

Shoreline search audit supports the repaired hybrid proof

A version-bounded audit of the logarithmic-spiral shoreline-search proof

SUPPORTED WITHIN STATED HYBRID SCOPE · FOUR AUDIT-AUTHORED REPAIRS

Published 2026-10-21 · Audit completed 21 October 2026 ·
HAIVIAN-SHORELINE-SEARCH-VALIDATION-V1-20261021

Haivian Technology · Commissioned by Richard Janutka, Qonduit Limited, New Zealand

Scope and disclosure

This report supports the repaired shoreline-search argument within its stated hybrid scope. It does not claim that the full theorem was formalized end to end in Lean, that C1-C4 are an author-issued corrigendum, or that the paper's author or AI provider endorsed or commissioned the report.

The audit used the official Lean compiler and cached compiled Mathlib dependencies; no full dependency rebuild from source was performed. Haivian's audit service and the Haivian Validation Journal share organizational control. Same-provider AI review had prior source exposure and was not external human accreditation.

The audit of Alexander Temerev's The logarithmic spiral is optimal for shoreline search: a computer-assisted proof supports the argument **within its stated hybrid scope, with four explicit repairs**. Numerical verification and Lean checks of supporting lemmas are complete. The full argument also depends on conventional mathematical reasoning; the main theorem has not been formalized end to end in Lean.

The audited edition is [arXiv:2609.24454v1](https://arxiv.org/abs/2609.24454v1). The four repairs described below are audit-authored corrections to that edition, rather than an author-issued corrigendum.

What the result concerns

The mathematical problem asks how a searcher starting at the origin should move to reach an unknown infinite straight line. The model allows deterministic planar paths whose speed is at most one, and compares hitting time with the line's distance from the origin. The claimed optimal ratio is approximately **13.811135179461**.

The audit supports the repaired argument in this idealized model. Paths may wind arbitrarily close to the origin. The conclusion does not establish optimality for physical vehicles with a minimum turning radius or other additional movement constraints.

What the numerical checks established

Two author-written verifier implementations accepted the same frozen projected certificate. They use different algorithms for polynomial construction, differentiation and subdivision, while sharing the Python-FLINT/Arb arithmetic backend. This provides algorithmic cross-checking with a shared arithmetic dependency.

Numerical section	First verifier boxes	Second verifier boxes
Surface	365,466	259,560
Interior	816,750	550,362
Small dives	650	4,284
Dives	15,224	11,154
Total	1,198,090	825,360

Both implementations reported all four sections verified, with empty failure lists and strictly positive minimum bounds. Their box counts differ because their subdivision methods, exclusion neighborhoods and small-dive cutoffs differ. Both used 128-bit precision for the covers. The second implementation retained the author's 400-bit internal setup for constants and correction construction.

The successful first run from 17 October was recovered and reused. Its original command, successful termination, full output and matching result file were located. That historical provenance was reconstructed; it cannot establish every historical environment variable or loaded library.

The second implementation completed locally on 21 October. Its counts and structural results reproduced the deposited second-verifier result. All four minimum bounds agreed within a relative tolerance of 10^{-11} and remained positive. An intentionally corrupted copy of the certificate was rejected. The accepted certificate and author source were preserved.

What the Lean checks established

All eight submitted helper modules compiled and passed explicit kernel replay using Lean 4.33.0 and the pinned dependencies. Statements and transitive axiom reports were captured for 41 public theorems and the `IsBestSource` predicate definition. Their combined axiom set was exactly `propext`, `Classical.choice` and `Quot.sound`.

These helpers establish supporting lemmas. Geometric constructions, maximizer attainment, measure and bounded-variation arguments, and the final theorem still rely on the conventional proof reviewed in the audit. Replay used the same Lean kernel against imported dependency environments; it did not replay the entire dependency chain from an empty environment or provide corroboration from a second kernel.

The audit used the official Lean compiler and cached compiled Mathlib dependencies; no full dependency rebuild from source was performed.

Four repairs required by the audit

- 1 C1 — Correct the dive comparison state.** The actual terminal dive state cannot be identified with the saturated comparison state. The repaired argument defines the quantity Q using the no-extra-finalization comparison endpoints. Verified monotonicity and dive inequalities then support the comparison, including the reduction in the actual endpoint state during finalization.
- 2 C2 — Include the full past in logarithmic time.** The dense-source restriction $[0, \tau)$ must become $(-\infty, \tau)$, retaining the current-source term. An exact radial counterexample demonstrates why negative logarithmic times matter. This correction is necessary for the source formula and its measurability argument.
- 3 C3 — Include the endpoint maximizer.** When $\beta < \alpha$, the maximum occurs at the endpoint and the lifted length is 2β . This omitted branch remains below 2π , so it leaves the competitiveness condition unchanged. The certified root lies in the interior branch $\beta > \alpha$.
- 4 C4 — Supply the continuously differentiable extension through zero.** Smoothness in β alone does not justify substitution by $\beta = \text{acos}(\exp(-v))$ at $v = 0$. The repair writes the first-cell spline as $G0 + B\beta^2/2 + C\beta^3/6$, extends it to negative v by $G0 + Bv$, and extends the fan term by zero. The replay supplies the coefficient identities, simple-knot checks and correction-support checks required by the saved analytic argument.

The audit found no remaining discrepancy that breaks the repaired argument within the reviewed scope. These repairs remain part of the conclusion: numerical acceptance does not make the deposited exposition flawless.

How to interpret the verification

At the spiral, an interval containing zero is a consistency check. Exact vanishing of the relevant derivatives depends on the symbolic contact and pitch argument and the implicit real correction in the conventional review. The numerical replay supplies the associated root isolation, sensitivity, Hessian and domain-cover obligations.

The separate conventional review used AI from the same provider, with prior source exposure disclosed. Final reconciliation reused that review and checked it against the completed execution evidence. This is not a new blind review or external human accreditation. The report makes no claim of endorsement or commissioning by the paper's author or the AI provider. Richard Janutka of Qonduit Limited, New Zealand, commissioned this audit. Haivian's audit service and the Haivian Validation Journal share organizational control.

All 20 checks in the final evidence reconciliation passed. The resulting disposition is **reproduced within the stated hybrid scope, with C1–C4 repairs**.

Source identity

Audited edition: [arXiv:2609.24454v1](https://arxiv.org/abs/2609.24454v1). Accepted article source SHA-256: d7afe60d25ea6a35ac590f002b570e30c71508760a60d9f76348b4d5063a4b22. Completed report SHA-256: 537761aee36fbbcb3fcfc4d63e9cfd334968f43382b0beaa28aeea2dc460dbaa. Final evidence SHA-256: 07a9c9d50f4faab3ab25ef5a49794a49a5ddcb9f5b44e4a897a75b33925dba7f.