

Boolean functions beyond the square-root degree bound: a mathematical and formal audit

A version-bounded audit of OpenAI/math Result 192

Corporate author: Haivian Technology

Commissioning owner: Richard Janutka, Qonduit Limited NZ

Audit date: 17 October 2026 Published: 2026-10-17

Version: HAIVIAN-BOOLEAN-DEGREE-VALIDATION-V1-20261017

Abstract. A Haivian audit supports unbounded violations of the square-root degree bound for Boolean functions. All 21 supplied Lean modules compiled; theorem and axiom inspection and the author-configured Comparator's default-kernel replay accepted the solution. The report states its exact scope, trust assumptions and AI-review limits.

What was audited

This audit examines OpenAI's Unbounded Violations of the Square-Root Degree Bound, dated 26 September 2026, and its supplied Lean proof, catalogued as Result192. The audited edition is repository commit `fd4aeeb2ee4fc729c18d98444fed42fd0529eeeb`: a 23-page paper and 21 original solution modules. The source was frozen before review and was not changed to obtain a successful build. [Frozen paper](#), [formalization scope](#).

Audit finding: the main result is supported by the mathematical review and reproduced formal verification for this frozen edition, within the trust and scope limits below. The main argument survived detailed reconstruction and a distinct AI review without a substantive gap or counterexample being found. All 21 original solution modules compiled, the actual printed theorem matched the paper, and seven inspected axiom closures contained only the permitted foundational axioms. The genuine author-configured Comparator and Lean default-kernel replay accepted the solution. Returned integrity and terminal-host evidence was reconciled with the execution record.

The claim, precisely

For every positive real number C , the theorem asserts that some positive finite dimension n and nonconstant Boolean function f on $\{-1, +1\}^n$ satisfy

$$\sum_i f(\{i\}) > C \sqrt{\deg(f)}$$

The Fourier coefficients use the uniform distribution on the cube. Degree means ordinary real multilinear, or equivalently real Fourier, degree. It does not mean degree over F_2 , approximate degree, threshold degree or decision-tree depth. The function need not be monotone or balanced. Both n and f may depend on C : the claim is not about one fixed function exceeding every factor.

The absolute version follows: the sum of the absolute linear Fourier coefficients divided by the square root of degree is unbounded over positive-degree Boolean functions. Positivity of degree is explicit, so no zero denominator is used.

This square-root claim must also be distinguished from the sharper majority benchmark, $E|X_1 + \dots + X_d|$. Earlier counterexamples to that stronger benchmark do not by themselves establish arbitrarily large violations of the square-root bound. The square-root conjecture and its conventions appear in the checked primary literature. [O'Donnell's problem collection](#), [Filmus and coauthors' open problems](#).

Why the argument works

The construction studies what can be learned about the sum of independent signs from a finite observation. Each possible observation has a cell indicator, whose Fourier degree gives a cost. The useful quantity is retained linear-score variance divided by that cost.

A rare event creates a gain through an exact cancellation. Let the central value select one leaf; let A mean all the unselected leaves pass their tests, but not all leaves pass. Its indicator is a sum of selected all-but-one products minus the all-pass product. Each term omits at least one input, despite the event as a whole referring to all of them. The subtraction is essential to the degree saving.

The paper's report reveals enough information to predict the total sum with a mean-square error strictly below one, while keeping the cell-degree cost controlled. Conditional expectation performs at least as well as this explicit predictor. Gaussian estimates establish a strictly positive gain for fixed, finite parameters. Finite central-limit approximations preserve that gain. Grouping and amplification then improve the variance-to-cost ratio by a fixed factor greater than one. For any requested C , only finitely many stages are needed. A final Boolean sign readout turns the retained score into the claimed large linear Fourier sum.

The audit specifically checked the order of these choices: Gaussian parameters and a positive margin are fixed first; each current finite observation then receives a sufficiently large finite grouping. No single central-limit estimate uniform over all future stages, or infinite-dimensional final witness, is required.

The formal route is sufficient, but different

The Lean proof reaches the same main theorem through softened Gaussian tests, deterministic table selection and switching. It is not a line-by-line formalization of every intermediate hard-interval construction in the paper.

Compact tent functions replace sharp leaf tests. Selector weights form a partition, with a residual assigned to a middle leaf. The noncompact residual moments are handled by exact full moments minus compact-test moments; they are not silently passed through an inapplicable bounded-test limit theorem.

Rounding assigns probabilities to entire selector and leaf tables. Averaging their deterministic event deficits reproduces the soft deficit, so one fixed table has positive gain. The chosen table

works for all inputs simultaneously; it is not freshly randomized at evaluation time, and adds no hidden random cube coordinates. Switching and finite amplification then preserve the necessary degree bounds. The final readout uses a positive universal constant, which suffices even though it is not the paper's particular numerical constant.

The actual printed formal target is `OAI.SquareRootDegree.main : SignedViolations  $\wedge$   $\neg$  BddAbove AbsoluteRatios`. We checked the expanded definitions against the paper, including the uniform real Fourier convention, finite positive dimension, nonconstancy and positive ratio denominator. [Original definitions](#), [main theorem](#).

Checks and formal evidence

The audit included visual review of all 23 paper pages, inspection of all 21 solution modules, primary-reference checks, independent derivations and exact finite calculations. A distinct reviewer reconstructed the argument and tested failure modes including a broken cancellation, incorrect whole-table weights, missing residual moments and changing rather than fixed amplification margins. No surviving mathematical objection was identified within this scope.

The small finite checks are not explicit counterexamples to the square-root conjecture. Their maximum retained variance minus cost was zero; a separate rational soft example had negative deficit. They test identities, degree accounting and error detection. Independent rational Gaussian bounds avoid floating-point underflow, but do not turn the enormous symbolic construction into a practical truth table.

The actual formal run used Lean 4.34.1 with pinned Mathlib revision `d13f23b723b8a846827a245b89c10fc7d3f11612`. All 21 original solution modules compiled freshly. Seven actual transitive axiom inspections—including the main theorem, Gaussian gain, table rounding, finite producer, amplification and readout—listed only `propext`, `Classical.choice` and `Quot.sound`. No additional assumption or `sorryAx` appeared in those closures.

The original author configuration selects the main theorem. The genuine Comparator built and exported the challenge and solution, ran Lean's default kernel on the solution and returned its final acceptance message with exit code zero. The challenge's intentional specification placeholder produced a warning; it is not the solution proof and does not occur in the inspected solution axiom closures. The thousands of dependency/build jobs reported by the tool are not thousands of separately audited candidate theorems: the fresh candidate closure contains 21 modules. [Unchanged author configuration](#).

The original configuration disables Nanoda, and no Nanoda check is claimed. This execution uses a retained compiler, compiled dependency cache and pinned comparison/export tools whose identities and isolation are checked. It is not a from-scratch rebuild of every dependency, a diverse-compiler experiment or an independently implemented second-kernel verification. Its trust boundary includes the Lean toolchain, libraries, exporter/comparison machinery and execution environment, as well as the three named foundational axioms.

The formal run completed in one attempt within its finite allocation. Checks preserved the original proof inputs and recorded outputs, retained tool/runtime identities and dependency contents outside normal generated cache metadata. The protected website and service baseline was unchanged at postflight, and no owned verification processes remained. The execution was not

operationally flawless: it recorded memory-high events without a hard-limit or out-of-memory kill, and a final watchdog cleanup command tried to stop an already unloaded successful unit. The returned journal distinguishes that cleanup error and a terminal sampling race from the successful proof commands. Neither anomaly was hidden or treated as mathematical evidence.

For version identification, the original paper's SHA-256 is 4b847d47e281e37df99e6c3c6561cfd1db55380e6504a73e30918ead56b74894; the original Comparator configuration's is d4e3e15f79b609ea5224bf8511d12e08344e9a5e684c392a2a2123601bcc35c9. The exact selected theorem and source revision above define the validation scope.

Limits and disclosure

Acceptance of the selected formal theorem does not automatically certify every auxiliary prose argument, historical assertion or appendix. This audit does not establish exhaustive novelty, a useful dimension-growth bound or a practical small violating function. Full primary texts for two contextual references—Wang and Kudin-Pašalić—were unavailable; their detailed historical assertions were not independently full-text-verified and are not premises of the main proof.

The mathematical owner and distinct reviewers used OpenAI-model agents from the same provider. Separate review contexts and independent calculations supply additional scrutiny, but not independent-provider corroboration or external human peer review. Haivian Technology authored the audit and operates the Validation Journal; this relationship is disclosed rather than presented as independent journal certification. The audit was commissioned by Richard Janutka, Qonduit Limited NZ. It is not an OpenAI endorsement.

The warranted conclusion concerns the exact frozen main theorem and the evidence actually obtained. It does not rest on the reputation of the paper's author, the number of agents involved or a count of automated build jobs.

Rights and disclosure

Copyright © 2026 Haivian Technology. Public article publication is authorized. Underlying private audit evidence, code, third-party manuscript assets and source bundles are not released by this publication.

This is same-provider AI review within an audit and journal operation controlled by Richard Janutka / Qonduit Limited NZ. It is not external human peer review, independent-provider verification, OpenAI endorsement or authorship of the original proof.

Canonical article: <https://haivian.com/validation-journal/boolean-functions-beyond-square-root-degree-bound-mathematical-formal-audit/>