

The Online Threshold for Contracting Hypergrid Vectors

An exact constant, $\sqrt{\pi/8}$, conjectured to separate what online algorithms can and cannot do

Status. Statement: AI-written from Andrej Bogdanov, Alon Rosen, Neekon Vafa and Vinod Vaikuntanathan, *Adaptive Robustness of Hypergrid Johnson-Lindenstrauss*, IACR ePrint 2025/666. Not yet checked by a human. This is the source’s Conjecture 1, the “online threshold conjecture”. The source gives an online algorithm achieving $\kappa = O(\sqrt{\alpha}/B)$ and an overlap-gap argument against stable algorithms; the conjecture pins the constant in front, $\sqrt{\pi/8} \approx 0.627$, as the exact place where online algorithms stop working.

Category. *Fine-Grained Cryptography.*

Conjecture (informal). *The Johnson-Lindenstrauss lemma says a random projection nearly preserves lengths. “Nearly” leaves room: some vectors do get badly contracted, and if the vector is allowed to have arbitrary real coordinates they are easy to find. Restrict the coordinates to a bounded integer grid and the problem gets interesting, with a wide gap between the vectors that exist and the vectors an efficient algorithm can find. The source charts that gap and finds a specific constant, $\sqrt{\pi/8}$, showing up as the fixed point of a drift in the analysis of its own online algorithm. The conjecture is that this constant is not an artefact of the analysis but the truth: below it, no online algorithm succeeds.*

1 The setting

Definition 1 (Contracting Hypergrid Vector, [BRVV25, Definition 2]). For $n, m, B \in \mathbb{N}$ and $\kappa \in \mathbb{R}_{>0}$ with $m < n$, the CHV problem with parameters n, m, B, κ is: given $A \sim$

$\mathcal{N}(0,1)^{m \times n}$, find $x \in ([-B, B] \cap \mathbb{Z})^n$ with

$$\|Ax\|_2 < \kappa \|x\|_2 \sqrt{m}.$$

Write $\alpha := m/n < 1$ for the aspect ratio.

Any non-zero x achieves $\kappa = \Theta(1)$ in expectation, so the question is how small κ can be pushed. The two thresholds the source identifies:

$$\kappa_{\text{stat}} = \Theta((2B+1)^{-1/\alpha}), \quad \kappa_{\text{comp}} = \tilde{\Theta}(\sqrt{\alpha}/B).$$

Above κ_{stat} solutions exist; below it they do not. κ_{comp} is the best κ the source's algorithms reach, and the gap between the two is what it exploits: it builds robust locality-sensitive hash functions from the hardness of CHV, which in turn imply collision resistance — so the statistical-computational gap yields cryptography beyond one-way functions.

Two algorithms establish the upper side. A kernel-rounding algorithm samples a random x in the kernel, scales, and rounds. An online algorithm — processing the columns of A in sequence and producing the entries of x one at a time, inspired by Bansal and Spencer — achieves $\kappa = O(\sqrt{m/n}/B)$ whenever $n \geq Km \log B$. Together they give $\kappa = O(\sqrt{\alpha} \log B/B)$ for all $n > m$.

2 Notation and parameters

- B bounds the coordinates of the solution: $x \in ([-B, B] \cap \mathbb{Z})^n$. It may be as small as 1 or polynomially large in n .
- $\alpha = m/n < 1$ is the aspect ratio; κ is the quality of the solution as in Definition 1.
- An algorithm is *online* if it processes the columns of A in sequence, fixing each entry of x before seeing the rest of the input.
- δ and ε are the conjecture's slack and success-fraction parameters.

3 The conjecture

Conjecture 1 (Online threshold conjecture, [BRVV25, Conjecture 1]). *For every δ and B there exists a sufficiently small α and ε so that every online algorithm fails to find $x \in ([-B, B] \cap \mathbb{Z})^n$ such that*

$$\frac{\|Ax\|}{\sqrt{m} \|x\|} < (\sqrt{\pi/8} - \delta) \frac{\sqrt{\alpha}}{B}$$

for at least an ε fraction of αn by n matrices A , for all sufficiently large n .

Remark 1 (Where the constant comes from). It is not chosen for convenience. In the analysis of the source’s online algorithm a process U has drift pushing it toward a fixed point $U = \frac{1}{2\mu} = \sqrt{\pi/8} \approx 0.627$, where $\mu = \sqrt{2/\pi}$ is the mean of $|\mathcal{N}|$. Conjecture 1 asserts that this fixed point is the genuine barrier for the whole class of online algorithms, not an artefact of one algorithm’s analysis. That is what makes it a sharp-threshold claim rather than an order-of-magnitude one, and it is why a resolution has to control the constant and not merely the exponent.

Remark 2 (The stability caveat is part of the claim). Immediately before stating the conjecture the source notes that its overlap-gap argument assumes the algorithm is committed to the approximate norm of the solution x before seeing A , that its own online algorithm and those of Bansal and Spencer satisfy this, and that “we believe that some assumption of this type is necessary for an OGP-based argument to ensure stability.” Conjecture 1 as printed quantifies over “every online algorithm” without repeating that condition. A resolution should say which reading it settles; a counterexample that is online but not norm-committed would be informative either way, and would bear directly on whether the assumption is necessary.

Remark 3 (Two lower bounds, and the gap between them). The source proves hardness twice, by different means, and the results are not comparable. Its overlap-gap theorem gives evidence that CHV is hard in the regime $\alpha \gg B^2 \kappa^2 \log(1/\kappa)$, against stable algorithms. Its lattice-based theorem gives, under the worst-case polynomial

hardness of standard lattice problems, that no polynomial-time algorithm solves CHV for $\kappa = O(1/(Bn^{1/2+\varepsilon}))$ — a statement about *all* efficient algorithms, but quantitatively weaker and for a more restricted parameter range. The source records: “We leave it as a fascinating open question as to whether this lower bound can be improved to match that of Theorem 3 or similar.” That is a second open problem in the same paper and it is not this statement, though a resolution of either informs the other.

Remark 4 (Neighbouring problems this is not). CHV is close to two studied problems and differs from each in a stated way. In the Symmetric Binary Perceptron and the Nearest Boolean Vector problem the domain of x is $\{-1, 1\}^n$ rather than $([-B, B] \cap \mathbb{Z})^n$, so the norm of x is fixed and x cannot have zero entries; and SBP bounds $\|Ax\|_\infty$ rather than $\|Ax\|_2$. In the special case $B = 1$ both the Bansal and Spencer algorithm and the overlap-gap bound of Gamarnik, Kızıldağ, Perkins and Xu match the source’s. So B is the axis along which this statement is new.

4 Bibliography

- [BRVV25] Andrej Bogdanov, Alon Rosen, Neekon Vafa and Vinod Vaikanathan. *Adaptive robustness of hypergrid Johnson-Lindenstrauss*. IACR ePrint 2025/666. The source. Definition 1 is on page 4 and Definition 2 on page 12; Theorem 2 (the online algorithm) and the informal statement of the conjecture are on page 6; Conjecture 1 is stated formally on page 16; Theorem 10 (the lattice-based bound) is on page 23. [UNVERIFIED] — the author list was not transcribed during this harvest and is deliberately not guessed here; a reviewer should fill it in from the ePrint record before citing.
- [BS20] Nikhil Bansal and Joel Spencer. *On-line balancing of random inputs*. Random Structures and Algorithms, 2020. The online discrepancy algorithm the source’s is a variant of, and whose guarantee matches the source’s in the case $B = 1$.
- [GKPX22] David Gamarnik, Eren C. Kızıldağ, Will Perkins and Changji Xu. *Algorithms and barriers in the symmetric binary perceptron model*. FOCS 2022. The overlap-gap bound that matches the source’s at $B = 1$.
- [APZ19] Benjamin Aubin, Will Perkins and Lenka Zdeborová. *Storage capacity in symmetric binary perceptrons*. Journal of Physics A, 2019. Introduces the Symmetric Binary Perceptron, the closest studied neighbour of CHV.

- [MRX20] Sidhanth Mohanty, Prasad Raghavendra and Jeff Xu. *Lifting sum-of-squares lower bounds: degree-2 to degree-4*. STOC 2020. Cited by the source as introducing the Nearest Boolean Vector problem. [UNVERIFIED] — cited in the source as [MRX20]; the bibliographic detail here is inferred.