solve Bounded Distance Decoding (BDD) in a lattice L given

1. CLWE oracle, and
2. Samples from a discrete Gaussian distribution over the dual L^*

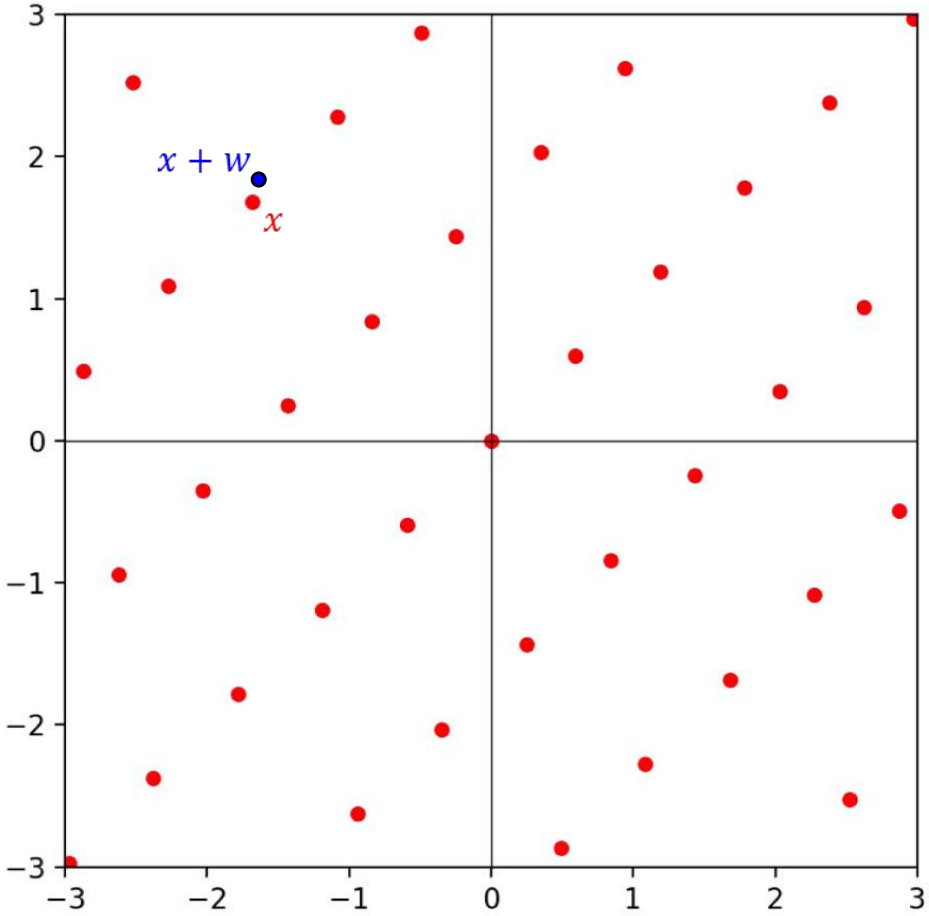

L

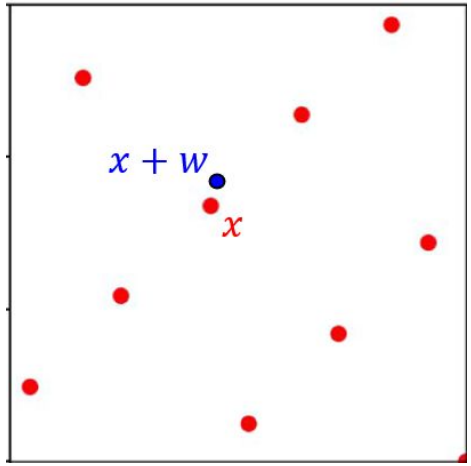

$y \sim L$

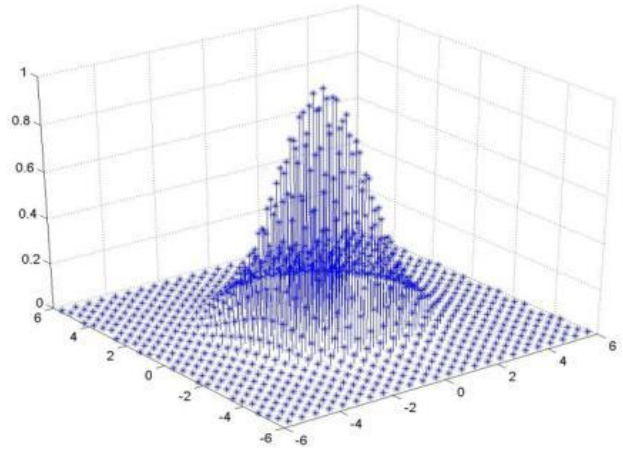

$$\langle y, x + w \rangle = \langle y, x \rangle + \langle y, w \rangle = \langle y, w \rangle \, mod \, 1$$

We can therefore use

$$(y + e_1, z = \langle y, x + w \rangle + e_2 \, mod \, 1)$$

as input to the CLWE oracle, thereby recovering w.

Learning with Errors (LWE)

Def: (α,q)-LWE: decide whether given samples of the form $(\mathbf{a},b)$ with $\mathbf{a} \sim (\mathbb{Z}/q\mathbb{Z})^n$ have either:

(1) periodic b along some secret direction $\mathbf{s} \in (\mathbb{Z}/q\mathbb{Z})^n$, i.e., $b = (\langle \mathbf{a},\mathbf{s} \rangle/q + e) \bmod 1$ where $e \sim \mathcal{N}(0,\alpha)$, or

(2) uniformly random $b \in [0,1)$.

Remark: By discretizing $b' = \lfloor q \cdot b \rfloor \in \mathbb{Z}/q\mathbb{Z}$, we obtain the more common definition of LWE. Specifically, the search version (to find secret $\mathbf{s}$ given periodic b) can be viewed as solving a system of linear equations with errors over $\mathbb{Z}/q\mathbb{Z}$, of the form $\langle \mathbf{a},\mathbf{s} \rangle \approx b'$.

Analogies between CLWE and LWE

(β,γ)-CLWE	(α,q)-LWE
secret $\mathbf{w} \in \mathbb{R}^n$, $\|\mathbf{w}\| = 1$	secret $\mathbf{s} \in (\mathbb{Z}/q\mathbb{Z})^n$
samples $(\mathbf{y},z)$	samples $(\mathbf{a},b)$
$\mathbf{y} \sim \mathcal{N}(0,I_n)$	$\mathbf{a} \sim (\mathbb{Z}/q\mathbb{Z})^n$
$z = (\gamma\langle\mathbf{y},\mathbf{w}\rangle + e) \bmod 1$ where $e \sim \mathcal{N}(0,\beta)$	$b = (\langle\mathbf{a},\mathbf{s}\rangle/q + e) \bmod 1$ where $e \sim \mathcal{N}(0,\alpha)$
reduce from $O(n/\beta)$-GapSVP for $\gamma \geq O(n^{1/2})$	reduce from $O(n/\alpha)$-GapSVP for $\alpha \cdot q \geq O(n^{1/2})$
noise rate β	noise rate α
inverse period γ	absolute noise $\alpha \cdot q$

Analogies between CLWE and LWE (cont.)

(β,γ)-CLWE	(α,q)-LWE
noiseless (β=0) is easy to solve by **LLL**	noiseless (α=0) is easy to solve by **Gaussian elimination**

Open Question 1

- Better algorithms for CLWE?

Open Question 2

- What happens with $< \sqrt{n}$ pancakes?
  - Subexponential algorithms are known
  - For instance, with $n^{1/4}$ pancakes, $\exp(n^{1/2})$ samples & time algorithm exists
  - Can we reduce the number of samples to $poly(n)$?
    - (Might shed light on hardness of low-sample LWE, important for cryptography)

Open Question 3

- Cryptographic applications?

Open Question 4

- Does hardness still hold for other moment-matching distributions in the hidden direction?
  - [BLPR19] used a slightly different distribution
- Alternatively, are there better algorithms for their distribution?