

Nine dimensions, two formal targets: auditing Borsuk's covering counterexample

A version-bounded audit of OpenAI/math Result 156

Corporate author: Haivian Technology

Commissioning owner: Richard Janutka, Qonduit Limited NZ

Audit completed: 16 October 2026 UTC Published: 2026-10-17

Version: HAIVIAN-BORSUK-NINE-VALIDATION-V1-20261017

Abstract. A Haivian audit supports the compact nine-dimensional counterexample to Borsuk's covering assertion. All 269 supplied Lean modules compiled; six inspected axiom closures used only the declared permitted foundations, and both the unchanged author Comparator and a separate literal-R9 Comparator accepted their targets with Lean's default kernel. The report distinguishes the conventional full-cycle argument from the two formal endpoints and records the retained-toolchain and operational limits.

Our audit supports the compact nine-dimensional counterexample in OpenAI's preprint "A nine-dimensional counterexample to Borsuk's covering assertion." The mathematical reconstruction, independent finite checks and actual formal execution agree on the stated result. All 269 original Lean modules compiled, and two separately scoped genuine Comparator checks accepted their selected theorems using Lean's default kernel. This is a result about the specified construction and frozen edition, not a certification of every claim in the manuscript or a claim about the smallest possible counterexample dimension.

The examined scientific edition is repository commit `fd4aeeb2ee4fc729c18d98444fed42fd0529eeeb`. The manuscript is dated September 23, 2026. A separately inspected later repository comparison through commit `1581d4cbdd69a0faf34b2089614fe2a814316a6a` changed only repository automation, not this scientific edition. We did not silently substitute a moving version.

The construction and its dimension

Borsuk's covering assertion asks whether every bounded set of positive diameter in Euclidean n -dimensional space can be covered by $n+1$ sets of strictly smaller diameter. The preprint gives a compact witness in nine dimensions that admits no such ten-set cover. The covering subsets are arbitrary; no measurability assumption is imposed.

The witness consists of matrices uu^T for unit vectors u in R^4 . Opposite vectors give the same matrix, making it a model of real projective three-space. Its intrinsic dimension is three, while its ambient affine space has dimension nine: symmetric 4-by-4 matrices have ten independent entries, and trace one removes one degree of freedom. The witness is the whole compact set, not an explicitly supplied finite point cloud.

Our exact algebra check recovered

$$\|uu^T - vv^T\|_F^2 = \|u\|^4 + \|v\|^4 - 2(u \cdot v)^2$$

For unit vectors the diameter is $\sqrt{2}$, attained by perpendicular vectors. A separate rational-polynomial calculation checked the nine-coordinate isometry on differences of trace-one symmetric matrices. These checks establish the geometric dimensions and distances used in the argument; they do not alone rule out a ten-part cover.

The difficult reduction and finite obstruction

The proof converts a hypothetical strict smaller-diameter cover into smooth support data on projective space, then derives constraints on a finite support complex. We checked the use of strict diameter gaps, the positive integral extension, its exact support, and the signed spectral extension. The topological hinge was reconstructed through compact regular fibers, bundle trivialization, degree parity and facet incidence. A distinct same-provider review context independently reconstructed these steps and found no concrete mathematical gap.

Independent executable checks tested the finite step. One exact checker verified 31,200 base clauses over 330 variables and replayed 4,093 supplied certificate statements, including 742 reverse-unit-propagation deductions, ending in an empty clause. Its contradiction requires both the TenSystem conditions and the existence of a tetrahedron. The all-false assignment satisfies the base conditions without that existence hypothesis; this control prevents an unjustified stronger claim.

A distinct enumeration examined all 184,756 choices of ten out of twenty triples and recovered exactly twelve qualifying six-label systems. It used a different search parametrization from the first checker. Negative controls rejected unsupported contradictions and a removed final proof body. These are substantive reproductions, not a vote among reviewers or a replacement for the topology.

What was actually verified formally

The audit used Lean 4.34.1 and the pinned Mathlib revision d13f23b723b8a846827a245b89c10fc7d3f11612. All 269 original solution modules, including 145 finite-certificate modules, completed successfully. The final diagnostic compiled an explicit statement in EuclideanSpace R (Fin9).

Actual axiom output was obtained for six declarations: `main_theorem`, `euclidean_nine_counterexample`, `no_admissible_ten`, `FiniteObstruction.no_ten_system`, `Admissible.maximal_support_complement_bound`, and `Admissible.no_free_ridge`, all under `OAI.BorsukNine`. Each printed closure contains exactly `propext`, `Classical.choice` and `Quot.sound`, the three permitted Lean axioms. This is an observed dependency result, not merely a search of source text for admissions.

The two final comparisons have different scopes:

1. The unchanged author configuration selects `OAI.BorsukNine.main_theorem`: the projector set is compact, lies in the trace-one symmetric slice, has diameter $\sqrt{2}$, and has no ten-set cover of strictly smaller diameter.

2. A separately labeled audit-authored challenge selects the original solution theorem `OAI.BorsukNine.euclidean_nine_counterexample`. It states the existence of such a compact set in literal EuclideanSpace R (Fin9).

Both genuine Comparator invocations built their actual targets, exported the selected declarations, ran Lean's default kernel and returned successful final acceptance. Their recorded command durations were 287.073 and 282.419 seconds. The second was not silently added to or substituted for the author's statement. Intentional holes in the two challenge templates are specifications, not solution proofs. Both configurations disable Nanoda: no second kernel implementation is claimed to have run. Lean's default printer elides some inferred types, so the explicit diagnostic and separate exact-R9 comparison matter to this conclusion.

The run reused the retained, identity-checked compiler and compiled dependency cache, together with the pinned retained Comparator and exporter. It did not rebuild all of Mathlib and its dependencies from source or perform diverse-compiler verification. The trust boundary includes Lean, its toolchain and libraries, the exporter, Comparator and execution environment, as well as the three listed axioms. The separate review contexts do not constitute an independent second kernel.

The outer verification then completed its source, compiled-output, tool, runtime and dependency-integrity checks. We rehashed 994 returned evidence files and reconciled all 269 terminal object digests with the build records. Compiled object bytes remain in the controlled execution environment; they were not downloaded as part of the returned log bundle. These distinctions are retained in the audit records rather than replaced by a blanket claim of public reproducibility.

Limits and operational observations

The Lean route uses purity, local parity and no-free-ridge consequences to reach the finite obstruction. It is not a literal formalization of every sentence of the manuscript's full facet-cycle argument. We support that argument through conventional mathematical review. The higher-dimensional corollary is also a conventional deduction here, not a separately executed formal target. We make no claim about an exact Borsuk number or minimal counterexample dimension.

The bounded run completed in 7,078.800 seconds, 21.200 seconds before its original 7,100-second scientific cutoff, without resetting the clock. No out-of-memory kill or active guard failure was observed. It was not operationally flawless: there were 7,906 MemoryHigh events, transient pressure, a largest valid fast-sample gap of 5.794 seconds, and one terminal telemetry race. Of 211 slow records, 210 contained successful health checks and one recorded that terminal race. After scientific completion, a cleanup command tried to stop an already unloaded unit and returned an error; the watchdog's main process itself exited successfully. Protected site/service state remained unchanged, no owned job remained, and recorded production mutations were zero. These caveats do not erase the successful proof checks, but they rule out describing execution as perfectly uninterrupted monitoring or flawless infrastructure.

Review, ownership and evidence access

The audit combines direct mathematical reconstruction, independent executable checks and separately commissioned same-provider AI review contexts. It is not external human peer review, certification by an independent publishing organization, or an OpenAI endorsement. Haivian Technology and the Haivian Validation Journal share organizational control under Richard Janutka / Qonduit Limited, New Zealand.

The linked manuscript and formalization remain their authors' work. This report is our original account, not a republication or relicensing of those materials. The accompanying identity manifest identifies the source edition and retained audit records; it does not claim that the complete internal logs, host records or compiled artifacts are publicly downloadable. No private operational paths are needed to understand the reported scientific scope.

Primary sources

[Frozen manuscript](#)

[Original author Comparator configuration](#)

[Original explicit nine-dimensional theorem](#)

[Later repository comparison](#)

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/nine-dimensions-two-formal-targets-auditing-borsuk-covering-counterexample/>