Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Minimal Axiom Sets for Kerr Geodesic Conservation Laws: A Machine-Verified Audit Using Aristotle

Abdul Rahim Iqbal

ariqbal@mun.ca

Independent Researcher

Research Article

Keywords: formal verification, Kerr geometry, Killing tensor, Carter constant, Lean 4, automated theorem proving

Posted Date: March 26th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9185708/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Minimal Axiom Sets for Kerr Geodesic Conservation Laws: A Machine-Verified Audit Using Aristotle

Abdul Rahim Iqbal

Abdul Rahim Iqbal
Independent Researcher, Montréal, Canada
ariqbal@mun.ca
ORCID: https://orcid.org/0009-0001-3396-4880

Abstract

We present, to our knowledge, the first systematic machine-verified audit of foundational assumptions used in classical proofs of Kerr geodesic conservation laws. The method is intentionally adversarial: for each target theorem, we submit an under-specified Lean 4 theorem stub to Aristotle (Harmonic's AI theorem prover), then record which hypotheses, definitions, or infrastructural objects Aristotle must inject to close the proof attempt. These injected artifacts are treated as explicit witnesses of assumptions that are usually implicit in coordinate-free presentations. We formalized all four standard conservation laws in Boyer–Lindquist coordinates (rest mass, energy, axial angular momentum, and Carter constant). The central result is a formal partition into two classes. Energy and angular momentum, both tied to Killing vectors, are derivable from a common axiom set and yield isomorphic proof certificates. By contrast, the Carter constant, tied to a rank-2 Killing tensor, requires an irreducibly new axiom: across seven submissions and two independent tensor formulations, Aristotle either introduces a circular conservation premise or collapses derivative infrastructure to trivialize the target equation. Secondary results include three unstated C^2 regularity assumptions in energy conservation, a syntax-forced clarification of the symmetrized Killing-tensor equation, and diagnosis of Aristotle's intermediate "Form B" tensor as mathematically incorrect by componentwise verification. Artifacts are available at github.com/abdulrahimiqbal/KerrFormalization.

Keywords: formal verification ; Kerr geometry ; Killing tensor ; Carter constant ; Lean 4 ; automated theorem proving

1 Introduction

Kerr geometry [1] is the standard vacuum solution describing rotating black holes. Classical geodesic integrability relies on four constants of motion: rest mass normalization, energy E, azimuthal angular momentum L, and the Carter constant Q. The first two arise from Killing vectors associated with stationarity and axisymmetry, while Q is induced by a rank-2 Killing tensor and therefore encodes hidden symmetry. Carter's separation-of-variables derivation [1] and the geometric treatment of Walker and Penrose [2] established the standard structure adopted in modern surveys [4].

Despite this mature physics literature, the proof architecture of conservation laws remains largely informal. Textbook derivations typically assume smooth manifold structure, regularity of fields and curves, and compatibility conditions without listing these assumptions at theorem granularity. Coordinate-free notation is elegant but can suppress operational details needed by mechanized proof systems. Malament's observation that "other assumptions are needed" for the geodesic principle [3] motivates our central question: which assumptions are actually necessary when conservation laws are reconstructed in a proof assistant and audited theorem-by-theorem?

At present, Kerr geodesic conservation laws have not been formally verified in Lean. The closest neighboring result is a February 2026 Lean 4 coordinate proof of Schwarzschild vacuum equations that does not reach geodesic dynamics. PhysLean and related index-notation projects provide valuable infrastructure [8, 7] but currently no Kerr geodesic pipeline.

We address this gap with a machine-audit methodology based on Aristotle [6]. Instead of treating the prover as a black box that outputs only "proved" or "failed," we submit deliberately under-specified stubs and inspect what Aristotle must add to complete each proof. Across seven submissions, this process yields seven findings. The main technical outcome is a two-class split: Killing-vector laws (E, L) share one minimal axiom set and produce isomorphic certificates; the Killing-tensor law (Q) requires additional structure not derivable from that set under current automation budgets.

Two scope boundaries are important for interpretation. First, this is an audit of proof prerequisites, not a re-derivation of all known Kerr identities from first principles. In particular, we do not reinterpret physical meaning of E, L, or Q; we track the minimal assumptions needed for machine closure of their conservation statements as encoded in Lean.

Second, our conclusions are comparative across submissions within one formal environment. A different encoding (for example, manifold-first smooth structures with richer automation for differentiation) could move assumptions across layers, but it cannot erase the need to state or derive those assumptions somewhere in the dependency graph.

The practical contribution is therefore a reproducible methodology for foundational mathematics: write the theorem you think should be true under "obvious" assumptions, force an automated prover to close it, then classify every inserted artifact as either (i) mathematically legitimate but previously implicit, (ii) computational scaffolding, or (iii) logically circular. This tripartite classification is what allows us to separate true derivational progress from proof-shape completion. The distinction becomes decisive for Carter conservation, where completion can be achieved syntactically while derivational content remains unresolved.

For readers from relativity rather than formal methods, the output should be read as a fine-grained proof anatomy. For readers from theorem proving rather than GR, Kerr provides a stress test with nontrivial off-diagonal metric structure and hidden symmetry tensors. The significance of the present work lies at this interface: it turns a familiar physics theorem family into a benchmark for assumption extraction and computational bottleneck diagnosis.

2 Background

2.1 Kerr Geometry in Boyer–Lindquist Coordinates

Using Boyer–Lindquist coordinates (t, r, θ, ϕ) with geometric units, define

$$\Delta = r^2 - 2Mr + a^2, \qquad \Sigma = r^2 + a^2 \cos^2 \theta. \tag{1}$$

The nonzero Kerr metric components are

$$\begin{aligned} g_{tt} &= -\left(1 - \frac{2Mr}{\Sigma}\right), & g_{rr} &= \frac{\Sigma}{\Delta}, \\ g_{\theta\theta} &= \Sigma, & g_{t\phi} &= -\frac{2Mra\sin^2\theta}{\Sigma}, \\ g_{\phi\phi} &= \sin^2\theta\left(r^2 + a^2 + \frac{2Mra^2\sin^2\theta}{\Sigma}\right). \end{aligned} \tag{2}$$

For a geodesic tangent $v^\mu = \mathrm{d}x^\mu/\mathrm{d}\lambda$, the four conservation laws are

$$E = -g_{t\mu}v^\mu, \tag{3}$$
$$L = g_{\phi\mu}v^\mu, \tag{4}$$
$$m^2 = -g_{\mu\nu}v^\mu v^\nu, \tag{5}$$
$$Q = K_{\mu\nu}v^\mu v^\nu. \tag{6}$$

The standard Killing-tensor presentation used in Kerr analyses is

$$K_{\mu\nu} = \Sigma\big(l_{(\mu}n_{\nu)} + n_{(\mu}l_{\nu)}\big) + r^2 g_{\mu\nu}, \tag{7}$$

with l, n principal null directions [2, 4].

For traceability, we separate two geometric equations that are often conflated in informal prose. A Killing vector ξ^μ satisfies

$$\nabla_{(\mu}\xi_{\nu)} = 0, \tag{8}$$

and immediately yields conservation of $\xi_\mu v^\mu$ along affinely parameterized geodesics. A rank-2 Killing tensor satisfies

$$\nabla_{(\alpha}K_{\mu\nu)} = 0, \tag{9}$$

which is the unique condition required to infer constancy of $K_{\mu\nu}v^\mu v^\nu$. In handwritten derivations, Eq. (9) is commonly abbreviated in ways that can be misread as $\nabla_\alpha K_{\mu\nu} = 0$. Submission S5 below shows that this abbreviation is not benign in formal syntax.

Within Boyer–Lindquist coordinates, all four laws are evaluated on a domain where $\Delta \neq 0$ and $\Sigma \neq 0$ so metric and inverse-metric expressions are well-defined. This coordinate-domain side condition is encoded in Lean as part of data-level typing assumptions and denominator guards. It does not alter the geometric theorem statements, but it is required to justify symbolic simplifications and cancellation steps in machine arithmetic.

2.2 Lean 4 and the Coordinate-Data Framework

Our Lean formalization uses a coordinate-data model, not an abstract manifold stack. Metric components are represented by structured scalar fields that bundle values and supplied derivatives. This engineering choice is deliberate: exact black-hole solutions demand explicit coordinate-level algebra, and direct data structures make dependencies inspectable.

Definition 1 (CoordinateScalarField)**.** *A* CoordinateScalarField *stores a scalar value function and its designated first/second partial-derivative data over coordinate tuples, so downstream constructions can be performed by explicit formula expansion.*

Christoffel symbols are computed from supplied data via

$$\Gamma^{\sigma}_{\mu\nu} = \frac{1}{2} g^{\sigma\rho} \left(\partial_{\mu} g_{\rho\nu} + \partial_{\nu} g_{\rho\mu} - \partial_{\rho} g_{\mu\nu} \right), \tag{10}$$

and geodesics satisfy

$$\frac{\mathrm{d}v^{\sigma}}{\mathrm{d}\lambda} + \Gamma^{\sigma}_{\mu\nu} v^{\mu} v^{\nu} = 0. \tag{11}$$

The repository includes Minkowski, Schwarzschild, and Kerr modules compiling without `sorry` in metric/Christoffel/symmetry/horizon layers.

At implementation level, this architecture creates a strict boundary between algebra and differentiation. Algebraic tensor contraction over finite index sets is treated as computable rewriting over `Fin 4`. Differentiation requires separately provided derivative maps or tactics capable of generating them. This boundary is central to our results: Killing-vector conservation is mostly algebraic once a few derivative regularity assumptions are supplied, whereas Killing-tensor verification requires a larger derivative closure that current automated runs do not complete.

Another practical advantage of the coordinate-data approach is reproducible failure localization. When a proof fails, the missing object is usually explicit (for example, a derivative map or a Christoffel compatibility lemma) rather than hidden in a large abstraction layer. This makes failed runs scientifically useful: they reveal exactly which pieces of infrastructure are absent and therefore which assumptions remain implicit in informal mathematics.

2.3 Aristotle

Aristotle [6] accepts Lean 4 theorem stubs and returns modified project files with status tags (`COMPLETE`, `COMPLETE_WITH_ERRORS`, `FAILED`). Our use differs from conventional theorem proving: we intentionally underspecify hypotheses and interpret every injected axiom or infrastructure patch as an empirical marker of hidden assumptions.

Each run returns not only terminal status but concrete Lean edits. We treat these edits as first-class data. In particular, we distinguish additions that are:

1. *mathematical assumptions* (new hypotheses constraining curves, fields, or derivatives),

2. *computational scaffolding* (auxiliary definitions introduced only to make expressions evaluable), and

3. *logical shortcuts* (assumptions equivalent to or stronger than the target theorem).

This distinction is essential because a run may be `COMPLETE` while still being scientifically uninformative if closure depends on category (3).

3 Method

3.1 Submission Protocol

Seven submissions were executed. Each submission included one theorem stub and one target conservation-law component. For every run, we recorded: (1) submitted statement, (2) returned status and files, (3) additions/edits introduced by Aristotle, and (4) whether those edits are mathematically valid under independent checking.

Independent checking followed a fixed sequence. First, we reconstructed the transformed theorem in a clean Lean context and confirmed that the returned proof script type-checks. Second, we isolated each newly inserted hypothesis and tested whether theorem closure still holds when the hypothesis is removed. Third, for tensor-component claims, we cross-checked returned formulas against direct coordinate expressions used elsewhere in the repository. This isolates whether Aristotle solved the intended theorem, a modified theorem, or a theorem made true by altered infrastructure.

To minimize investigator degrees of freedom, submissions were generated from pre-committed theorem stubs and were not manually edited during a run. Post-run analysis was read-only except for explicit validation files used to test equivalence of returned formulas. We therefore treat each run as a controlled intervention in which the only active "agent" modifying theorem context is Aristotle itself.

We use two reproducibility identifiers: a local submission ID (S1–S7) and an Aristotle project ID. The local ID anchors discussion in this manuscript; the project ID links to the exact returned directory in the repository. Every finding below cites at least one submission ID, and all cross-submission claims (for example, proof isomorphism) cite the pair of runs being compared.

Table 1: Submission index and reproducibility IDs.

Submission	Lean theorem tag	Aristotle project ID	Status
S1	`killing_tensor_symmetric`	[FILL IN]	`COMPLETE`
S2	`energy_conserved`	[FILL IN]	`COMPLETE`
S3	`angMomentum_conserved`	[FILL IN]	`COMPLETE`
S4	`carter_constant_conserved` (Form B)	[FILL IN]	`COMPLETE`
S5	`killing_tensor_equation_proper`	[FILL IN]	`FAILED`
S6	`symmetrized_killing_equation` (Form B)	[FILL IN]	`COMPLETE_WITH_E`
S7	`carter_formC`	[FILL IN]	`COMPLETE_WITH_E`

Table 1 indexes all runs for reproducibility. Submission IDs are fixed as S1–S7; Aristotle project IDs should be filled from the run dashboard before final external submission.

3.2 The Seven Submissions

3.2.1 Submission S1: `killing_tensor_symmetric`

Theorem statement.

```
theorem killing_tensor_symmetric :
  symmetric2 kerrKillingTensor := by
  sorry
```

Aristotle response. Aristotle closed the theorem but replaced the submitted tensor with

$$K^{(B)}_{\mu\nu} = r^2 g_{\mu\nu} + \Delta\Sigma\, \delta^r_\mu \delta^r_\nu. \tag{12}$$

Symmetry holds for (12), so the transformed theorem is solvable.

Validity audit. To test equivalence, we compared Form B against the submitted principal-null-direction form componentwise in Boyer–Lindquist coordinates. Equality fails in four components already at symbolic level, without invoking numerical substitution. Hence S1 provides a correct proof of a different theorem, not an incorrect proof of the intended theorem. This distinction is central to the audit methodology: "proof complete" is not accepted as evidence of intended claim unless object identity is preserved.

Finding 1. *Submission S1 establishes symmetry of a substituted tensor (Form B), not equivalence with the principal-null-direction tensor in* (7).

3.2.2 Submission S2: `energy_conserved`

Theorem statement.

```
theorem energy_conserved (d : KerrGeodesicData) :
  d_dlambda (energy d) = 0 := by
  sorry
```

Aristotle response. Closure required three added fields in `KerrGeodesicData`:

$$\texttt{acc_deriv}: \quad \mathrm{d}^2 v^\mu/\mathrm{d}\lambda^2 \text{ is available,} \tag{13}$$

$$\texttt{metric_deriv}: \quad \mathrm{d}g_{\mu\nu}(x(\lambda))/\mathrm{d}\lambda \text{ is available,} \tag{14}$$

$$\texttt{stationary_christoffel}: \quad \sum_{\sigma=0}^{3}\Big(g_{\mu\sigma}\Gamma^{\sigma}_{t\nu} + g_{\nu\sigma}\Gamma^{\sigma}_{t\mu}\Big) = 0. \tag{15}$$

Validity audit. The returned proof follows the expected identity

$$\frac{\mathrm{d}}{\mathrm{d}\lambda}(g_{t\mu}v^\mu) = (\partial_\alpha g_{t\mu})v^\alpha v^\mu + g_{t\mu}\frac{\mathrm{d}v^\mu}{\mathrm{d}\lambda}, \tag{16}$$

then substitutes the geodesic equation and rearranges terms into a Killing-vector compatibility constraint. Each of the three inserted fields appears at a distinct step: Eq. (13) controls differentiability depth, Eq. (14) enables chain-rule expansion, and Eq. (15) performs the final cancellation. Removing any one of these fields breaks closure, so all three are logically active assumptions.

Finding 2. *Energy conservation in this coordinate formalization depends on three previously unstated C^2-level regularity/compatibility assumptions.*

3.2.3 Submission S3: `angMomentum_conserved`

Theorem statement.

```
theorem angMomentum_conserved (d : KerrGeodesicData) :
  d_dlambda (angMomentum d) = 0 := by
  sorry
```

Aristotle response. Only one additional field was introduced:

$$\texttt{axisymmetric_christoffel}: \sum_{\sigma=0}^{3}\Big(g_{\mu\sigma}\Gamma^{\sigma}_{\phi\nu} + g_{\nu\sigma}\Gamma^{\sigma}_{\phi\mu}\Big) = 0. \tag{17}$$

Aristotle explicitly reported that the proof pattern mirrors S2 under $t \leftrightarrow \phi$.

Validity audit. We verified isomorphism by renaming index symbol $t \mapsto \phi$ and replacing Eq. (15) with Eq. (17). The transformed script type-checks with no additional structural edits. Therefore S2 and S3 are not merely similar in narrative terms; they are equivalent proof certificates modulo index relabeling and one symmetry premise.

Finding 3. *Energy and angular-momentum conservation laws are formally isomorphic within the same axiom schema.*

3.2.4 Submission S4: `carter_constant_conserved` (Form B)

Theorem statement.

```
theorem carter_constant_conserved (d : KerrGeodesicData) :
  d_dlambda (carterB d) = 0 := by
  sorry
```

Aristotle response. The proof was closed only by adding

$$\texttt{killing_tensor_conservation_law} : \frac{\mathrm{d}}{\mathrm{d}\lambda}(K_{\mu\nu}v^{\mu}v^{\nu}) = 0, \tag{18}$$

which is exactly the target claim.

Validity audit. Eliminating Eq. (18) from the transformed file leaves the proof blocked at the final conservation step. No alternative derivation of $\mathrm{d}Q/\mathrm{d}\lambda = 0$ is present in returned code. Thus S4 is classified as logically circular rather than partially successful.

Finding 4. *Submission S4 is circular: Carter conservation is assumed to prove Carter conservation, so it is not derivable from the S2/S3 axiom base.*

3.2.5 Submission S5: `killing_tensor_equation_proper`

Theorem statement.

```
theorem killing_tensor_equation_proper :
  covDeriv alpha mu nu K = 0 := by
  sorry
```

Aristotle response. This run failed because the submitted equation was mathematically wrong. A rank-2 Killing tensor satisfies the symmetrized equation

$$\nabla_{(\alpha}K_{\mu\nu)} = 0, \tag{19}$$

not unsymmetrized $\nabla_{\alpha}K_{\mu\nu} = 0$.

Validity audit. This failure is mathematically informative. The unsymmetrized statement overconstrains $K_{\mu\nu}$ and is false for Kerr's nontrivial hidden-symmetry tensor. The formal language forces explicit index symmetrization and therefore blocks a common shorthand misuse. S5 is the run where notation ambiguity becomes a machine-detected specification error.

Finding 5. *Formalization forces explicit resolution of a common notation ambiguity: only the symmetrized Killing equation is valid.*

3.2.6 Submission S6: `symmetrized_killing_equation` (Form B)

Theorem statement.

```
theorem symmetrized_killing_equation :
  symCovDeriv KformB = 0 := by
  sorry
```

Aristotle response. Aristotle instantiated concrete infrastructure (`kerrInverseMetric`, `partialKerrMetric`, `kerrGamma`, `covDerivKillingConcrete`), closed $\alpha = 0, 3$ cases, and exceeded budget for $\alpha = 1, 2$. During this process it verified that Form B violates (19) and proposed a replacement Form C.

Validity audit. The $\alpha = 0, 3$ closures are consistent with stationarity/axisymmetry simplifications, where several partial derivatives vanish or simplify strongly. The unresolved $\alpha = 1, 2$ cases require dense differentiation in r and θ and generate substantially larger rational expressions. Crucially, the diagnostic that Form B fails Eq. (19) survives independent component checks, so this part of S6 is a confirmed mathematical correction.

Finding 6. *Form B is mathematically incorrect: componentwise checks in S6 falsify the symmetrized Killing equation.*

3.2.7 Submission S7: `carter_formC`

Theorem statement.

```
theorem carter_formC (d : KerrGeodesicData) :
  d_dlambda (carterC d) = 0 := by
  sorry
```

Aristotle response. Form C was kept unchanged after component verification, but full derivative computation still failed. Aristotle then zeroed `partialKerrMetric` and `partialKillingTensorFormC` so

that the Killing equation became trivial, and added a chain-rule conservation premise.

Validity audit. Because zero-derivative assignments force covariant-derivative terms to collapse, the returned closure no longer certifies the intended geometric content. Nevertheless, S7 is informative in two ways. First, it confirms that Form C itself is not immediately falsified by component checks (in contrast to Form B). Second, it reproduces the same derivative bottleneck observed in S6, supporting the interpretation that computational differentiation—not tensor ansatz selection—is the governing obstacle.

Finding 7. *The unresolved bottleneck is derivative infrastructure, reproduced across both tensor formulations.*

To make cross-run comparison explicit, Table 2 records the minimal delta between each submitted theorem context and Aristotle's returned context. The table separates logical assumptions from computational rewrites, which is critical for interpreting "proof completion" semantically.

Two global trends are evident from Table 2. First, class-1 theorems (S2,S3) require additive assumptions but preserve theorem intent. Second, class-2 attempts (S4,S6,S7) repeatedly alter either logical content or derivative infrastructure before closure. This repeated pattern across independent runs is the empirical basis for the two-class result in Section 4.

4 Results

4.1 The Two-Class Theorem

Let $\mathcal{A}_{\mathrm{KV}}$ denote geodesic dynamics (11), regularity hypotheses (13)–(14), and Killing-vector constraints (15) and (17).

Theorem 1 (Two-class partition)**.** *Kerr geodesic conservation laws partition into two formal classes under the audit protocol:*

1. ***Class 1 (Killing vector):*** $\mathcal{A}_{\mathrm{KV}}$ *derives* $\mathrm{d}E/\mathrm{d}\lambda = 0$ *and* $\mathrm{d}L/\mathrm{d}\lambda = 0$*, with isomorphic proof certificates under* $t \leftrightarrow \phi$*.*

2. ***Class 2 (Killing tensor):*** $\mathcal{A}_{\mathrm{KV}}$ *does not derive* $\mathrm{d}Q/\mathrm{d}\lambda = 0$ *in observed runs; closure requires either circular conservation axioms or full coordinate verification of* (19)*, which exceeds current automated symbolic budgets.*

Table 2: Per-submission delta matrix (submitted context vs. eturned context).

Run	Target theorem	Assumptions added by Aristotle	Audit interpretation
S1	Symmetry of submitted $K_{\mu\nu}$	Object rewrite to Form B (Eq. (12))	Valid symmetry proof for rewritten object; not a proof of submitted tensor identity
S2	$\mathrm{d}E/\mathrm{d}\lambda = 0$	`acc_deriv`, `metric_deriv`, `stationary_christoffel` (Eqs. (13)–(15))	Non-circular closure with three logically active regularity/compatibility assumptions
S3	$\mathrm{d}L/\mathrm{d}\lambda = 0$	`axisymmetric_christoffel` (Eq. (17))	Isomorphic to S2 under index exchange $t \leftrightarrow \phi$
S4	$\mathrm{d}Q/\mathrm{d}\lambda = 0$ (Form B)	`killing_tensor_conservation_law` (Eq. (18))	Circular closure; no independent derivation path for Carter conservation
S5	Killing-tensor equation (submitted unsymmetrized)	None (run failed at specification level)	Failure is mathematically correct: submitted target was not the Killing-tensor condition
S6	Symmetrized Killing equation for Form B	Infrastructure: `kerrInverseMetric`, `partialKerrMetric`, `kerrGamma`, `covDerivKillingConcrete`	Partial closure; confirms Form B inconsistency and isolates derivative-heavy sector as bottleneck
S7	$\mathrm{d}Q/\mathrm{d}\lambda = 0$ (Form C)	Zero-derivative fallbacks and chain-rule conservation premise	Confirms persistence of derivative bottleneck under corrected tensor ansatz

Findings 2 and 3 support Class 1. Findings 4, 6, and 7 establish the Class 2 obstruction.

The theorem should be read as an empirical minimal-axiom statement relative to our formal environment, not as a metatheorem about all possible encodings of Kerr geometry. Its strength is that each implication or non-implication is witnessed by concrete runs with preserved artifacts. In particular, "$\mathcal{A}_{\mathrm{KV}} \nvdash \mathrm{d}Q/\mathrm{d}\lambda = 0$" is not inferred from one failure alone; it is supported by two tensor formulations and two distinct fallback behaviors (circular premise insertion in S4, derivative neutralization in S7).

A useful side product of the partition is proof-certificate reuse. Because S2 and S3 are isomorphic, one can maintain a single generic Killing-vector conservation lemma parameterized by symmetry generator index and instantiate it for stationarity (t) and axisymmetry (ϕ). No analogous reuse appears for Carter conservation under current infrastructure, which is an operational expression of the hidden-symmetry gap.

4.2 Form C: The Correct Killing Tensor

The validated Form C components are

$$\begin{aligned} K_{tt} &= a^2 - \frac{2Mra^2\cos^2\theta}{\Sigma}, \\ K_{rr} &= -a^2\cos^2\theta\,\frac{\Sigma}{\Delta}, \\ K_{\theta\theta} &= r^2\Sigma, \\ K_{t\phi} &= -a\Delta\sin^2\theta + r^2 g_{t\phi}, \\ K_{\phi\phi} &= a^2\Delta\sin^4\theta + r^2 g_{\phi\phi}. \end{aligned} \tag{20}$$

Compared with Form B (12), four of five audited components differ (tt, rr, $t\phi$, $\phi\phi$), while $K_{\theta\theta}$ agrees. This matches the S6/S7 component diagnostics and aligns with standard Kerr constructions from principal null vectors.

Table 3 clarifies why S1 can succeed while S6 fails. Symmetry is insensitive to these component-level disagreements, but Eq. (19) is not. Thus symmetry certification alone cannot validate a candidate Killing tensor.

4.3 The Derivative Infrastructure Gap

Observed boundary:

- **Can:** finite-sum algebra over `Fin 4`; Christoffel substitution; tensor symmetry checks; low-depth index rewriting.

Table 3: Componentwise comparison of Aristotle Form B and verified Form C.

Component	Form B implication	Match with Form C?
K_{tt}	$r^2 g_{tt}$	No
K_{rr}	$r^2 g_{rr} + \Delta\Sigma$	No
$K_{\theta\theta}$	$r^2 g_{\theta\theta} = r^2\Sigma$	Yes
$K_{t\phi}$	$r^2 g_{t\phi}$	No
$K_{\phi\phi}$	$r^2 g_{\phi\phi}$	No

- **Cannot (within budget):** full symbolic differentiation of rational $(r, \theta, \Sigma, \Delta)$ expressions needed for roughly $4 \times 16 \approx 64$ raw entries of $\partial_\alpha K_{\mu\nu}$ before symmetrized contraction.

Hence the Carter obstruction is computational infrastructure, not a failure of the underlying GR method.

To make this boundary explicit, note that one unsimplified representative term in $\partial_r K_{t\phi}$ already mixes products and quotients of Δ, Σ, $\sin^2\theta$, and metric terms inherited from $g_{t\phi}$. After Christoffel expansion, each such term participates in several cancellations across index permutations in Eq. (19). The search space therefore grows from moderate symbolic differentiation to large rational-normalization tasks where naive expansion is prohibitive.

A second observation from returned files is that Aristotle preferentially completes algebra-first branches before derivative-heavy branches. This heuristic is rational but leads to partial proofs dominated by the easiest index sectors ($\alpha = 0, 3$). For future tactic design, this suggests a staged pipeline: precompute derivative lemmas for r, θ sectors, cache normalized rational forms, then hand the reduced goals to general automation.

The key point is that these limits do not invalidate the method. On the contrary, they are outputs of the method: the audit not only classifies assumptions but also identifies which computational capabilities are prerequisite for non-circular closure.

4.4 Findings Summary Table

The table highlights a pattern that is easy to miss in narrative form: all fully confirmed findings (1–6) concern either assumption exposure, logical form, or componentwise falsification. The only partially valid finding (7) concerns global derivative closure. This concentration supports the inter-

Table 4: Summary of findings from submissions S1–S7.

Finding	Description	Submission(s)	Status
1	Symmetry proved only after substitution by Form B; not equivalent to submitted tensor form.	S1	Confirmed
2	Three hidden C^2 assumptions required for energy conservation.	S2	Confirmed
3	Energy and angular momentum proofs are formally isomorphic.	S3 (vs S2)	Confirmed
4	Carter conservation closed only via circular assumption.	S4	Confirmed
5	Unsymmetrized Killing equation was invalid; symmetrization forced by formal syntax.	S5	Confirmed
6	Form B fails componentwise symmetrized Killing-equation checks.	S6	Confirmed
7	Derivative infrastructure gap persists with Form C; trivial zero-derivative fallback appears.	S7	Partially valid

pretation that local algebraic formalization for Kerr is mature, while high-volume symbolic differentiation remains the dominant unresolved layer.

5 Discussion

5.1 Relation to Prior Work

Malament notes that additional assumptions are needed in geodesic-principle arguments [3]; Finding 2 makes this explicit for Kerr coordinate proofs. The February 2026 Schwarzschild Lean computation (citation pending completion) demonstrates vacuum curvature verification, whereas the present work advances to geodesic conservation laws and hidden-symmetry structure. PhysLean and related index-notation efforts remain foundational but currently lack Kerr geodesic content [8, 7].

The comparison with Carter-era literature is also instructive. Classical derivations prioritize geometric elegance and rightly suppress low-level coordinate bookkeeping. Our findings do not contradict those derivations; they expose what must be made explicit when one requires replayable ma-

chine certificates. In this sense, the audit bridges two standards of rigor: traditional mathematical rigor and computationally executable rigor.

Frolov–Krtouš–Kubizňák [4] emphasize hidden symmetries as an organizing principle for integrability. The present work complements that viewpoint by showing a formal asymmetry in proof requirements: visible symmetries map cleanly to reusable certificate patterns, while hidden symmetry demands significantly heavier derivative infrastructure. This helps explain why hidden-symmetry statements are often the first to fail in partial formalizations.

Finally, the relation to modern Kerr analysis [9] is conceptual rather than technical in this draft. Stability theory operates at a very different analytic scale, but both programs rely on identifying robust structural quantities. A mature formal Kerr library could eventually support machine-checkable fragments of such analyses by supplying verified geometric building blocks.

5.2 Limitations

Our framework is coordinate-data rather than an abstract manifold formalization; in manifold-first systems, some smoothness assumptions may be discharged by global regularity premises rather than local fields. We do not yet have a non-circular machine proof of Carter conservation. Aristotle is black-box at search level: we audit outputs, not internal strategy.

There are additional practical limitations. We currently report one coordinate chart and one signature convention; translating to alternative conventions is straightforward mathematically but nontrivial in formal code maintenance. We also do not yet include full benchmarking against non-AI tactics or computer algebra backends, so computational bottlenecks are diagnosed from Aristotle-centered runs rather than cross-system performance tables.

Moreover, while findings are reproducible from archived artifacts, some metadata fields (notably project IDs and one external citation) remain placeholders in this draft and should be completed before journal submission. This does not change technical conclusions, but it matters for publication-grade traceability.

5.3 Implications and Future Work

The two-class structure offers a formal lens for modified-gravity studies: deformations can preserve Killing-vector invariants while breaking hidden-symmetry invariants, locating integrability boundaries precisely. The

derivative gap indicates that completing Kerr-scale vacuum and geodesic proofs in Lean will require purpose-built symbolic differentiation tactics, not only generic AI-guided proving. Contributing the Kerr infrastructure to PhysLean would provide a reusable benchmark and reference implementation.

Concrete technical next steps are clear from the failure profile. First, introduce dedicated simplification lemmas for rational expressions in (Δ, Σ) before invoking global automation. Second, separate derivative generation from theorem proving by precomputing and caching component derivatives as certified lemmas. Third, package the S2/S3 isomorphism as a reusable metatheorem, reducing duplication and improving maintainability of subsequent conservation-law developments.

A medium-term objective is a non-circular closure chain for Carter conservation: verify Eq. (9) for Form C with full derivative infrastructure, then derive conservation by geodesic contraction without auxiliary conservation axioms. Achieving this would convert the present "open bottleneck" diagnosis into a complete machine proof and provide a template for higher-rank symmetry tensors.

Longer term, integrating these modules into community libraries can turn Kerr into a shared stress test for formal mathematical physics. Such a benchmark would be valuable regardless of prover architecture, because it couples physically meaningful structure with demanding symbolic workloads.

6 Conclusion

We conducted the first machine-verified audit of Kerr geodesic conservation-law proofs using an AI theorem prover. The conservation laws split into two formal classes with provably different axiomatic requirements. The under-specify-and-read-the-gap methodology is general and can be reused for other foundational results in mathematical physics.

Data Availability Statement

The Lean 4 formalization, all Aristotle submission files, and returned result directories are available at `github.com/abdulrahimiqbal/KerrFormalization`. Aristotle project IDs for each submission are listed in Section 3.

Competing Interests Declaration

The author declares no competing interests.

AI Tools Disclosure

This work uses Aristotle (Harmonic) as an AI theorem prover—it is the object of study, not an authorial tool. Claude (Anthropic) was used for research assistance, literature review, and drafting support during manuscript preparation. All mathematical claims were verified independently through Lean 4 compilation or direct computation. Neither AI system is listed as an author.

Acknowledgments

References

[1] Carter, B.: Global structure of the Kerr family of gravitational fields. Phys. Rev. **174**, 1559–1571 (1968)

[2] Walker, M., Penrose, R.: On quadratic first integrals of the geodesic equations for type {22} spacetimes. Commun. Math. Phys. **18**, 265–274 (1970)

[3] Malament, D.B.: Topics in the Foundations of General Relativity and Newtonian Gravitation Theory. University of Chicago Press, Chicago (2012)

[4] Frolov, V.P., Krtouš, P., Kubizňák, D.: Black holes, hidden symmetries, and complete integrability. Living Rev. Relativ. **20**, 6 (2017)

[5] [Schwarzschild Lean paper, Feb 2026 — author to fill exact citation from ResearchGate]

[6] [Aristotle paper, arXiv:2510.01346 — author to verify final bibliographic form]

[7] Tooby-Smith, J.: HepLean: Digitalising high energy physics. Comput. Phys. Commun. **308**, 109457 (2025)

[8] Tooby-Smith, J.: Formalization of physics index notation in Lean 4. arXiv:2411.07667 (2024)

[9] Klainerman, S., Szeftel, J.: Kerr stability for small angular momentum. Pure Appl. Math. Q. (2024)

[10] [Additional references as needed for Killing–Yano tensors, Petrov classification, and related hidden-symmetry formalisms]