Hardness and Algorithms for Batch LPN under Dependent Noise

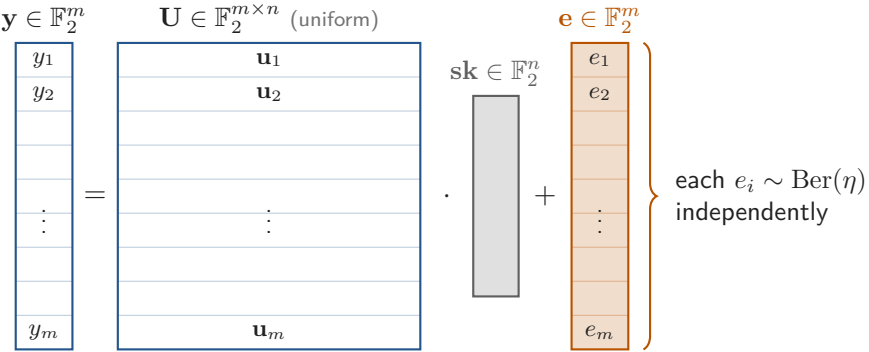
Xin Li[1] Songtao Mao[1] Zhaienhe Zhou[2]

[1]Johns Hopkins University [2]University of Michigan

TCC 2026

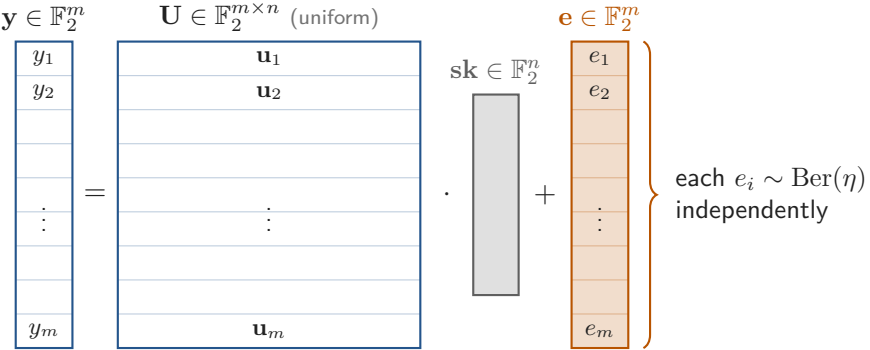
Learning Parity with Noise (LPN)

m: number of samples
n: dimension of the LPN problem

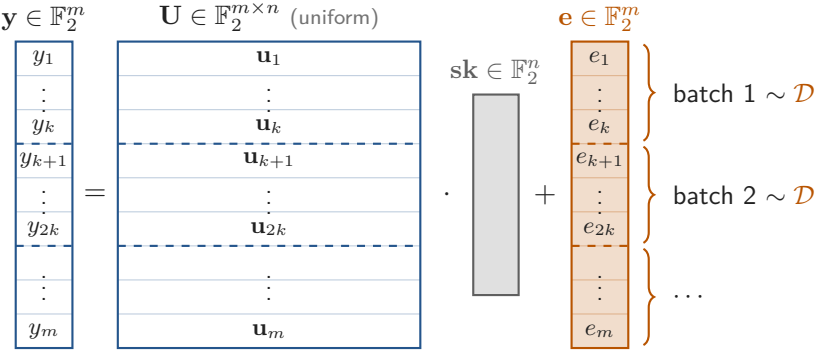


- $\mathrm{LPN}_{n,\eta}(\mathbf{sk})$: given $(\mathbf{U},\ \mathbf{y} = \mathbf{U} \cdot \mathbf{sk} + \mathbf{e})$, recover $\mathbf{sk} \in \mathbb{F}_2^n$.
- For constant η, best algorithm takes $2^{O(n/\log n)}$ time. [Blum–Kalai–Wasserman'03]

Batch LPN


$\mathbf{y} \in \mathbb{F}_2^m$
$\mathbf{U} \in \mathbb{F}_2^{m \times n}$ (uniform)
$\mathbf{e} \in \mathbb{F}_2^m$
$\mathbf{sk} \in \mathbb{F}_2^n$
y_1
y_2
y_m
$\mathbf{u}_1$
$\mathbf{u}_2$
$\mathbf{u}_m$
e_1
e_2
e_m
each $e_i \sim \mathrm{Ber}(\eta)$
independently

Batch LPN



- $\mathrm{LPN}_{n,k,\mathcal{D}}(\mathbf{sk})$: batches of k samples.
- Noise correlated inside each batch: $(e_1, \ldots, e_k) \sim \mathcal{D}$ (known); i.i.d. across batches.
- Standard $\mathrm{LPN}_{n,\eta}$ is the case $\mathcal{D} = \mathrm{Ber}(\eta)^{\otimes k}$.

Question: what $\mathcal{D}$ makes Batch LPN hard?

Motivation

- Real noise is rarely i.i.d. [Bellizia et al.'21]
- Correlations can break hardness: $\leq k/3$ errors per batch $\Rightarrow n^{O(k)}$ time. [Arora–Ge'11]
- Dependent noise separates RL from supervised learning. [Golowich–Moitra–Rohatgi'24]

Question: what $\mathcal{D}$ makes Batch LPN hard?

Motivation

- Real noise is rarely i.i.d. [Bellizia et al.'21]
- Correlations can break hardness: $\leq k/3$ errors per batch $\Rightarrow n^{O(k)}$ time. [Arora–Ge'11]
- Dependent noise separates RL from supervised learning. [Golowich–Moitra–Rohatgi'24]

Prior work

Easy: structured noise [AG'11]

Fixed nonzero polynomial P of degree d:

$P(\mathbf{e}) = 0$ for every $\mathbf{e} \in \mathrm{supp}(\mathcal{D})$

$\Rightarrow n^{O(d)}$ time (linearization)

?

Hard: SV source [GMR'24]

Santha–Vazirani: each noise bit is weakly dependent on the previous ones:

$\Pr[e_i = 1 \mid e_1, \ldots, e_{i-1}] \in \frac{1}{2} \pm \delta$

$\delta \lesssim 2^{-k}(1-2\eta) \Rightarrow$ as hard as $\mathsf{LPN}_{n,\eta}$

Our results

Fourier coefficient of $\mathcal{D}$: $\widehat{P}_{\mathcal{D}}(\mathbf{s}) = \mathbb{E}_{\mathbf{e}\sim\mathcal{D}}\left[(-1)^{\langle \mathbf{s},\, \mathbf{e}\rangle}\right]$

	Hard: as hard as $\mathsf{LPN}_{n,\eta}$	
High noise η near $\frac{1}{2}$	**Fourier-bounded** (implies GMR's SV) $\sum_{\mathbf{s}\neq\mathbf{0}} \lvert\widehat{P}_{\mathcal{D}}(\mathbf{s})\rvert \le 1-2\eta$ δ-**SV source** (GMR: $\delta \lesssim 2^{-k}(1-2\eta)$) $\delta \lesssim 2^{-k/2}(1-2\eta)$ via Parseval	
Low noise $\eta \ll 1/k$	**Dense** $P_{\mathcal{D}}(\mathbf{z}) \gtrsim k\eta\, 2^{-k}$ for every $\mathbf{z}$	

Reductions take $\mathrm{poly}(2^k, n, m)$ time; $\Delta = d_{\mathrm{TV}}(\mathcal{D}, \mathrm{Unif})$.

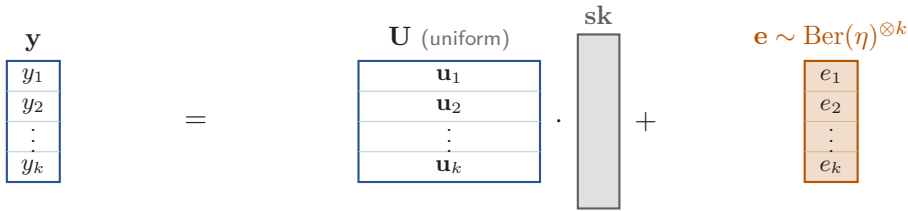
Our results

Fourier coefficient of $\mathcal{D}$: $\widehat{P}_{\mathcal{D}}(\mathbf{s}) = \mathbb{E}_{\mathbf{e}\sim\mathcal{D}}\left[(-1)^{\langle \mathbf{s},\, \mathbf{e}\rangle}\right]$

	Hard: as hard as $\mathsf{LPN}_{n,\eta}$	**Easy: efficient algorithm**
High noise η near $\frac{1}{2}$	**Fourier-bounded** (implies GMR's SV) $\sum_{\mathbf{s}\neq\mathbf{0}} \lvert\widehat{P}_{\mathcal{D}}(\mathbf{s})\rvert \le 1-2\eta$ δ-**SV source** (GMR: $\delta \lesssim 2^{-k}(1-2\eta)$) $\delta \lesssim 2^{-k/2}(1-2\eta)$ via Parseval	**Large Fourier coefficient** $\lvert\widehat{P}_{\mathcal{D}}(\mathbf{s})\rvert = 1-\gamma$ for some $\mathbf{s}\neq\mathbf{0}$ $\Rightarrow$ reduces to $\mathsf{LPN}_{n,\gamma/2}$, time $e^{O(\gamma n)}$
Low noise $\eta \ll 1/k$	**Dense** $P_{\mathcal{D}}(\mathbf{z}) \gtrsim k\eta\, 2^{-k}$ for every $\mathbf{z}$	**Rare point** (AG: $\tau = 0$) $P_{\mathcal{D}}(\mathbf{z}_0) \le \tau$ for some $\mathbf{z}_0$ $\Rightarrow$ time $e^{O(\tau n^k 4^k)} \cdot \mathrm{poly}(n^k, 1/\Delta)$

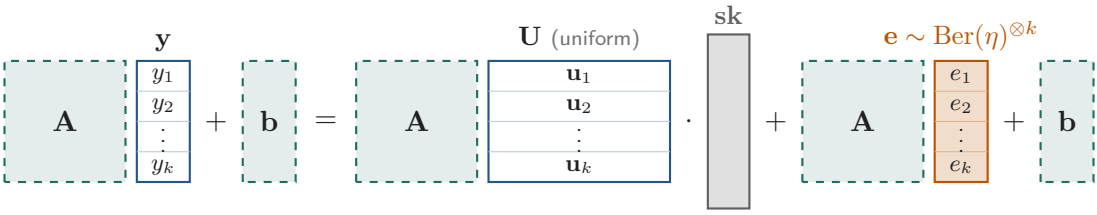
Reductions take $\mathrm{poly}(2^k, n, m)$ time; $\Delta = d_{\mathrm{TV}}(\mathcal{D}, \mathrm{Unif})$.

Simulate Batch LPN with LPN: affine reduction



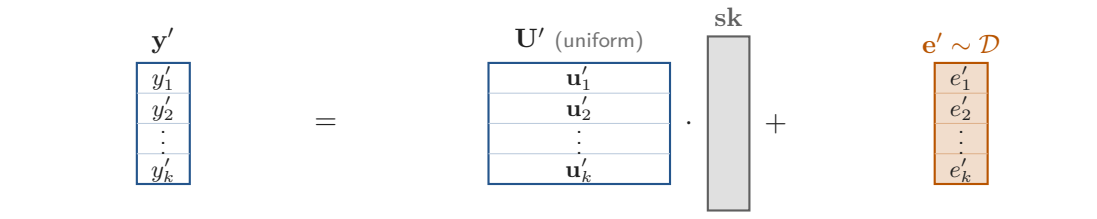
- Take k samples from $\mathsf{LPN}_{n,\eta}(\mathbf{sk})$.

Simulate Batch LPN with LPN: affine reduction



- Take k samples from $\mathsf{LPN}_{n,\eta}(\mathbf{sk})$.
- Sample random $(\mathbf{A} \in \mathrm{GL}(k,2), \mathbf{b} \in \mathbb{F}_2^k)$ from a distribution we design.

Simulate Batch LPN with LPN: affine reduction



- Take k samples from $\mathsf{LPN}_{n,\eta}(\mathbf{sk})$.
- Sample random $(\mathbf{A} \in \mathrm{GL}(k,2), \mathbf{b} \in \mathbb{F}_2^k)$ from a distribution we design.
- $\mathbf{U}' = \mathbf{A}\mathbf{U}$ uniform and independent of $\mathbf{e}'$
- Goal: make $\mathbf{e}' = \mathbf{A}\mathbf{e} + \mathbf{b} \sim \mathcal{D}$.

$\Rightarrow$ a sample of $\mathsf{LPN}_{n,k,\mathcal{D}}(\mathbf{sk})$

Fourier view of the affine reduction

Adding $\mathbf{b}$ shifts the whole distribution. While in Fourier, it is only a pointwise product.

Recall

- Fourier Transform: $\widehat{P}_{\mathcal{D}}(\mathbf{s}) = \mathbb{E}_{\mathbf{e}\sim\mathcal{D}}\left[(-1)^{\langle \mathbf{s}, \mathbf{e}\rangle}\right]$
- Inverse Fourier Transform: $P_{\mathcal{D}}(\mathbf{x}) = \frac{1}{2^k}\sum_{\mathbf{s}\in\mathbb{F}_2^k} \widehat{P}_{\mathcal{D}}(\mathbf{s})(-1)^{\langle \mathbf{s}, \mathbf{x}\rangle}$
- $\widehat{P}_{\mathrm{Ber}(\eta)^{\otimes k}}(\mathbf{s}) = (1-2\eta)^{\mathrm{wt}(\mathbf{s})}$

Fourier view of the affine reduction

Adding $\mathbf{b}$ shifts the whole distribution. While in Fourier, it is only a pointwise product.

Recall

- Fourier Transform: $\widehat{P}_{\mathcal{D}}(\mathbf{s}) = \mathbb{E}_{\mathbf{e}\sim\mathcal{D}}\left[(-1)^{\langle \mathbf{s},\, \mathbf{e}\rangle}\right]$
- Inverse Fourier Transform: $P_{\mathcal{D}}(\mathbf{x}) = \frac{1}{2^k}\sum_{\mathbf{s}\in\mathbb{F}_2^k} \widehat{P}_{\mathcal{D}}(\mathbf{s})(-1)^{\langle \mathbf{s},\, \mathbf{x}\rangle}$
- $\widehat{P}_{\mathrm{Ber}(\eta)^{\otimes k}}(\mathbf{s}) = (1-2\eta)^{\mathrm{wt}(\mathbf{s})}$

Lemma $\mathbf{e} \sim \mathrm{Ber}(\eta)^{\otimes k}$

$$\widehat{P}_{\mathbf{Ae}+\mathbf{b}}(\mathbf{s}) = \mathbb{E}_{(\mathbf{A},\mathbf{b})}\Big[\underbrace{(1-2\eta)^{\mathrm{wt}(\mathbf{A}^\top\mathbf{s})}}_{\text{coefficient of } \mathbf{e} \text{ at } \mathbf{A}^\top\mathbf{s}} \cdot \underbrace{(-1)^{\langle \mathbf{s},\, \mathbf{b}\rangle}}_{\text{bias of } \mathbf{b}}\Big]$$

Proof: $(-1)^{\langle \mathbf{s},\, \mathbf{Ae}+\mathbf{b}\rangle} = (-1)^{\langle \mathbf{A}^\top\mathbf{s},\, \mathbf{e}\rangle} \cdot (-1)^{\langle \mathbf{s},\, \mathbf{b}\rangle}$, then use the formula for $\mathrm{Ber}(\eta)^{\otimes k}$ above.

Fourier view of the affine reduction

Adding $\mathbf{b}$ shifts the whole distribution. While in Fourier, it is only a pointwise product.

Recall

- Fourier Transform: $\widehat{P}_{\mathcal{D}}(\mathbf{s}) = \mathbb{E}_{\mathbf{e}\sim\mathcal{D}}\left[(-1)^{\langle \mathbf{s},\, \mathbf{e}\rangle}\right]$
- Inverse Fourier Transform: $P_{\mathcal{D}}(\mathbf{x}) = \frac{1}{2^k}\sum_{\mathbf{s}\in\mathbb{F}_2^k} \widehat{P}_{\mathcal{D}}(\mathbf{s})(-1)^{\langle \mathbf{s},\, \mathbf{x}\rangle}$
- $\widehat{P}_{\mathrm{Ber}(\eta)^{\otimes k}}(\mathbf{s}) = (1-2\eta)^{\mathrm{wt}(\mathbf{s})}$

Lemma $\mathbf{e} \sim \mathrm{Ber}(\eta)^{\otimes k}$

$$\widehat{P}_{\mathbf{Ae}+\mathbf{b}}(\mathbf{s}) = \mathbb{E}_{(\mathbf{A},\mathbf{b})}\Big[\underbrace{(1-2\eta)^{\mathrm{wt}(\mathbf{A}^\top\mathbf{s})}}_{\text{coefficient of } \mathbf{e} \text{ at } \mathbf{A}^\top\mathbf{s}} \cdot \underbrace{(-1)^{\langle \mathbf{s},\, \mathbf{b}\rangle}}_{\text{bias of } \mathbf{b}}\Big]$$

Proof: $(-1)^{\langle \mathbf{s},\, \mathbf{Ae}+\mathbf{b}\rangle} = (-1)^{\langle \mathbf{A}^\top\mathbf{s},\, \mathbf{e}\rangle} \cdot (-1)^{\langle \mathbf{s},\, \mathbf{b}\rangle}$, then use the formula for $\mathrm{Ber}(\eta)^{\otimes k}$ above.

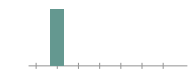
Goal: choose distribution of $(\mathbf{A}, \mathbf{b})$ so that $\widehat{P}_{\mathbf{Ae}+\mathbf{b}}(\mathbf{s})$ matches $\widehat{P}_{\mathcal{D}}(\mathbf{s})$

Hardness of Fourier-bounded noise

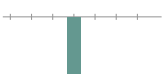
Theorem. If $\sum_{\mathbf{s}\neq\mathbf{0}}|\widehat{P}_{\mathcal{D}}(\mathbf{s})| \leq 1-2\eta$, then $\mathsf{LPN}_{n,k,\mathcal{D}}$ is as hard as $\mathsf{LPN}_{n,\eta}$.

Proof idea

1. Split the random choice of $(\mathbf{A}, \mathbf{b})$ into branches, each identified by a vector $\mathbf{s} \in \mathbb{F}_2^k$.
2. Branch $\mathbf{s} \neq \mathbf{0}$: $\mathbf{Ae} + \mathbf{b}$ has a one-hot Fourier spectrum, only $\pm(1-2\eta)$ at $\mathbf{s}$.



spectrum of branch $\mathbf{s}_1$



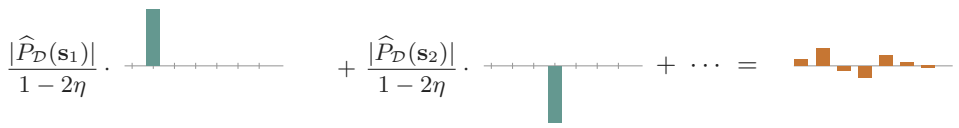
spectrum of branch $\mathbf{s}_2$

Hardness of Fourier-bounded noise

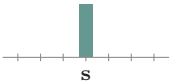
Theorem. If $\sum_{\mathbf{s}\neq\mathbf{0}}|\widehat{P}_{\mathcal{D}}(\mathbf{s})| \leq 1-2\eta$, then $\mathsf{LPN}_{n,k,\mathcal{D}}$ is as hard as $\mathsf{LPN}_{n,\eta}$.

Proof idea

1. Split the random choice of $(\mathbf{A},\mathbf{b})$ into branches, each identified by a vector $\mathbf{s}\in\mathbb{F}_2^k$.
2. Branch $\mathbf{s}\neq\mathbf{0}$: $\mathbf{Ae}+\mathbf{b}$ has a one-hot Fourier spectrum, only $\pm(1-2\eta)$ at $\mathbf{s}$.
3. Pick branch $\mathbf{s}$ w.p. $|\widehat{P}_{\mathcal{D}}(\mathbf{s})|/(1-2\eta)$: by linearity, the mixture's spectrum is $\widehat{P}_{\mathcal{D}}$.

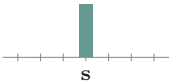


One branch: a one-hot Fourier spectrum



$$\widehat{P}_{\mathbf{Ae}+\mathbf{b}}(\mathbf{s}') = (1-2\eta)^{\mathrm{wt}(\mathbf{A}^\top \mathbf{s}')} \qquad \cdot \qquad \mathbb{E}_{\mathbf{b}}\left[(-1)^{\langle \mathbf{s}', \mathbf{b}\rangle}\right]$$

One branch: a one-hot Fourier spectrum

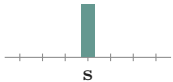


$$\widehat{P}_{\mathbf{Ae}+\mathbf{b}}(\mathbf{s}') = (1-2\eta)^{\mathrm{wt}(\mathbf{A}^\top \mathbf{s}')} \cdot \underbrace{\mathbb{E}_{\mathbf{b}}\left[(-1)^{\langle \mathbf{s}', \mathbf{b}\rangle}\right]}_{\mathbf{b}\text{: mask every } \mathbf{s}' \neq \mathbf{s}}$$

Branch $\mathbf{s} \neq \mathbf{0}$

- Sample $\mathbf{b}$ uniformly subject to $(-1)^{\langle \mathbf{s}, \mathbf{b}\rangle} = \operatorname{sgn} \widehat{P}_{\mathcal{D}}(\mathbf{s})$: factor $\operatorname{sgn} \widehat{P}_{\mathcal{D}}(\mathbf{s})$ at $\mathbf{s}$, 0 elsewhere.

One branch: a one-hot Fourier spectrum

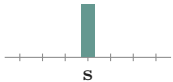


$$\widehat{P}_{\mathbf{Ae}+\mathbf{b}}(\mathbf{s}') = \underbrace{(1-2\eta)^{\mathrm{wt}(\mathbf{A}^\top \mathbf{s}')}}_{\mathbf{A}\text{: maximize it at } \mathbf{s}'=\mathbf{s}} \cdot \underbrace{\mathbb{E}_{\mathbf{b}}\left[(-1)^{\langle \mathbf{s}', \mathbf{b}\rangle}\right]}_{\mathbf{b}\text{: mask every } \mathbf{s}'\neq\mathbf{s}}$$

Branch $\mathbf{s} \neq \mathbf{0}$

- Sample $\mathbf{b}$ uniformly subject to $(-1)^{\langle \mathbf{s}, \mathbf{b}\rangle} = \operatorname{sgn} \widehat{P}_{\mathcal{D}}(\mathbf{s})$:
  factor $\operatorname{sgn} \widehat{P}_{\mathcal{D}}(\mathbf{s})$ at $\mathbf{s}$, 0 elsewhere.
- Fixed invertible $\mathbf{A}$ with $\mathbf{A}^\top \mathbf{s} = (1, 0, \ldots, 0)$:
  Hamming weight 1, achieving largest factor $1-2\eta$.

One branch: a one-hot Fourier spectrum



$$\widehat{P}_{\mathbf{Ae}+\mathbf{b}}(\mathbf{s}') = \underbrace{(1-2\eta)^{\mathrm{wt}(\mathbf{A}^\top \mathbf{s}')}}_{\mathbf{A}\text{: maximize it at } \mathbf{s}'=\mathbf{s}} \cdot \underbrace{\mathbb{E}_{\mathbf{b}}\left[(-1)^{\langle \mathbf{s}', \mathbf{b}\rangle}\right]}_{\mathbf{b}\text{: mask every } \mathbf{s}'\neq\mathbf{s}}$$

Branch $\mathbf{s} \neq \mathbf{0}$

- Sample $\mathbf{b}$ uniformly subject to $(-1)^{\langle \mathbf{s}, \mathbf{b}\rangle} = \operatorname{sgn} \widehat{P}_{\mathcal{D}}(\mathbf{s})$:
  factor $\operatorname{sgn} \widehat{P}_{\mathcal{D}}(\mathbf{s})$ at $\mathbf{s}$, 0 elsewhere.
- Fixed invertible $\mathbf{A}$ with $\mathbf{A}^\top \mathbf{s} = (1, 0, \ldots, 0)$:
  Hamming weight 1, achieving largest factor $1-2\eta$.
- Together: $\operatorname{sgn} \widehat{P}_{\mathcal{D}}(\mathbf{s}) \cdot (1-2\eta)$ at $\mathbf{s}$, and 0 at every other $\mathbf{s}' \neq \mathbf{0}$.

Branch $\mathbf{s} = \mathbf{0}$ (the rest): $\mathbf{A} = I$, $\mathbf{b}$ uniform $\Rightarrow$ every nonzero coefficient is 0.

Matching algorithm

Large Fourier coefficient (opposite of Fourier-bounded)

- Take the largest Fourier coefficient $|\widehat{P}_{\mathcal{D}}(\mathbf{s})| = 1 - \gamma$ for some $\mathbf{s} \neq \mathbf{0}$
- XOR the ith samples in each batch with $s_i = 1$,
  resulting samples have noise rate $\gamma/2 \Rightarrow \mathsf{LPN}_{n,\gamma/2}$: time $e^{O(\gamma n)}$.

Summary

Random affine maps + Fourier analysis

- Hard: Fourier mass $\leq 1 - 2\eta$; dense; δ-SV with $\delta \lesssim 2^{-k/2}(1 - 2\eta)$.
- Easy: one large Fourier coefficient; one rare noise pattern.

Open problems

- Which $\mathcal{D}$ can affine reductions reach? (entropy and Fourier barriers)
- Noise correlated across batches.

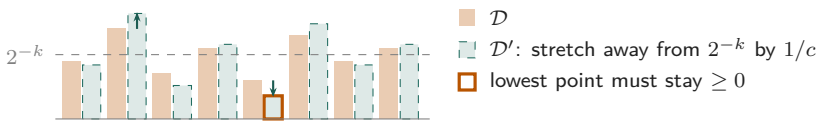
Thank you!

Backup: hardness of dense noise

Theorem. If $P_{\mathcal{D}}(\mathbf{z}) \gtrsim k\eta\, 2^{-k}$ for all $\mathbf{z}$, then $\mathsf{LPN}_{n,k,\mathcal{D}}$ is as hard as $\mathsf{LPN}_{n,\eta}$.

$$\widehat{P}_{\mathbf{Ae}+\mathbf{b}}(\mathbf{s}) = \underbrace{\mathbb{E}_{\mathbf{A}}\left[(1-2\eta)^{\mathrm{wt}(\mathbf{A}^\top \mathbf{s})}\right]}_{\mathbf{A}\text{: shrink all coeff equally}} \cdot \underbrace{\widehat{P}_{\mathcal{D}'}(\mathbf{s})}_{\mathbf{b}\text{: stretch back}}$$

- Uniform $\mathbf{A} \in \mathrm{GL}(k,2)$: $\mathbf{A}^\top \mathbf{s}$ is uniform, so every $\mathbf{s} \neq \mathbf{0}$ shrink by $c \approx 1 - k\eta$.
- $\mathbf{b} \sim \mathcal{D}' = \mathrm{Unif} + \frac{1}{c}(\mathcal{D} - \mathrm{Unif})$, an intensified version of $\mathcal{D}$: then $c \cdot \widehat{P}_{\mathcal{D}'} = \widehat{P}_{\mathcal{D}}$.
- $\mathcal{D}'$ must be a distribution: its lowest point must stay $\geq 0 \iff \mathcal{D}$ is dense.



Backup: algorithm for a rare noise pattern

Rare point $P_{\mathcal{D}}(\mathbf{z}_0) \leq \tau$ (opposite of dense)

- As for low-noise LPN: draw $M \approx n^k 4^k$ batches and hope that no noise equals $\mathbf{z}_0$. This happens with probability $e^{-O(\tau M)}$.
- Then all noise lies in the zero set of a degree-k polynomial $\Rightarrow$ AG linearization.
- Repeat $e^{O(\tau M)}$ times; test the candidates on fresh samples.

Backup: a better bound for SV sources

GMR: $\delta \lesssim 2^{-k}(1-2\eta)$ $\longrightarrow$ ours: $\delta \lesssim 2^{-k/2}(1-2\eta)$.

GMR's simulation build the noise one bit at a time (write $\varepsilon = \frac{1}{2} - \eta$)

$$e_i = e' + F_0 + \sum_{j<i} F_j\, e_j, \qquad e' \sim \mathrm{Ber}(\eta) \text{ fresh}, \quad (F_0, \ldots, F_{i-1}) \sim \mu$$

Need: for every prefix $\mathbf{z} = (e_1, \ldots, e_r)$, $r = i-1$, the bit e_i gets the right conditional law.

Our lemma Let $g(\mathbf{z}) = (\frac{1}{2} - \Pr[e_i = 1 \mid \mathbf{z}])/\varepsilon$. Then μ exists if $\sum_{\mathbf{z}} g(\mathbf{z})^2 \le 1$:

$$\text{weight of } \mu = 2^{-r}\sum_{\mathbf{s}} |\widehat{g}(\mathbf{s})| \overset{\text{C–S}}{\le} \Big(2^{-r}\sum_{\mathbf{s}} \widehat{g}(\mathbf{s})^2\Big)^{1/2} \overset{\text{Parseval}}{=} \Big(\sum_{\mathbf{z}} g(\mathbf{z})^2\Big)^{1/2} \le 1$$

SV $|g(\mathbf{z})| \le \delta/\varepsilon$ on 2^{i-1} prefixes, so the sum is $\le 2^{k-1}\delta^2/\varepsilon^2 \le 1$ when $\delta \le 2^{-(k-1)/2}\varepsilon$.

Backup: examples of Fourier-bounded noise

Each has $\sum_{\mathbf{s}\neq\mathbf{0}}|\widehat{P}_{\mathcal{D}}(\mathbf{s})| \le 1-2\eta$, so it is as hard as $\mathsf{LPN}_{n,\eta}$.

Noise $\mathcal{D}$	**In words**	**Hard when**
δ-biased	all nonzero coefficients in $[-\delta,\delta]$	$\delta \le \frac{1-2\eta}{2^k-1}$
$\bigotimes_i \mathrm{Ber}(\frac{1}{2}-\gamma_i)$	independent bits, different rates	$\prod_i(1+2\lvert\gamma_i\rvert)-1 \le 1-2\eta$
$(1-\rho)\,\mathrm{Unif}+\rho\,Q$	uniform, mixed with any Q	$\rho \le \frac{1-2\eta}{2^k-1}$
$e_i = e_{i-1}+\mathrm{Ber}(\frac{1}{2}-\gamma)$	Markov chain, e_1 uniform	$(1+2\lvert\gamma\rvert)^{k-1}-1 \le 1-2\eta$

A δ-SV source is 2δ-biased, but not the other way around.

Backup: limits of affine reductions

Entropy $H(\mathbf{Ae}+\mathbf{b}) \geq H(\mathbf{e})$: an invertible affine map never lowers entropy.
So noise with less entropy than $\mathrm{Ber}(\eta)^{\otimes k}$ is out of reach.

Fourier Every nonzero coefficient of $\mathbf{Ae}+\mathbf{b}$ has absolute value $\leq 1-2\eta$.
So uniform noise on an affine hyperplane (one coefficient is ± 1) is out of reach.

Not tight $\mathrm{Ber}(\eta)^{\otimes k}$ is reachable, yet its Fourier mass is $(2-2\eta)^k - 1 > 1-2\eta$.
Our L1 condition is sufficient, not necessary.

Across batches We assume fresh noise for each batch.
E.g. one uniform $\mathbf{c}$ reused in all batches looks uniform, yet is easy to solve.

Backup: precise statements

Hardness: as hard as $\mathrm{LPN}_{n,\eta}$ (q batches from qk samples, plus $q \cdot \mathrm{poly}(2^k, n)$ time)

Fourier-bounded	$\sum_{\mathbf{s}\neq\mathbf{0}} \lvert\widehat{P}_{\mathcal{D}}(\mathbf{s})\rvert \le 1-2\eta$
Dense	$P_{\mathcal{D}}(\mathbf{z}) \ge \frac{1-(1-\eta)^k}{2^k-1}$ for every $\mathbf{z}$
δ-SV	$\delta \le 2^{-(k-1)/2}\,\varepsilon$ with $\varepsilon = \frac{1}{2}-\eta$ (GMR: $2^{-k-3}\varepsilon$)

Algorithms

Large coefficient	$\max_{\mathbf{s}\neq\mathbf{0}} \lvert\widehat{P}_{\mathcal{D}}(\mathbf{s})\rvert = 1-\gamma$: time $O(k2^k) + e^{O(\gamma n)} \cdot \mathrm{poly}(n, k, \frac{1}{1-\gamma})$
Rare point	$P_{\mathcal{D}}(\mathbf{z}_0) \le \tau$, $\Delta > 0$: time $\exp\big(O(\tau N_k 4^k)\big) \cdot \mathrm{poly}(N_k, \frac{1}{\Delta})$

Here $N_k = \sum_{i=1}^{k} \binom{n}{i}$ and $\Delta = d_{\mathrm{TV}}(\mathcal{D}, \mathrm{Unif})$.