Counterexample to Majority Optimality in NICD with Erasures

Paata Ivanisvili [1] **Xinyuan Xie** [1]

Abstract

We asked GPT-5 Pro to look for counterexamples among a public list of open problems (the Simons "Real Analysis in Computer Science" collection). After several numerical experiments, it suggested a counterexample for the *Non-Interactive Correlation Distillation (NICD) with erasures* question: namely, a Boolean function on 5 bits that achieves a strictly larger value of $\mathbb{E}|f(z)|$ than the 5-bit majority function when the erasure parameter is $p = 0.40$. In this very short note we record the finding, state the problem precisely, give the explicit function, and verify the computation step by step *by hand* so that it can be checked without a computer. In addition, we show that for each fixed odd n the majority is optimal (among unbiased Boolean functions) in a neighborhood of $p = 0$. We view this as a "little spark" of an AI contribution in Theoretical Computer Science: while modern Large Language Models (LLMs) often assist with literature and numerics, here a concrete finite counterexample emerged.

A Short Preface

Finding a successful instance in which an AI suggests a viable counterexample to a decade-plus-old open question appears, to our knowledge, to be the first publicly reported case. In most problems we tested, the models contributed little beyond literature pointers and small-scale numerics. Here, by contrast, the outcome is a clear win; had someone shown us this instance beforehand, we would have been impressed. The result also echoes the now-classical counterexamples (Gopi, 2013; Jain, 2017; Biswas and Sarkar, 2023) related to "Majority is Least Stable" for linear threshold functions (LTFs): even if an AI merely applied a known LTF counterexample pattern in a new context and verified it, that would still merit credit.

[1]Department of Mathematics, University of California, Irvine, CA 92697, USA. Correspondence to: Paata Ivanisvili <pivanisv@uci.edu>, Xinyuan Xie <xinyuax7@uci.edu>.

The 3rd *AI for Math Workshop* at the 43rd International Conference on Machine Learning (ICML), Seoul, South Korea, 2026.
Copyright 2026 by the author(s).

1. Problem statement

We work on the hypercube $\{-1,1\}^n$ equipped with the uniform probability measure. A Boolean function is $f : \{-1,1\}^n \to \{-1,1\}$. We say f is *unbiased* if $\mathbb{E}[f] = 0$. A sufficient condition is *oddness*: $f(-x) = -f(x)$.

Multilinear extension. Every f has a unique multilinear polynomial representation

$$f(z) = \sum_{S\subseteq[n]} \widehat{f}(S) \prod_{i\in S} z_i, \qquad z \in [-1,1]^n,$$

where $\widehat{f}(S) = \mathbb{E}[f(x)\prod_{i\in S} x_i]$, and $x \sim \mathrm{unif}(\{-1,1\}^n)$.

Erasure model. Fix $p \in (0,1)$. Let $z \in \{-1,0,1\}^n$ have independent coordinates with

$$\Pr[z_i = 1] = \Pr[z_i = -1] = \tfrac{p}{2}, \qquad \Pr[z_i = 0] = 1 - p.$$

This is the standard *binary erasure* NICD model (cf. Yang (Yang, 2007)).[1]

Objective. For unbiased Boolean f, define

$$\Phi_p(f) \;:=\; \mathbb{E}_z\big[\,|f(z)|\,\big],$$

where $f(z)$ is the multilinear extension evaluated at the erasure sample z.

Open question (as recorded in the Simons list). For $p < \frac{1}{2}$, is $\Phi_p(f)$ maximized (over unbiased f) by a majority function? (When $p \geq \frac{1}{2}$, dictators are optimal; see (O'Donnell and Wright, 2012; O'Donnell, 2012).) We refer to the Simons compilation (Filmus et al., 2014) for a concise statement.

[1]Yang introduced and analyzed NICD for various noise models, including the erasure model; the maximization question we study below is how it is recorded in later Boolean-analysis sources and in the Simons open-problems list.

2. A Finite Counterexample at $p = 0.40$

Set $n = 5$ and $p = \frac{2}{5} = 0.40$. Consider the linear threshold function

$$f(x_1, \ldots, x_5) \;=\; \mathrm{sgn}\big(x_1 - 3x_2 + x_3 - x_4 + 3x_5\big)$$

(with the convention $\mathrm{sgn}(t) = 1$ for $t > 0$, -1 for $t < 0$; note $t \neq 0$ in our case). Then $f(-x) = -f(x)$, hence f is unbiased.

Theorem 1. *For $p = 0.40$ one has*

$$\Phi_{0.40}(f) = \frac{2689}{6250} = 0.43024,$$
$$\Phi_{0.40}(\mathrm{Maj}_5) = \frac{5363}{12500} = 0.42904,$$

hence $\Phi_{0.40}(f) > \Phi_{0.40}(\mathrm{Maj}_5)$. In particular, majority is not optimal at $p = 0.4$ already for $n = 5$.

Proof. See Appendix A. □

3. A Related (Solved) Conjecture: Majority Is Least Stable

The "Majority-is-Least-Stable" conjecture states: for all linear threshold functions $f : \{-1, 1\}^n \to \{-1, 1\}$, all odd n, and all $p \in [0, 1]$, the noise stability

$$\mathrm{Stab}_p[f] \;:=\; \sum_{S \subseteq [n]} p^{|S|}\, \widehat{f}(S)^2$$

is minimized by the majority function Maj_n. Equivalently, with z drawn from the erasure model at rate p, one has $\mathrm{Stab}_p(f) = \mathbb{E}\big[f(z)^2\big]$, so the conjecture asserts that the majority functions minimize the L_2-norm of $f(z)$ among LTFs. By contrast, the NICD-with-erasures question asks whether majority *maximizes* the L_1-norm $\Phi_p(f) = \mathbb{E}|f(z)|$ among unbiased f. Although $\mathbb{E}|f(z)| \leq \sqrt{\mathbb{E}f(z)^2}$, to the best of our knowledge there is no implication in either direction between these two optimization problems. After identifying the counterexample presented here, we noted that it is closely related in spirit to the counterexample in (Gopi, 2013; Jain, 2017), e.g. to $\mathrm{sgn}(2x_1 + 2x_2 + x_3 + x_4 + x_5)$; both have the "two heavy + three light" LTF structure, and up to permutation and flips of variables, they are the same function.

Remark 2. *Comparison between L_1 norms is often more delicate than comparison between L_2 norms due to the lack of orthogonality. This is reflected in Appendix A: the manual verification for L_1 occupies more casework than the corresponding L_2 verification in (Jain, 2017). This may help explain why a counterexample familiar from the "Majority-is-Least-Stable" setting had not, to our knowledge, been tested against the NICD-with-erasures question until now.*

Remark 3. *In the Majority-is-Least-Stable question (L_2 norm), the counterexample works in a neighborhood of $p = 0$. On the other hand, in the NICD-with-erasures question there is no counterexample in a neighborhood of $p = 0$ for each fixed n. Indeed, in this case we show that majority is optimal (see Section 4).*

Remark 4. *A brute-force search is feasible for this finite question; indeed, GPT-5 Pro noted in the chat history that it used such a method. Still, being able to program against an open-problems list and to uncover a concrete counterexample seems noteworthy to us.*

4. For Each Fixed n Majority Is Optimal in a Neighborhood of $p = 0$

Lemma 5. *For unbiased $f : \{\pm 1\}^n \to \{\pm 1\}$ we have*

$$\Big|\Phi_p(f) - p \sum_{i=1}^{n} |\hat{f}(i)|\Big| \;\leq\; \Big((n-1)\sqrt{n} + \tbinom{n}{2}\Big)p^2.$$

Proof. Let $K := |\{i : Z_i \neq 0\}|$, so $K \sim \mathrm{Bin}(n, p)$. By conditioning on K,

$$\Phi_p(f) = \sum_{k=0}^{n} \Pr[K = k]\; \mathbb{E}[\,|f(Z)| \mid K = k]\,.$$

We treat three cases.

Case $K = 0$. Then $Z = 0$ and $|f(Z)| = |\hat{f}(\emptyset)| = 0$ since f is unbiased.

Case $K = 1$. We have

$$\mathbb{E}[\,|f(Z)| \mid K = 1] = \frac{1}{n} \sum_{i=1}^{n} |\hat{f}(i)|,$$

hence the unconditional contribution from $K = 1$ equals

$$\Pr[K = 1] \cdot \mathbb{E}[\,|f(Z)| \mid K = 1] = p(1-p)^{n-1} \sum_{i=1}^{n} |\hat{f}(i)|.$$

Therefore

$$\begin{aligned}
&\big|\, p(1-p)^{n-1} - p \,\big| \sum_{i=1}^{n} |\hat{f}(i)| \\
&\leq p\big(1 - (1-p)^{n-1}\big) \sum_{i=1}^{n} |\hat{f}(i)| \\
&\leq (n-1)\sqrt{n}\, p^2,
\end{aligned}$$

using $1 - (1-p)^{n-1} \leq (n-1)p$ and $\sum_i |\hat{f}(i)| \leq \sqrt{n}$ (by Cauchy–Schwarz).

Case $K \geq 2$**.** Since $|f(Z)| \leq 1$,

$$\sum_{k\geq 2} \Pr[K=k]\, \mathbb{E}[\,|f(Z)| \mid K=k] \leq \Pr[K \geq 2] \leq \sum_{1\leq i<j\leq n} \Pr[Z_i \neq 0,\, Z_j \neq 0] = \binom{n}{2} p^2.$$

Combining the three cases gives

$$\left|\Phi_p(f) - p\sum_{i=1}^{n} |\hat{f}(i)|\right| \leq (n-1)\sqrt{n}\, p^2 + \binom{n}{2} p^2,$$

as claimed. □

Lemma 6 ((O'Donnell, 2021), Thm. 2.33)**.** *Let n be odd. Among all Boolean functions $f : \{-1,1\}^n \to \{-1,1\}$, the quantity*

$$\sum_{i=1}^{n} \widehat{f}(i)$$

is uniquely maximized by the majority function

$$\mathrm{Maj}_n(x) \;=\; \mathrm{sgn}(x_1 + \cdots + x_n).$$

Theorem 7. *Fix odd n. There exists $p_0 = p_0(n) > 0$ such that for all $p \in (0, p_0)$, the function Maj_n is the unique maximizer of Φ_p over unbiased Boolean f; i.e.,*

$$\Phi_p(\mathrm{Maj}_n) > \Phi_p(f) \qquad \textit{for all unbiased } f \neq \mathrm{Maj}_n.$$

Proof. By Lemma 5,

$$\Phi_p(g) = p\sum_{i=1}^{n} |\widehat{g}(i)| + E_g(p), \qquad |E_g(p)| \leq C_n\, p^2,$$

for every unbiased g, where $C_n := (n-1)\sqrt{n} + \binom{n}{2}$. Since Φ_p and $\sum_i |\hat{g}(i)|$ are invariant under flipping individual coordinates, we may assume w.l.o.g. that all level-1 coefficients are nonnegative for the competitor f, so $\sum_i |\widehat{f}(i)| = \sum_i \widehat{f}(i)$ and likewise for Maj_n.

Let

$$\delta_n \;:=\; \min_{f\neq \mathrm{Maj}_n} \left(\sum_{i=1}^{n} \widehat{\mathrm{Maj}_n}(i) - \sum_{i=1}^{n} \widehat{f}(i)\right) \;>\; 0,$$

which is strictly positive by Lemma 6 and finiteness of the domain $\{-1,1\}^n$ for fixed n (uniqueness of the maximizer implies a positive gap). Then for any unbiased $f \neq \mathrm{Maj}_n$,

$$\begin{aligned}\Phi_p(\mathrm{Maj}_n) - \Phi_p(f) &= p\left(\sum_i \widehat{\mathrm{Maj}_n}(i) - \sum_i \widehat{f}(i)\right) \\ &\quad + \big(E_{\mathrm{Maj}_n}(p) - E_f(p)\big) \\ &\geq p\,\delta_n - 2C_n\, p^2.\end{aligned}$$

Choose $p_0 := \min\{\delta_n/(4C_n),\, 1\}$. Then for all $p \in (0, p_0)$, we have $p\,\delta_n - 2C_n p^2 \geq \frac{1}{2} p\,\delta_n > 0$, proving the claim and the uniqueness. □

Acknowledgments

P.I. thanks EpochAI for encouraging broad testing of models, and University of Würzburg for its warm hospitality. The P.I. was supported in part by NSF CAREER grant DMS-2152401 and a Simons Fellowship. This work benefited from interactions with ChatGPT (GPT-5 Pro, OpenAI), which was used to find the counterexample. Following arXiv policy, the model is acknowledged here but not listed as an author. The authors thank Sivakanth Gopi for providing the reference of his counterexample of the "Majority-is-Least-Stable" conjecture.

References

A. Biswas, P. Sarkar. On the "majority is least stable" conjecture. *Information Processing Letters*, Volume 179, 2023.

Y. Filmus, H. Hatami, S. Heilman, E. Mossel, R. O'Donnell, S. Sachdeva, A. Wan, K. Wimmer. Real Analysis in Computer Science: A collection of Open Problems. Simons Institute (2014). https://simons.berkeley.edu/sites/default/files/openprobsmerged.pdf.

S. Gopi. Stability of Linear Threshold Functions. Undergraduate Thesis. 2013.https://gopisivakanth.github.io/files/UndergraduateThesis_IITBombay.pdf.

V. Jain. A Counterexample to the "Majority is Least Stable" Conjecture. Manuscript, 2017. Available at https://arxiv.org/abs/1703.07657.

R. O'Donnell and J. Wright. A new point of NP-hardness for Unique Games. *STOC '12: Proceedings of the Forty-Fourth Annual ACM Symposium on Theory of Computing*, ACM, New York, pp. 289–306, 2012.

R. O'Donnell. Open Problems in the Analysis of Boolean Functions. 2012. https://www.cs.cmu.edu/~odonnell/papers/aobf-open-problems.pdf.

R. O'Donnell. Analysis of Boolean Functions. 2021. https://arxiv.org/abs/2105.10386.

K. Yang. On the (im)possibility of non-interactive correlation distillation. *Theoretical Computer Science*, 382(2–3):157–166, 2007. (Erasure NICD model and bounds.)

Appendix A

The full chat conversation between P.I. and GPT-5 Pro (ChatGPT) is available here:

https://chatgpt.com/share/68e80fbd-1f94-8001-8eb3-32be3ccd1bf2

Manual verification.

Multilinear expansions of majority functions. The majority function $\mathrm{Maj}_n : \{-1,1\}^n \to \{-1,1\}$ is defined for all odd n as

$$\mathrm{Maj}_n(x_1,\dots,x_n) = \mathrm{sgn}(x_1+\cdots+x_n),$$

i.e., $\mathrm{Maj}_n(x) = 1$ if a majority of the coordinates are 1, and -1 otherwise. Its unique multilinear expansions for $n = 3, 5$ are

$$\mathrm{Maj}_3(x_1,x_2,x_3) = \tfrac{1}{2}(x_1+x_2+x_3) - \tfrac{1}{2}x_1x_2x_3,$$

$$\mathrm{Maj}_5(x) = \tfrac{3}{8}\sum_{i=1}^{5} x_i - \tfrac{1}{8}\sum_{1\le i<j<k\le 5} x_ix_jx_k + \tfrac{3}{8}\,x_1x_2x_3x_4x_5.$$

See Theorem 5.19 in (O'Donnell, 2021) for the general formula for the Fourier coefficients of majority when n is odd.

Multilinear expansion of the counterexample f. The counterexample found by GPT-5 Pro,

$$f(x_1,\dots,x_5) \;=\; \mathrm{sgn}\big(x_1 - 3x_2 + x_3 - x_4 + 3x_5\big),$$

admits the following factorized multilinear form:

$$f(x) = \tfrac{1}{4}\Big(2(-x_2+x_5) + (1+x_2x_5)\,(x_1+x_3-x_4+x_1x_3x_4)\Big).$$

Indeed, f acts as follows: if $-x_2+x_5 = 0$, the output is the majority of $(x_1, x_3, -x_4)$; otherwise it outputs the majority of $(-x_2, x_5)$. Thus

$$f(x) = \mathrm{Maj}_3\big(-x_2,\, x_5,\, \mathrm{Maj}_3(x_1,x_3,-x_4)\big).$$

Plugging in $\mathrm{Maj}_3(a,b,c) = \frac{1}{2}(a+b+c-abc)$ gives the displayed formula.

Expected value of $|\mathrm{Maj}_5(z)|$. Extend functions on $\{-1,1\}^5$ to $\{-1,0,1\}^5$ by evaluating their multilinear expansions at $z \in \{-1,0,1\}^5$. Let z be drawn from the erasure model and set

$$K = \#\{i : z_i \neq 0\} \sim \mathrm{Bin}(5,p),$$

$$\Pr(K=k) = \binom{5}{k} p^k (1-p)^{5-k}.$$

Then

$$\begin{aligned}
\mathbb{E}\,|\mathrm{Maj}_5(z)| &= \sum_{k=0}^{5} \mathbb{E}[\,|\mathrm{Maj}_5(z)| \mid K=k]\,\Pr(K=k)\\
&= 0\cdot(1-p)^5 + \tfrac{3}{8}\cdot 5p(1-p)^4\\
&\quad + \tfrac{3}{8}\cdot 10p^2(1-p)^3 + \tfrac{5}{8}\cdot 10p^3(1-p)^2\\
&\quad + \tfrac{5}{8}\cdot 5p^4(1-p) + 1\cdot p^5\\
&= \tfrac{15}{8}p - \tfrac{15}{4}p^2 + \tfrac{25}{4}p^3\\
&\quad - \tfrac{45}{8}p^4 + \tfrac{9}{4}p^5.
\end{aligned}$$

Here are the conditional expectations, by symmetry (we may fix which coordinates are nonzero and average over their signs). **Case** $K=0$: $z=(0,0,0,0,0)$, so $|\mathrm{Maj}_5(z)| = 0$. **Case** $K=1$: without loss of generality (by symmetry) $z_1 \neq 0$, so $|\mathrm{Maj}_5(z)| = \frac{3}{8}$. **Case** $K=2$: assume $z_1, z_2 \neq 0$, then $\mathrm{Maj}_5(z) = \frac{3}{8}(z_1+z_2)$, hence the conditional expectation is $\frac{3}{8}$. **Case** $K=3$: assume $z_1,z_2,z_3 \neq 0$, then $\mathrm{Maj}_5(z) = \frac{3}{8}(z_1+z_2+z_3) - \frac{1}{8}z_1z_2z_3$; a check over $\{\pm 1\}^3$ gives $\frac{5}{8}$. **Case** $K=4$: assume $z_1,\dots,z_4 \neq 0$, $z_5 = 0$. Writing m for the number of -1's among $(z_1,\dots,z_4)$, one gets $|\mathrm{Maj}_5(z)| = 0$ iff $m = 2$ (6 of 16 points) and $= 1$ otherwise; hence $\frac{5}{8}$. **Case** $K=5$: trivially the value is 1.

Expected value of $|f(z)|$. Set

$$M := \tfrac{1}{2}\big(z_1 + z_3 - z_4 + z_1z_3z_4\big).$$

Using the factorized form,

$$f(z) = \tfrac{1}{2}(-z_2+z_5) + \tfrac{1}{2}(1+z_2z_5)M.$$

Taking the conditional expectation over (z_2, z_5) yields

$$\mathbb{E}|f(z)| = \frac{1-2p+2p^2}{2}\,\mathbb{E}|M| + p - \frac{p^2}{2},$$

since the four cases are: both zero (prob. $(1-p)^2$), giving $\frac{1}{2}|M|$; exactly one zero (prob. $2p(1-p)$), giving $\frac{1}{4}(|M+1|+|M-1|) = \frac{1}{2}$ because $M \in [-1,1]$; both nonzero and equal (prob. $p^2/2$), giving $|M|$; and both nonzero and opposite (prob. $p^2/2$), giving 1.

It remains to compute $\mathbb{E}|M|$ by the number of nonzeros among (z_1,z_3,z_4): with probability $(1-p)^3$ there are none and $|M| = 0$; with probability $3p(1-p)^2$ there is exactly one, yielding $|M| = \frac{1}{2}$; with probability $3p^2(1-p)$ there are exactly two, again averaging to $\frac{1}{2}$; and with probability p^3 there are three, in which case M is the majority of three bits and $|M| = 1$. Therefore,

$$\mathbb{E}|M| = \tfrac{3}{2}p - \tfrac{3}{2}p^2 + p^3,$$

and hence

$$\mathbb{E}|f(z)| = \frac{7}{4}p - \frac{11}{4}p^2 + \frac{7}{2}p^3 - \frac{5}{2}p^4 + p^5.$$

At $p = 0.4$ this gives $\mathbb{E}|f(z)| = \frac{2689}{6250} = 0.43024$, whereas the earlier computation yields $\mathbb{E}\,|\mathrm{Maj}_5(z)| = \frac{5363}{12500} = 0.42904$.