

Seymour's second-neighborhood conjecture: mathematical audit and completed formal verification

A version-bounded audit of OpenAI/math Result 173

Corporate author: Haivian Technology

Commissioning owner: Richard Janutka, Qonduit Limited NZ

Published: 2026-10-16 **Version:** HAIVIAN-SEYMOUR-VALIDATION-V1-20261016

Abstract. A Haivian audit supports the specified finite oriented-graph theorem. All 43 supplied Lean modules compiled, permitted axiom dependencies were inspected, and the author's Comparator accepted default-kernel replay. The report states the exact edition, trust assumptions and same-provider AI-review limits.

Result

Our direct audit supports the main proof in OpenAI's *A proof of Seymour's second-neighborhood conjecture*. The formal-verification follow-up has now also completed: all 43 supplied proof modules compiled under the pinned Lean environment, the actual theorem and axiom dependencies were inspected, and the supplied Comparator accepted the solution after Lean default-kernel replay. This conclusion concerns one precisely identified manuscript edition and its unweighted theorem, under the trust assumptions stated below—not every possible formulation or an independent second-kernel certification.

The reviewed edition is the 15-page September 23, 2026 manuscript at OpenAI/math commit `fd4aeeb2ee4fc729c18d98444fed42fd0529eeeb`, Result 173. Its PDF SHA256 is `20be859799f83e3412dba6fd645a72a0d929cbe05ca546807c591f001f9e6e82`. The authors' [pinned manuscript](#) and [formal scope note](#) identify the original work. We have not changed the reviewed edition or claimed authorship of its proof.

What was proved and checked

In a nonempty finite directed graph with no loops and no pair of opposite arcs, the conjecture asserts that some vertex has at least as many second out-neighbors as first out-neighbors. A second out-neighbor must be reachable in two steps but must not be the starting vertex or an immediate out-neighbor. Multiple paths to the same endpoint count once. A sink qualifies automatically. Connectivity and completeness are not required.

These conventions are reflected in the actual printed Lean definitions. The compiled theorem quantifies over an arbitrary finite nonempty vertex type and arbitrary relation satisfying the orientation conditions. Its conclusion compares finite-set cardinalities. The decidable-equality instance is a classical representation device, not an additional graph assumption. This matters: verifying a nearby statement with overlapping neighborhoods or stronger hypotheses would not verify the conjecture under review. The exact [solution theorem](#) is `OAI.SeymourSecondNeighborhood.exists_goodVertex`.

Mathematical investigation

The initial paper audit inspected all 15 rendered pages and used complementary, procedurally separated AI investigators before direct challenge and reconciliation. It examined the minimal-counterexample and arc-deletion reduction, the transitive-tournament blowup, the arbitrary-relation pruning lemma, the maximum-rank and matching argument, and the final strict extremal inequality.

Several distinctions were decisive. Arc deletion can both create and remove exact second neighbors, so the accounting must include both effects. The blowup's full-support case must be separated from proper support. The rank construction needs one coefficient assignment satisfying all required conditions, while its perturbation contradiction is local to a single matrix. Finally, the pruning argument charges against different uncovered sets before and after pruning; interchanging them would invalidate the strict improvement.

The investigators found no concrete proof gap or counterexample in that claim chain. Small exhaustive graph and relation tests, supplemented by a separately implemented pruning optimizer, supported checks of finite interfaces and boundary cases. Those tests do not prove an unrestricted theorem, and overlapping test sets were not counted as independent mathematical evidence. The conclusion rests on the reconstructed arguments and, now, the actual formal check—not on model agreement or a test total.

The earlier conventional audit also supported the manuscript's weighted and directed-cycle consequences. Those consequences are not separate targets of the supplied Comparator configuration and should not be described as newly formally verified by this execution.

Formal execution and its limits

The follow-up used Lean 4.34.1, mathlib commit d13f23b723b8a846827a245b89c10fc7d3f11612 and the recorded locked dependency set. Every one of the 43 OAI modules was freshly compiled. The actual axiom printouts for the final theorem, pruning theorem, counterexample reduction and extremal contradiction each contained exactly `propext`, `Classical.choice` and `Quot.sound`—the three axioms permitted by the [author's unchanged configuration](#).

The real Comparator separately built and exported the challenge and solution, compared the configured declaration, checked axiom compliance and replayed the solution with Lean's default kernel. It returned zero and reported final acceptance. The challenge's expected `sorry` is a specification placeholder; it was not accepted as the solution's proof. The solution's separate axiom closure contains no `sorryAx`.

Nanoda was disabled in the supplied configuration and did not run. Default-kernel replay is not an independent second-kernel result. The procedure also depends on the identified compiler/checker binaries, operating system and retained external dependency cache. We did not clean-bootstrap the whole compiler and library stack or manually audit every transitive external source file. Source identity and manuscript correspondence were reviewed separately from kernel acceptance.

Earlier execution attempts encountered read-only cache metadata, missing empty Git directories and a pressure-guard stop before proof compilation. These were recorded as technical failures, not mathematical refutations. Reviewed, source-preserving integration corrections allowed the completed run without editing the proof, manufacturing cache metadata, weakening safety controls or silently resetting its time allocation. Historical reports remain intact rather than being retroactively credited with later success.

The completed run's terminal evidence records unchanged sources and compiled outputs after the Comparator, a successful protected-host postflight and no remaining owned processes. A watchdog cleanup warning after the successful unit had unloaded, and a terminal monitoring read race, are disclosed in the technical addendum; they are not concealed as a flawless infrastructure run. The final result is supported by actual command outputs and linked evidence, not just a success label.

Evidence identifiers and review boundary

The actual theorem stdout is identified by SHA256

4294f400af5b0a04b28f2112673ee0ae34302a9f9a506b8aeec4f85f098057fe. The completed formal receipt is d682bd464c2f1a9e76eda6f0b9ccc31a27da059ea030d763479954794bf6fdce; the actual Comparator receipt is 49ed69bf8c9f3c6c1c30da3d25ff76137e29897a80f907c91a14343a6574c3b. These identifiers bind the retained evidence; they are not themselves publicly downloadable evidence or substitutes for inspection.

Local reconciliation checked all 365 returned files against their recorded identities and matched the 43 build records to terminal object-hash records. A distinct final scientific/editorial review assesses the actual evidence and this article separately from the implementation work. The audit does not purport to settle historical priority, establish general external consensus, or certify every introductory literature claim.

Commissioning and disclosure

Richard Janutka of Qonduit Limited, New Zealand, commissioned this audit and controls the audit/journal operation. Haivian Technology is the user-directed corporate author brand. OpenAI is both the reviewed manuscript's author and the provider of the Astra models used in the review. OpenAI is not represented as sharing ownership of the commissioning operation, and this article is not an OpenAI endorsement.

The investigators and reviewers were configured as Astra with high reasoning in procedurally separated contexts. This is same-provider AI review, not external human or institutional peer review, independent-provider corroboration, or independently attested model-runtime identity. The evidence and stated limits—not the number of agreeing agents—support the conclusion. The original manuscript remains its authors' work and is cited rather than republished or relicensed here.

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

Canonical article:

<https://haivian.com/validation-journal/seymours-second-neighborhood-conjecture-mathematical-audit-formal-verification/>