

Nondeterministic Interactive Refutations for Nearest Boolean Vector

Andrej Bogdanov* Alon Rosen†

Abstract

Most n -dimensional subspaces \mathcal{A} of \mathbb{R}^m are $\Omega(\sqrt{m})$ -far from the Boolean cube $\{-1, 1\}^m$ when $n < cm$ for some constant $c > 0$. How hard is it to certify that the Nearest Boolean Vector (NBV) is at least $\gamma\sqrt{m}$ far from a given random \mathcal{A} ?

Certifying NBV instances is relevant to the computational complexity of approximating the Sherrington-Kirkpatrick Hamiltonian, i.e. maximizing $x^T Ax$ over the Boolean cube for a matrix A sampled from the Gaussian Orthogonal Ensemble. The connection was discovered by Mohanty, Raghavendra, and Xu (STOC 2020). Improving on their work, Ghosh, Jeronimo, Jones, Potechin, and Rajendran (FOCS 2020) showed that certification is not possible in the sum-of-squares framework when $m \ll n^{1.5}$, even with distance $\gamma = 0$.

We present a non-deterministic interactive certification algorithm for NBV when $m \gg n \log n$ and $\gamma \ll 1/mn^{1.5}$. The algorithm is obtained by adapting a public-key encryption scheme of Ajtai and Dwork.

1 Introduction

When can we expect to have a reduction from problem A to problem B? Complexity theory can be used not only to show existence of reductions but also to argue separations. For example, one reason an oracle for factoring is not considered an imminent threat to SAT is that the correctness of prime factorizations can be both proved and refuted, that is (the decision version of) factoring is in $\text{NP} \cap \text{coNP}$.

In general, there cannot be a reduction (of sufficiently low complexity) from A to B if there is a complexity class that (conjecturally) separates the two. For worst-case problems in NP the separating class is often $\text{NP} \cap \text{coNP}$ or one of its close relatives ($\text{NP} \cap \text{coAM}$ or Statistical Zero-Knowledge).

It is natural to wonder whether analogous separations in average-case complexity can clarify the landscape of reductions within distributional NP; a class of particular importance to cryptography and learning theory. Reductions among non-NP-complete distributional problems do exist, but are few and far between. Notable examples include lattice problems [MR07, Reg09, Pei09, LM09]. More recently, a web of reductions was developed to explain the hardness of various statistical inference problems [BB20].

A handful of average-case NP complete problems were found in the 1980-90s [Lev86, Gol11]. All these problems are closely related to simulation of Turing Machines, perhaps necessarily

*University of Ottawa. Email: abogdano@uottawa.ca

†Bocconi University and Reichman University. Email: alon.rosen@unibocconi.it

so [Tre10]. The conjectured hardness of combinatorial problems like random SAT or planted clique still lacks satisfactory explanation.

In the context of random SAT, Feige, Kim, and Ofek [FKO06] showed that random 3CNF instances with n variables and $m \gg n^{1.4}$ equations admit efficient nondeterministic refutations of satisfiability, that is, belong to Avg-coNP.¹ Although most such instances are unsatisfiable, it is not known how to efficiently certify the lack of a satisfying assignment in the regime $n^{1.4} \ll m \ll n^{1.5}$. On the other hand, when $m \ll n^{1.4}$ not even nondeterministic refutations are known. Thus we do not expect a reduction from random 3SAT with clause-to-variable density $n^{0.41}$ to random 3SAT with density $n^{0.39}$ barring a major algorithmic advance.

Our contribution is an analogous result for the distributional Nearest Boolean Vector to a Subspace problem which was introduced by Mohanty, Raghavendra and Xu [MRX20]. In Theorem 1 we show that for a certain parameter regime in which this problem may be intractable, the problem is in average-case statistical zero-knowledge (Avg-SZK) and therefore admits *interactive* nondeterministic refutations.

1.1 The Nearest Boolean Vector problem

We work with the following formulation of the Nearest Boolean Vector problem:

Nearest Boolean Vector (NBV):

Input: An n -dimensional subspace \mathcal{A} of \mathbb{R}^m .

Yes instances: There exists a $v \in \{-1, 1\}^m$ such that $\text{dist}(v, \mathcal{A}) \leq \gamma\sqrt{m}$.

No instances: For all $v \in \{-1, 1\}^m$, $\text{dist}(v, \mathcal{A}) > \sqrt{m}/2$.

When $n < cm$ for a sufficiently small constant c , most subspaces \mathcal{A} (chosen from the uniform Haar measure) are no instances [MRX20]. We are interested in the errorless average-case complexity of NBV. An efficient average-case algorithm for distributional NBV can be viewed as an efficiently computable certificate that most subspaces are far from the Boolean cube.

When $\gamma < 1/2$, NBV is in NP. Several works [MRX20, GJJ⁺20, PTVW22] provide evidence that it is intractable on average in the regime $m \ll n^2$.

1.2 Our Result

We give a reduction from distributional NBV to the Statistical Distance to Uniform (SDU) problem. The input to SDU is a sampler of outputs in $\{0, 1\}^n$, the YES instances are samplers whose outputs are $1 - \delta$ far from uniform, and NO instances are samplers whose values are δ close to uniform. For $\delta = 1/3$ SDU is in the class Statistical Zero Knowledge (SZK) [SV03], which is a subclass of coAM.²

Theorem 1. *Let C be a sufficiently large constant and $\epsilon \geq 2^{-n/C}$. For all but an ϵ -fraction of instances, NBV with parameters $m = Cn \log n$ and $\gamma = 1/Cmn^{3/2} \log^{1/2}(n/\epsilon)$ is in AvgSZK.*

The proof is given in Section 3. In Section 5 we outline a tentative approach for improving the completeness error γ .

¹Their result was recently extended to the semi-random model [GKM22] in which the formula is arbitrary and only the literals are polarized randomly.

²When $\delta = 1/n$, SDU is in the more restricted class Non-Interactive Statistical Zero-Knowledge (NISZK) [GSV99]. AvgNISZK membership can also be obtained for smaller γ .

When ϵ is polynomial in n , SZK membership holds for all but a $n^{-O(1)}$ fraction of instances and the approximation factor γ has value $\tilde{\Theta}(1/mn^{3/2})$. When ϵ is $2^{-\Omega(n)}$ then the fraction of instances is exponential, but $\gamma = \Theta(1/mn^2)$.

2 Background and Overview

2.1 Average-case refutations

Refutations come up naturally in the study of combinatorial optimization. A worst-case approximation algorithm A for a minimization problem P is required to output a value within a factor of c of the optimum on all instances. Such an algorithm provides an efficient refutation of the claim

$$x \text{ has a solution of value at most } A(x)/c \tag{1}$$

for every instance x .

When efficient refutations are hard to obtain for all x it may be natural to relax the condition to hold for most x . An average-case refutation should still certify (1), but it is now allowed to fail on some small fraction of inputs x .

For many natural distributions, the optimum is tightly concentrated around its expectation. For example, the maximum number of satisfiable clauses in random 3SAT with sufficiently large clause-to-variable density is close to $7/8$ on most instances. In particular, an average-case refutation must certify that most instances are not satisfiable, but it should be allowed to output “I don’t know” on a small fraction of inputs. This motivates the following definition:

Definition 2. *A refutation R with failure rate ϵ for distributional (promise) problem f is an algorithm that outputs “no” or “I don’t know”, is always correct ($R(x) = f(x)$ or “I don’t know”), and outputs “I don’t know” on at most an ϵ -fraction of inputs.*

While efficient deterministic or randomized refutations are needed for the design of approximation algorithms, in this work we are interested in the existence of nondeterministic (coNP-type) refutations. Such refutations yield efficiently *verifiable* certificates of (1) on most inputs. As a consequence of Theorem 1 we have

Corollary 3. *There is a efficient nondeterministic interactive refutation for NBV with failure rate $\epsilon \geq 2^{-n/C}$ and parameters $m = Cn \log n$, $\gamma = 1/Cmn^{3/2} \log^{1/2}(n/\epsilon)$.*

2.2 Refutations in the Sum-of-Squares Framework

The sum-of-squares (SoS) framework is an incomplete but powerful framework for refuting optimization problems. It has been used to argue efficient refutations do not exist for problems such as clique [BHK⁺19]. The most notable incorrect prediction of SoS is on random 3LIN with perfect completeness [Gri01, Sch08]. In that case not only do refutations exist but they can be found by Gaussian elimination.

In contrast, the nondeterministic refutations of Feige, Kim, and Ofek arise as solutions to the level- $O(n^{2\delta})$ SoS relaxation of random 3SAT with n variables and m constraints. This may be viewed as evidence that SoS correctly predicts refutability in problems that are immune to Gaussian elimination “attacks”.

2.3 Sherrington-Kirkpatrick and Nearest Boolean Vector

The negative energy of the Sherrington-Kirkpatrick Hamiltonian at zero-temperature is the value

$$SK(M) = \min \frac{1}{\sqrt{n}} \cdot x^T M x \quad \text{subject to } x \in \{\pm 1/\sqrt{n}\}^n$$

for a matrix M sampled from the Gaussian Orthogonal Ensemble. It can be efficiently certified that $SK(M) \leq 2 + \epsilon$ for every $\epsilon > 0$ and most matrices M via the relaxation

$$SK(M) \leq \min_{\|u\|=1} \frac{1}{\sqrt{n}} \cdot u^T M u = \lambda_1(M), \quad (2)$$

where $\lambda_1(M)$ is the largest eigenvalue of M , which is known to not exceed $2 + \epsilon$ for most matrices M .

Parisi [Par79] conjectured and Talagrand [Tal11] proved that $SK(M)$ is in fact strictly smaller than 2 for most matrices M . The true value for most M is concentrated around Parisi's constant $P_* \approx 1.526$. More recently Montanari [Mon19] found an algorithm that finds a solution x for which $x^T M x \leq P_* - \epsilon$ for most matrices M and proved its correctness under some plausible conjecture.

Mohanty, Raghavendra, and Xu [MRX20] ask whether Montanari's algorithm can be matched with an efficient certificate that $SK(M) \leq P_* + \epsilon$ for most matrices M . Together with Montanari's algorithm, this would give an errorless heuristic for calculating $SK(M)$ up to lower-order terms. As a first step they show that SK reduces to the potentially more tractable Nearest Boolean Vector Problem.

Mohanty, Raghavendra, and Xu prove that for all $c, \gamma > 0$ there exists an $\epsilon > 0$ such that if NBV with parameters $m/n = c$ and γ admits efficient refutations then so does the claim $SK(M) \leq 2 - \epsilon$ for most M . Moreover, for sufficiently small c , most subspaces \mathcal{A} are no-instances of NBV.

However, their main evidence for refutability of NBV is negative: They show that no refutations can be obtained from the natural degree-4 SoS relaxation of NBV for any constant c , even in case of perfect completeness $\gamma = 0$. A refutation algorithm for $\gamma = 0$ is merely required to certify that no Boolean vector belongs to the subspace \mathcal{A} . The SoS hardness regime was later extended to $m \ll n^{3/2}$ and to degree- $n^{\Omega(1)}$ SoS by Ghosh et al. [GJJ+20]. It is believed that it can be further extended up to $m < n^2/4$, as (heuristically) suggested by calculations of the low-degree likelihood ratio (see Potechin et al. [PTVW22]).

Theorem 1 has no bearing on the complexity of certifying that $SK(M) \leq 2$ for most M . To obtain an improvement over the spectral certificate (2) the completeness error γ would have to be constant, or at least $m^{-\epsilon}$ for some small ϵ .

2.4 Algorithms for NBV

When $m \gg n^2$ and γ is a sufficiently small constant it is plausible that NBV can be efficiently solved by linearization. Represent \mathcal{A} as the column span of B for some $m \times n$ matrix B . Consider the objective

$$\text{minimize } \sum_{i=1}^m (\langle B_i, x \rangle^2 - 1)^2 \quad \text{over } x \in \mathbb{R}^n, \quad (3)$$

where B_i is the i -th row of B . If \mathcal{A} had a Boolean vector $\langle B_i, x \rangle = \pm 1$ the value of this objective would be zero. We suspect that for most matrices B it should be lower bounded by $\Omega(m)$. If (3)

were efficiently computable its value would be the required certificate. Although this is unlikely, the same argument can be applied to its linearization in which degree-2 monomials $x_i x_j$ are represented by variables y_{ij} :

$$\text{minimize } \sum_{i=1}^m \left(\sum_{j,k=1}^n B_{ij} B_{ik} y_{jk} - 1 \right)^2 \quad \text{over } y \in \mathbb{R}^{n(n+1)/2}, \quad (4)$$

which is a convex quadratic objective and therefore efficiently minimizable.

In the case of perfect completeness, $\gamma = 0$ NBV reduces to the Shortest Vector Problem (SVP) in lattices with approximation factor exponential in the dimension and can therefore be solved by the LLL algorithm [LLL82] for any $m > n$. Here is an outline of the (standard) reduction R . Let the columns of $C \in \mathbb{R}^{m \times (m-n)}$ be a random orthonormal basis of the dual subspace \mathcal{A}^\perp . Consider the lattice \mathcal{L} spanned by the rows of the $m \times (2m-n)$ matrix $C' = [\delta I_m | C]$ for $\delta = 2^{-2m^2}$. If \mathcal{A} contained a Boolean vector x then $C'x$ would be a vector of length $\delta\sqrt{m}$ in \mathcal{L} . If not, by a union bound there is unlikely to exist a vector $x \in \{-2^m, \dots, 2^m\}^m$ for which $\|Cx\| < 2^m \delta$ so the shortest vector in \mathcal{L} has length at least $2^m \delta$.

2.5 Nondeterministic refutations for NBV

This reduction R extends to almost-perfect completeness $\gamma = 2^{-\Theta(m^2)}$. It is tempting to conjecture for $m \gg n \log n$ that there is a constant d such that R reduces NBV with parameter $\gamma = m^{-d}$ to SVP with approximation factor \sqrt{m} , which is a coNP problem [AR05]. Should such a reduction exist it would imply efficient nondeterministic refutations for NBV.

We were unable to prove the soundness of R in this parameter regime. Our preliminary calculations indicate that \mathcal{L} may contain unusually short vectors for most instances \mathcal{A} of NBV.

Instead, we prove Theorem 1 by adapting a public-key encryption scheme of Ajtai and Dwork [AD97] (see [BCHR22] for a “modern” description) into the desired reduction from NBV to SDU.

2.6 Refutations, SZK, and Public-key Encryption

The *chosen plaintext attack* security notion for one-bit encryption with public key PK and encryption algorithm Enc posits that the distributions $(PK, Enc(PK, 0))$ and $(PK, Enc(PK, 1))$ are computationally indistinguishable. In contrast, functionality requires that they be statistically distinguishable by the decryption algorithm.

The security of several public-key encryption candidates is argued using a model (fake) public-key distribution FK with the property that PK and FK are computationally indistinguishable while $(FK, Enc(FK, 0))$ and $(FK, Enc(FK, 1))$ are statistically indistinguishable. This proof strategy yields a reduction from distinguishing real and model public keys to SDU.

The security proof for the Ajtai-Dwork (AD) and Bogdanov et al.’s (BCHR) pancake encryptions are of this type. In BCHR, the model public key FK is a sequence m of independent standard n -dimensional Gaussians, while in the real public key PK an almost-periodic component is planted in a secret direction s of \mathbb{R}^n . If the almost-periodic component is concentrated around the values -1 and 1 , the row-span of PK can be viewed as a yes-instance of NBV.

To turn this distinguisher between PK and FK into a refutation, we observe that the encryption remains functional even for a worst-case choice of PK that satisfies some efficiently verifiable conditions (the largest and smallest singular values of PK are pseudorandom). By verifying these conditions the reduction from NBV to SDU ensures that *all* yes-instances of

NBV map to yes-instances of SDU, while affecting only a small fraction of no-instances, thus providing interactive remoteness certificates for most instances of NBV.

The BCHR encryption and security proof suggest the following visualization of the remoteness certificates. If a random matrix FK is multiplied on the right by a random $x \sim \{\pm 1\}^m$ the output $FK \cdot x$ is close to a random Gaussian point in \mathbb{R}^n (see Fact 17). On the other hand, $PK \cdot x$ is concentrated around “pancakes” perpendicular to the secret direction s . To certify remoteness, the verifier asks the prover to furnish an $x \in \{\pm 1\}^m$ close to a random Gaussian point g in \mathbb{R}^n . Unless g happens to land close to a pancake the prover will fail on an no-instance PK of NBV .

A fatal weakness of BCHR encryption is that it is insecure unless $m \gg n^2$, a setting of parameters in which NBV is tractable. In contrast, security of AD can be proved when $m = O(n \log n)$. This improvement is obtained by modifying the encryption from $\text{round}(PK \cdot x)$ to $\text{round}(\sigma A \cdot x) \bmod \mathcal{P}(B)$, where $PK = [A|B]$ with $A \in \mathbb{R}^{(m-n) \times n}$ and $B \in \mathbb{R}^{n \times n}$ is the public key matrix, $\mathcal{P}(B)$ is the parallelepiped spanned by the columns of B , and σ is a suitable scaling factor. Reduction 1 in Section 3 implements this security proof (in a different basis which is more suitable for analysis), again by imposing some efficiently verifiable conditions that hold for typical yes-instances but for none of the no-instances of NBV .

3 Refutation via Lattice Smoothing

We represent the random subspace \mathcal{A} as the row space of a random $n \times m$ matrix $[A|B']$ of independent normal entries. It is sufficient to specify these entries up to $O(\log n)$ bits of precision. We carry out our analyses assuming infinite precision. It will be clear from the calculations that the additional effect of rounding the entries of A does not affect correctness.

For a real number x let $x = \lfloor x \rfloor + \{x\}$ be its unique representation with $\lfloor x \rfloor \in \mathbb{Z}$ and $\{x\} \in [-1/2, 1/2)$. Let $\{x\}_p$ be the multiple of $1/p$ in $[-1/2, 1/2)$ closest to $\{x\}$. The notation extends to vectors and matrices entrywise.

Fact 4. (a) $|\{x\}| \leq |x|$ and (b) $|\{x+y\}| \leq |\{x\}| + |\{y\}|$.

We choose the modulus p to equal $Cn\sqrt{m}$ for a sufficiently large constant C . Let $\sigma = (1/\pi)\sqrt{n \ln(12mn/\epsilon + 2n)}$.

Reduction 1: On input $[A|B']$, $A \in \mathbb{R}^{n \times (m/2)}$, $B' \in \mathbb{R}^{n \times (m/2)}$,

- 1 Find a submatrix B of B' with smallest singular value at least $1/\sqrt{n}$.
- 2 If step 1 is unsuccessful, fail.
- 3 If any column of A has norm more than $2\sqrt{n}$, fail.
- 4 Otherwise, output the sampler S that maps $x \sim \{\pm 1\}^{m/2}$ to $\{\{\sigma B^{-1}A\}_p x\} \in \frac{1}{p}\mathbb{Z}_p^n$.

A naive implementation of step 1 would split B' into m/n candidate matrices B and attempt to find one with singular value $1/\sqrt{n}$, resulting in failure rate $\epsilon = 2^{-O(m/n)}$ which is $n^{-O(1)}$ when $m = O(n \log n)$. In Section 4 we design a greedy procedure for choosing B that improves the failure rate to $2^{-\Omega(n)}$.

Theorem 1 follows from Claims 5 and 8.

Claim 5. Assume $\epsilon > 2^{-\Omega(n)}$. For all but an ϵ -fraction of instances $[A|B']$ the output of S is $1/3$ -close to a uniformly random element of $\frac{1}{p}\mathbb{Z}_p^n$.

Fact 6 (Smoothing). [MR07, Lemmas 3.3 and 4.1] If all columns of $B \in \mathbb{R}^{n \times n}$ have norm at most b , g is standard normal in \mathbb{R}^n , and $\sigma \geq (b/2\pi)\sqrt{\ln(n/\epsilon + 2n)}$, then $\{\sigma B^{-1}g\}$ is ϵ -close to a uniform random point in $[-1/2, 1/2)^n$.

Fact 7 (Leftover hash lemma). [ILL89] If $C \sim \mathbb{Z}_p^{n \times m}$ is a random matrix and $x \in \mathbb{Z}_p^m$ be a random vector uniformly distributed on some set of size M then (C, Cx) is $\sqrt{p^n/M}$ -close to uniformly random.

Proof of Claim 5. By Proposition 9 B can be found (efficiently) except with probability $\exp(-\Omega(m))$. By large deviation bounds all columns of B have norm at most $2\sqrt{n}$ except with probability $2^{-\Omega(n)}$. By our choice of parameters, both conditions are satisfied except with probability $2^{-\Omega(m)} + 2^{-\Omega(n)} \leq \epsilon/2$. Assuming this we argue the conclusion holds even when conditioning on B .

For each column a_i of A , $\sigma a_i \in \mathbb{R}^n$ is a normal vector of zero mean and covariance σI . By smoothing Fact 6, $\{\sigma B^{-1}a_i\}$ is $\epsilon/4m$ -close to a uniform point in $[-1/2, 1/2]^n$. Therefore $C = \{\sigma B^{-1}A\}_p$ is $\epsilon/12$ -close to a random matrix in $\frac{1}{p}\mathbb{Z}_p^{(m/2) \times n}$. By Fact 7, (C, Cx) is $\epsilon/12 + \sqrt{p^n/2^{m/2}}$ -close to random. By our choice of parameters, $\epsilon/12 + \sqrt{p^n/2^{m/2}} \leq \epsilon/6$. By Markov's inequality the output of the sampler is $1/3$ -close to random except with probability $\epsilon/2$ over the choice of A , and therefore except with probability ϵ over the choice of A and B' . \square

Claim 8. If $[A|B']$ is a yes instance of NBV with parameters $m > Cn \log n$ and $\gamma < 1/Cmn^{3/2} \log^{1/2}(n/\epsilon)$, either the reduction fails, or the output of S is $2/3$ -far from random.

Proof. As $[A|B']$ is a yes instance of NBV there exists a witness $w \in \mathbb{R}^n$ such that $w[A|B'] = v + e$, where $v \in \{\pm 1\}^m$ and $\|e\| \leq \gamma\sqrt{m}$. Let D be the distinguisher that on input $y \in \frac{1}{p}\mathbb{Z}_p^n$ accepts if $|\langle wB, y \rangle| < 1/24$.

Assume y is uniform in $\frac{1}{p}\mathbb{Z}_p^n$. We show D accepts y with probability at most $1/6$. We can write y as $\{u\}_p$ where u is uniform in $[0, 1)^n$. Let $e' = y - u$ and let v_B and e_B be the projections of v and e on the columns indexed by B . Then

$$\langle wB, y \rangle = \langle v_B + e_B, u + e' \rangle = \langle v_B, u \rangle + \langle v_B, e' \rangle + \langle e_B, y \rangle$$

The random variable $\{\langle v_B, u \rangle\}$ is uniform in $[-1/2, 1/2)$, so $|\langle v_B, u \rangle| > 1/12$ with probability $5/6$. If this happens, by the triangle inequality,

$$\begin{aligned} |\langle wB, y \rangle| &\geq |\langle v_B, u \rangle| - |\langle v_B, e' \rangle| - |\langle e_B, y \rangle| \\ &\geq 1/12 - \|v_B\| \|e'\| - \|e_B\| \|y\| \\ &\geq 1/12 - n/p - \gamma\sqrt{mn} \\ &> 1/24 \end{aligned}$$

and D rejects y .

Now assume the reduction does not fail so that $\|B^{-1}\| \leq \sqrt{n}$ and all columns of A and B have norm at most $2\sqrt{n}$. We will show that D accepts $y = \{\{\sigma B^{-1}A\}_p x\}$ with probability at least $5/6$. Therefore D distinguishes this distribution from the uniform one, so the two must be $2/3$ -far.

Let $E = \{\sigma B^{-1}A\}_p - \{\sigma B^{-1}A\}$. Then

$$\{\sigma B^{-1}A\}_p = \{\sigma B^{-1}A\} + E = \sigma B^{-1}A - \lfloor \sigma B^{-1}A \rfloor + E.$$

Since x is integral,

$$y = \{\{\sigma B^{-1}A\}_p x\} = \{\sigma B^{-1}Ax + Ex\}.$$

Therefore

$$\langle wB, y \rangle = \langle wB, \sigma B^{-1}Ax \rangle + \langle wB, Ex \rangle - \langle wB, f \rangle,$$

where $f = \lfloor \sigma B^{-1}Ax + Ex \rfloor$. The first term equals

$$\langle wB, \sigma B^{-1}Ax \rangle = \sigma \langle wA, x \rangle = \sigma \langle v_A, x \rangle + \sigma \langle e_A, x \rangle,$$

where v_A and e_A are the projections of v and e on the coordinates indexed by the columns of A . The third term equals

$$\langle wB, f \rangle = \langle v_B, f \rangle + \langle e_B, f \rangle.$$

As $\sigma \langle v_A, x \rangle$ and $\langle v_B, f \rangle$ are integers,

$$\begin{aligned} |\{\langle wB, y \rangle\}| &\leq |\sigma \langle e_A, x \rangle| + |\langle wB, Ex \rangle| + |\langle e_B, f \rangle| \\ &\leq \sigma \|e_A\| \|x\| + \|wB\| \|Ex\| + \|e_B\| \|f\| \\ &\leq \sigma \|e_A\| \|x\| + (\|v_B\| + \|e_B\|) \|Ex\| + \|e_B\| (\sigma \|B^{-1}\| \|Ax\| + \|Ex\| + \sqrt{n}) \\ &\leq \sigma \gamma m + (\sqrt{n} + \gamma \sqrt{m})(\sqrt{mn}/p) + \gamma \sqrt{m} (\sigma \cdot \sqrt{n} \cdot \|Ax\| + \sqrt{mn}/p + \sqrt{n}). \end{aligned}$$

As Ax is a random ± 1 sum of vectors of norm at most $2\sqrt{n}$, its expected squared norm is mn , so its norm is at most $3\sqrt{mn}$ with probability at least $5/6$. Since $\gamma < 1/Cmn^{3/2} \log^{1/2}(n/\epsilon)$, $p > Cn\sqrt{m}$, and so $p > C\gamma m\sqrt{n}$, each term on the right hand side is less than $1/72$ (if C is sufficiently large). Then the left hand side is less than $1/24$ and D accepts y . \square

4 Well-conditioned submatrices of random matrices

We now present and analyze the simple greedy algorithm used in step 1 in Reduction 1.

Proposition 9. *Let $B \in \mathbb{R}^{m \times n}$ be a random Gaussian matrix with $m > Cn$. The probability that B contains a square submatrix with smallest singular value at least $1/\sqrt{n}$ is $1 - \exp(-\Omega(m))$. Moreover this submatrix can be found efficiently.*

Think of the column vectors of B as a stream of random normal vector samples. The matrix A is constructed incrementally column by column, starting with the empty matrix. After $k - 1$ columns of A have been chosen, the next sample from the stream is considered as a candidate for the k -th column. It is rejected unless

$$\rho = \sum_{i=1}^k \frac{1}{\sigma_i^2} \leq \frac{k}{n - k + 1}, \quad (5)$$

where $\sigma_1, \dots, \sigma_k$ are the singular values of A .

Once all n columns of A have been chosen, (5) guarantees that the sum of inverse squares of the singular values is at most n , so the smallest singular value will be at least $1/\sqrt{n}$ as desired. It remains to argue that no more than $m - n$ rejections happen except with probability $\exp(-\Omega(m))$.

Evolution of ρ We analyze the evolution of ρ as columns are being added to A . Let A_k be any non-singular $n \times k$ matrix. Then

$$\rho(A_k) = \frac{\sum_{i=1}^k \prod_{j \neq i} \sigma_j^2}{\prod_{i=1}^k \sigma_i^2} = -\frac{\chi'_k(0)}{\chi_k(0)},$$

where $\chi_k(\lambda) = \det(A_k^\top A_k - \lambda I)$. Given A_k , let A_{k+1} be the random matrix obtained by appending a random normal column x to A_k .

Let $L \in \mathbb{R}^{k \times k}$ be an orthogonal matrix such that $L^\top A_k^\top A_k L = \text{diag}(\sigma_1^2, \dots, \sigma_k^2)$. It can be obtained from the singular value decomposition of A_k . The matrix $L' \in \mathbb{R}^{(k+1) \times (k+1)}$ given by $L' = \text{diag}(L, 1)$ is also orthogonal and

$$A_{k+1} L' = \begin{bmatrix} A_k & x \end{bmatrix} \cdot \begin{bmatrix} L & \\ & 1 \end{bmatrix} = \begin{bmatrix} A_k L & x \end{bmatrix}$$

Since the columns of $A_k L$ are orthogonal of length $\sigma_1, \dots, \sigma_k$, the columns of

$$A_k L \text{diag}(\sigma_1^{-1}, \dots, \sigma_k^{-1})$$

can be completed to an orthonormal basis C . The change of variables

$$y^\top = x^\top C$$

is then an isometry, so y_1, \dots, y_n are independent standard normals, and $\|y\| = \|x\|$. Then

$$L'^\top A_{k+1}^\top A_{k+1} L' = \begin{bmatrix} \sigma_1^2 & & & & \sigma_1 y_1 \\ & \sigma_2^2 & & & \sigma_2 y_2 \\ & & \ddots & & \vdots \\ & & & \sigma_k^2 & \sigma_k y_k \\ \sigma_1 y_1 & \sigma_2 y_2 & \dots & \sigma_k y_k & \|y\|^2 \end{bmatrix}$$

Therefore

$$\begin{aligned} \chi_{k+1}(\lambda) &= \det(A_{k+1}^\top A_{k+1} - \lambda I) \\ &= \det(L'^\top A_{k+1}^\top A_{k+1} L' - \lambda I) \\ &= (\|y\|^2 - \lambda) \prod_{i=1}^k (\sigma_i^2 - \lambda) - \sum_{i=1}^k \sigma_i^2 y_i^2 \prod_{j \neq i} (\sigma_j^2 - \lambda) \\ &= \chi_k(\lambda) \left(\|y\|^2 - \lambda - \sum_{i=1}^k \frac{\sigma_i^2 y_i^2}{\sigma_i^2 - \lambda} \right). \end{aligned}$$

We obtain the following recurrences:

$$\begin{aligned} \chi_{k+1}(0) &= \chi_k(0) \|y^{\perp k}\|^2 \\ \chi'_{k+1}(0) &= \chi'_k(0) \|y^{\perp k}\|^2 - \chi_k(0) \left(1 + \sum_{i=1}^k \frac{y_i^2}{\sigma_i^2} \right), \end{aligned}$$

where $y^{\perp k} = (y_{k+1}, \dots, y_n)$.

Claim 10. *If $(n - k + 1)\chi'_k(0) + k\chi_k(0) \geq 0$ then*

$$\mathbb{E}[(n - k)\chi'_{k+1}(0) + (k + 1)\chi_{k+1}(0) | A_k] \geq 0.$$

The claim follows from linearity of expectation using the facts $\mathbb{E}[y_i^2] = 1$ and $\mathbb{E}\|y^{\perp k}\|^2 = k$.

Proof of Proposition 9. We show that the number of samples required for each column of A is dominated by a geometric random variable whose success probability is some absolute constant p_\star . The expected number of samples required is then at most n/p_\star . By large deviation bounds

for geometric random variables [Jan18] the probability that more than Cn samples are required is then at most $\exp(-\Omega(Cnp_\star))$, assuming $C > 1/p_\star$.

For the first column of A to fulfill (5) its squared norm needs to be at least n . This is at least p_\star by Corollary 12 (with $a_1 = \dots = a_n = 1$ and $b = 0$).

Now suppose (5) holds after the k -th column was added. Fix A_k and let X be the random variable $(n-k)\chi'_{k+1}(0) + (k+1)\chi_{k+1}(0)$. By Claim 10 $\mathbb{E}[X] \geq 0$. The random variable X is of the form in Corollary 12 so $\Pr(X > \mathbb{E}[X]) \geq p_\star$. Once a column x has been picked so that $X \geq 0$, the invariant (5) will hold for the matrix $A_{k+1} = [A_k \ x]$. \square

4.1 Anticoncentration

The concentration Q of a real-valued random variable X is $Q(X, h) = \sup_x \Pr(x \leq X \leq x+h)$.

Proposition 11. *There exists an absolute constant C such that if X_1, \dots, X_n are independent mean zero, unit variance random variables such that $Q(X_i, h) \leq 3/4$ for all i and some $h \leq 1/4C$ then*

$$\Pr(a_1 X_1 + \dots + a_n X_n > 0) \geq \frac{h^2}{32 + 4h^2}.$$

for all a_1, \dots, a_n .

Corollary 12. *There is an absolute constant p_\star so that for every n and a_1, \dots, a_n, b ,*

$$\Pr(a_1 Z_1^2 + \dots + a_n Z_n^2 + b \geq \mu) \geq p_\star,$$

where Z_1, \dots, Z_n are independent normals and $\mu = a_1 + \dots + a_n + b$.

Proof. Apply Proposition 11 to the random variables $Y_i = (X_i^2 - 1)/\sqrt{2}$ which have mean zero and unit variance. The condition $Q(Y_i, h) \leq 3/4$ is satisfied for all $h \leq 0.2$. \square

Proof of Proposition 11. Let $X = a_1 X_1 + \dots + a_n X_n$. We may assume X has unit variance. By Rogozin's inequality [Rog61],

$$Q(X, H) \leq CH \left(\sum a_i^2 (1 - Q(a_i X_i, a_i h)) \right)^{-1/2} = 2CH \leq 2Ch,$$

where $H = h \max_i |a_i| \leq h$. Applying Claim 13 we get

$$\Pr[X > 0] \geq \frac{1}{t+h} (h(1 - 2Ch) - 2/t) = \frac{h/2 - 2/t}{t+h}.$$

Choosing $t = 8/h$ we get $\Pr(X > 0) \geq h^2/(32 + 4h^2)$. \square

Claim 13. *For every zero-mean, unit-variance X , every $\lambda > 0$, and every $t \geq 1$*

$$\Pr[X > 0] \geq \frac{1}{t+h} (h \cdot \Pr(-h < X \leq 0) - 2/t).$$

Proof. Let $p = \Pr(X \in (0, t])$ and $q = \Pr(X \in (-h, 0])$. Then

$$\begin{aligned} \mathbb{E}[X] &\leq -h \Pr(X \leq -h) + 0 \Pr(-h < X \leq 0) + t \Pr(0 < X \leq t) + \mathbb{E}[X1(X > t)] \\ &\leq -h \cdot (1 - q - p) + t \cdot p + \mathbb{E}[X1(X > t)]. \end{aligned}$$

As $\mathbb{E}[X] = 0$,

$$p \geq \frac{1}{t+h} (h(1 - q) - \mathbb{E}[X1(X > t)]).$$

By Claim 14, $\mathbb{E}[X1(X > t)] \leq \mathbb{E}[|X|1(|X| > t)] \leq 2/t$. \square

Claim 14. For every zero-mean, unit-variance X and every $t \geq 1$,

$$\mathbb{E}[|X|1(|X| > t)] \leq 2/t.$$

Proof.

$$\begin{aligned} \mathbb{E}[|X|1(|X| > t)] &= \int_0^\infty \Pr(|X|1(|X| > t) > x)dx \\ &= \int_0^t \Pr(|X| > t)dx + \int_t^\infty \Pr(|X| > x)dx \\ &\leq \int_0^t (1/t^2)dx + \int_t^\infty (1/x^2)dx \\ &= 2/t. \end{aligned}$$

The inequality is Chebyshev's. □

5 Refutation via Boolean combinations

Theorem 1 was proved by adapting the Ajtai-Dwork encryption scheme into a refutation algorithm for NBV. In this Section we carry out an analogous analysis for the “pancake encryption” of Bogdanov, Cueto Noval, Hoffmann, and Rosen (BCHR).

Their public key is also computationally indistinguishable from a random subspace of \mathbb{R}^m . The dimension of this subspace is, however, only $o(\sqrt{m})$. As a consequence, the resulting refutation only applies to a regime of NBV that is efficiently tractable.

While BCHR becomes insecure when $n \gg \sqrt{m}$, we believe that a modification of it may be secure up to $n = m^{1-o(1)}$. The advantage of the BCHR-based reduction over Theorem 1 is that it applies to larger completeness error γ .

Theorem 15. For every constant ϵ there exists a constant C such that for all but an ϵ -fraction of instances, average-case NBV with parameters $m = C(n \log n)^2$ and $\gamma = 1/C\sqrt{m}$ is in SZK.

Let Z be a normal random variable and let $\zeta_1 < \dots < \zeta_r$ be the unique numbers such that $\Pr(Z \leq \zeta_i) = (2i + 1)/2r$. The Gaussian rounding $\text{round}_r: \mathbb{R} \rightarrow \{\zeta_1, \dots, \zeta_r\}$ is the function $\text{round}_r(z) = \zeta_i$ where i is the unique index for which $\lceil r \cdot \Pr(Z \leq z) \rceil = \lceil r \cdot \Pr(Z \leq \zeta_i) \rceil$ (see Figure 1). For $z \in \mathbb{R}^n$ let $\text{round}_r: \mathbb{R}^n \rightarrow \{\zeta_1, \dots, \zeta_r\}^n$ be given by $\text{round}_r(z) = (\text{round}_r(z_1), \dots, \text{round}_r(z_n))$. Set $r = \max\{Cm, Cn^2/\gamma^2\}$.




Figure 1: The function round_r for $r = 7$. All intervals have equal Gaussian measure. The values in the i -th interval round to ζ_i .

Reduction 2: On input $A \in \mathbb{R}^{m \times n}$,

- 1 If the largest singular value of A is more than $2\sqrt{m}$, fail.
- 2 If the smallest singular value of A is less than $\sqrt{m}/4$, fail.
- 3 Otherwise, output the sampler S that maps $x \sim \{\pm 1/\sqrt{m}\}^m$ to $\text{round}_r(Ax)$.

Theorem 15 follows from Claims 16 and 19.

Claim 16. For every ϵ there is a C so that for a $1 - \epsilon$ fraction of instances $A \in \mathbb{R}^{m \times n}$, where $m = (Cn \log n)^2$, the output of S is $2/3$ -close to random.

Fact 17. [BCHR22] The distribution $(A, \text{round}_r(Ax))$ is $\sqrt{4en \ln r / \sqrt{m}}$ -close to (A, ζ) , where ζ is uniform over rounded values and independent of A .

Fact 18. [DS01] Assume $m > 2n$. The largest and smallest singular values of A is at most $2\sqrt{n}$ and at least $\sqrt{n}/4$, except with probability $\exp(-\Omega(n))$.

Proof of Claim 16. By Fact 17, the joint distribution of A and the output of the sampler is $O(C^{-1/2})$ -close to uniform. Therefore for all but $O(C^{-1/2})$ choices of A the output is $2/3$ -close to uniform. By a Chernoff bound and Fact 18 at most $2^{-\Omega(m)}$ other inputs A cause the reduction to fail. \square

Claim 19. If A is a yes instance of NBV with $\gamma < 1/C\sqrt{m}$, either Reduction 2 fails, or the output of S is $2/3$ -far from random.

Fact 20. [BCHR22] For sufficiently large r , $\text{round}_r(z)$, $z \in \mathbb{R}$ is $r^{-1/2}$ -close to z unless $|z| > t$ for t such that $\Pr(|Z| > t) \leq 3(r \ln r)^{-1/2}$, where Z is normal in \mathbb{R} .

Proof of Claim 19. Let $w \in \mathbb{R}^n$ be the witness for which $wA = v + e$ where $v \in \{\pm 1\}^m$ and $\|e\| \leq \gamma\sqrt{m}$. Let D be the distinguisher that, given ζ , accepts if $|\{\sqrt{m}\langle w, \zeta \rangle\}| \leq 1/48$.

Assuming the reduction did not fail, by the assumption on singular values,

$$\frac{1}{4} \leq \frac{\|v\| - \|e\|}{2\sqrt{m}} \leq \|w\| \leq \frac{\|v\| + \|e\|}{\sqrt{m}/4} \leq 8.$$

If ζ is random, we argue that D rejects with probability at least $5/6$. we can write $\zeta = \text{round}_r(g)$ for a normal $g \in \mathbb{R}^n$. Let $e = \text{round}_r(g) - g$. Then $\sqrt{m}\langle w, \zeta \rangle = \sqrt{m}\langle w, g \rangle + \sqrt{m}\langle w, e \rangle$. The random variable $\sqrt{m}\langle w, g \rangle$ is a univariate normal with standard deviation at least $\sqrt{m}\|w\| \geq \sqrt{m}/4$. By Fact 6, $\{\sqrt{m}\langle w, g \rangle\}$ is $2^{-\Omega(m)} < 1/24$ close to uniform in $[-1/2, 1/2]$. In particular, $|\{\sqrt{m}\langle w, g \rangle\}| > 1/24$ except with probability $11/12 - 1/24$. By Fact 20, $\|e\|_\infty \leq r^{-1/2}$ except with probability $3n(r \ln r)^{-1/2} < 1/24$. Both events happen with probability at least $5/6$. Assuming this,

$$|\{\sqrt{m}\langle w, \zeta \rangle\}| > 1/24 - |\{\sqrt{m}\langle w, e \rangle\}| \geq 1/24 - \sqrt{m}\|w\|\|e\| > 1/48$$

because $\sqrt{m}\|w\|\|e\| \leq 8\sqrt{m}r^{-1/2}$ and D rejects.

If ζ is the output of the sampler we argue that the distinguisher accepts it with probability at least $8/9$:

$$|\{\sqrt{m}\langle w, Ax \rangle\}| = |\{\sqrt{m}\langle v + e, x \rangle\}| = |\{\sqrt{m}\langle v, x \rangle + \sqrt{m}\langle e, x \rangle\}| = \sqrt{m}|\langle e, x \rangle| \quad (6)$$

because v and $\sqrt{m}x$ are integral. As x is random, $\mathbb{E}[\langle e, x \rangle^2] = \|e\|^2/m$. By Markov's inequality, $|\langle e, x \rangle| \leq 3\|e\|/\sqrt{m}$ except with probability $1/9$. If this holds (6) is at most $3\|e\| \leq 3\gamma\sqrt{m}$.

As the largest singular value of A is at most $2\sqrt{m}$, all entries of Ax are between -2 and 2 . By Fact 20, $\|\text{round}_r(Ax) - Ax\|_\infty \leq nr^{-1/2}$. Therefore

$$|\{\sqrt{m}\langle w, \text{round}_r(Ax) - Ax \rangle\}| \leq \sqrt{m}\|w\|\|\text{round}_r(Ax) - Ax\| \leq 8\sqrt{m}nr^{-1/2} \leq \gamma\sqrt{m}.$$

Together with (6), $|\{\sqrt{m}\langle w, Ax \rangle\}| \leq 4\gamma\sqrt{m} \leq 1/48$. \square

Acknowledgements

We are grateful to Chris Jones for useful advice and feedback. Part of this work was done while the first author was affiliated with and the second author visited the Chinese University of Hong Kong.

Work supported by European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (Grant agreement No. 101019547). The first author was additionally supported by Hong Kong RGC GRF grant CUHK 14209920 and NSERC grant RGPIN-2023-05006. The second author was additionally supported by Cariplo CRYPTONOMEX grant.

References

- [AD97] Miklós Ajtai and Cynthia Dwork. A public-key cryptosystem with worst-case/average-case equivalence. In *Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing*, 1997.
- [AR05] Dorit Aharonov and Oded Regev. Lattice problems in $\text{NP} \cap \text{coNP}$. *J. ACM*, 52(5):749–765, sep 2005.
- [BB20] Matthew S. Brennan and Guy Bresler. Reducibility and statistical-computational gaps from secret leakage. In *Conference on Learning Theory, COLT 2020*, volume 125 of *Proceedings of Machine Learning Research*, pages 648–847. PMLR, 2020.
- [BCHR22] Andrej Bogdanov, Miguel Cueto Noval, Charlotte Hoffmann, and Alon Rosen. Public-key encryption from homogeneous clwe. Cryptology ePrint Archive, Paper 2022/093, 2022. <https://eprint.iacr.org/2022/093>.
- [BHK⁺19] Boaz Barak, Samuel Hopkins, Jonathan Kelner, Pravesh K. Kothari, Ankur Moitra, and Aaron Potechin. A nearly tight sum-of-squares lower bound for the planted clique problem. *SIAM Journal on Computing*, 48(2):687–735, 2019.
- [DS01] Kenneth R. Davidson and Stanislaw J. Szarek. Chapter 8 - local operator theory, random matrices and banach spaces. In W.B. Johnson and J. Lindenstrauss, editors, *Handbook of the Geometry of Banach Spaces*, volume 1 of *Handbook of the Geometry of Banach Spaces*, pages 317–366. Elsevier Science B.V., 2001.
- [FKO06] U. Feige, J. H. Kim, and E. Ofek. Witnesses for non-satisfiability of dense random 3CNF formulas. In *47th Annual IEEE Symposium on Foundations of Computer Science*, 2006.
- [GJJ⁺20] Mrinalkanti Ghosh, Fernando Granha Jeronimo, Chris Jones, Aaron Potechin, and Goutham Rajendran. Sum-of-squares lower bounds for sherrington-kirkpatrick via planted affine planes. In Sandy Irani, editor, *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*, pages 954–965. IEEE, 2020.
- [GKM22] Venkatesan Guruswami, Pravesh K. Kothari, and Peter Manohar. Algorithms and certificates for boolean CSP refutation: smoothed is no harder than random. In *54th Annual ACM SIGACT Symposium on Theory of Computing*, 2022.

- [Gol11] Oded Goldreich. *Average Case Complexity, Revisited*, pages 422–450. Springer Berlin Heidelberg, Berlin, Heidelberg, 2011.
- [Gri01] Dima Grigoriev. Linear lower bound on degrees of positivstellensatz calculus proofs for the parity. *Theor. Comput. Sci.*, 259(1):613–622, may 2001.
- [GSV99] Oded Goldreich, Amit Sahai, and Salil Vadhan. Can statistical zero knowledge be made non-interactive? or on the relationship of szk and nisz. In Michael Wiener, editor, *Advances in Cryptology — CRYPTO’99*, pages 467–484, Berlin, Heidelberg, 1999. Springer Berlin Heidelberg.
- [ILL89] R. Impagliazzo, L. A. Levin, and M. Luby. Pseudo-random generation from one-way functions. In *Proceedings of the Twenty-First Annual ACM Symposium on Theory of Computing, STOC ’89*, page 12–24, New York, NY, USA, 1989. Association for Computing Machinery.
- [Jan18] Svante Janson. Tail bounds for sums of geometric and exponential variables. *Statistics & Probability Letters*, 135:1–6, 2018.
- [Lev86] Leonid A. Levin. Average case complete problems. *SIAM J. Comput.*, 15(1):285–286, 1986.
- [LLL82] H.W. Jr. Lenstra, A.K. Lenstra, and L. Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261:515–534, 1982.
- [LM09] Vadim Lyubashevsky and Daniele Micciancio. On bounded distance decoding, unique shortest vectors, and the minimum distance problem. In Shai Halevi, editor, *Advances in Cryptology - CRYPTO 2009, 29th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2009. Proceedings*, volume 5677 of *Lecture Notes in Computer Science*, pages 577–594. Springer, 2009.
- [Mon19] Andrea Montanari. Optimization of the Sherrington-Kirkpatrick Hamiltonian. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, 2019.
- [MR07] Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on gaussian measures. *SIAM Journal on Computing*, 37(1):267–302, 2007.
- [MRX20] Sidhanth Mohanty, Prasad Raghavendra, and Jeff Xu. Lifting sum-of-squares lower bounds: Degree-2 to degree-4. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020*, page 840–853, New York, NY, USA, 2020. Association for Computing Machinery.
- [Par79] G. Parisi. Infinite number of order parameters for spin-glasses. *Phys. Rev. Lett.*, 43:1754–1756, Dec 1979.
- [Pei09] Chris Peikert. Public-key cryptosystems from the worst-case shortest vector problem: Extended abstract. In *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing, STOC ’09*, page 333–342, New York, NY, USA, 2009. Association for Computing Machinery.
- [PTVW22] Aaron Potechin, Paxton Turner, Prayaag Venkat, and Alexander S. Wein. Near-optimal fitting of ellipsoids to random points, 2022.

- [Reg09] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6), 9 2009.
- [Rog61] B. A. Rogozin. On the increase of dispersion of sums of independent random variables. *Theory of Probability & Its Applications*, 6(1):97–99, 1961.
- [Sch08] Grant Schoenebeck. Linear level lasserre lower bounds for certain k-csp. In *Proceedings of the 2008 49th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '08, page 593–602, USA, 2008. IEEE Computer Society.
- [SV03] Amit Sahai and Salil Vadhan. A complete problem for statistical zero knowledge. *J. ACM*, 50(2):196–249, Mar 2003.
- [Tal11] Michel Talagrand. *The Parisi Formula*, pages 349–474. Springer Berlin Heidelberg, Berlin, Heidelberg, 2011.
- [Tre10] Luca Trevisan. The program-enumeration bottleneck in average-case complexity theory. In *2010 IEEE 25th Annual Conference on Computational Complexity*, pages 88–95, 2010.