

Gauge Freedom Journal: Proof \otimes Presence for Human–AI Science

Marcelo Maciel Amaral^{1,2} and Raymond Aschheim^{1,2}

¹Gauge Freedom, Inc. (California Benefit Corporation), Los Angeles, CA, USA

²Gauge Freedom Journal, Editorial Office

EDITORIAL GAUGE FREEDOM JOURNAL (2026) 1:1

<https://doi.org/10.65323/gfj.2026.001>



Correspondence: Editorial Office (editor@gaugefreedom.org)

Received: 01 January 2026 / *Accepted:* 09 January 2026

Published online: 20 January 2026

© 2026 The Author(s). Licensed under CC BY 4.0.

Abstract

Science is entering an era in which Human–AI collaboration is not an exception but a default mode of work. This change is already producing legitimate novelty, including peer-reviewed work where AI-assisted exploration contributed to key steps later subjected to independent verification. However, it also brings a new class of failure modes: opaque AI influence, fragile reasoning masked by fluent prose, and a growing difficulty in distinguishing rigorous advances from plausible artifacts. *Gauge Freedom Journal* (GFJ) launches to help establish a new norm: provenance-first, open, energy-aware physics publishing. We encourage a dual-layer disclosure standard for AI-assisted research: **AI Integrity Reports** as concise human-readable summaries and **Content-Addressable Receipts (CARs)** as machine-verifiable audit trails. Our aim is simple: as AI accelerates discovery, the scientific method must become more visible, not less.

Keywords— gauge theory; information geometry; agentic science; provenance; reproducibility; AI Integrity; CAR; human–AI symbiosis

1 The Signal: From Black Box to Glass Box

The research workflow has changed faster than the institutions that validate it. LLMs now assist with derivations, literature synthesis, code generation, and the rapid exploration of alternative hypotheses. We have already seen peer-reviewed work where authors describe structured AI-assisted workflows and emphasize independent verification of key steps before elevating them into scientific claims.

This is not a crisis; it is a transition. The deeper risk is not that AI is used, but that it is used invisibly. The risk is that the scientific record fails to show how those ideas were produced, challenged, corrected, and refined.

In earlier eras, journals matured by standardizing what “good method” looks like: clear derivations, reproducible experiments, and shared data. The Human–AI era needs an equivalent standard for cognitive provenance.

2 The Problem: The Next Reproducibility Wave

The reproducibility crisis is about to collide with a generative flood. A single-pass AI draft can look persuasive while hiding realistic failure modes: misstated assumptions, subtle logical leaps, missing citations, or computational mistakes. These errors are not always malicious; they are structural properties of probabilistic tools used without sufficient constraints.

We therefore propose a simple institutional upgrade: When AI materially contributes to a scientific claim, the method of AI involvement should be as inspectable as the claim itself. For clarity, “material contribution” refers to substantive involvement in ideas, derivations, analysis, code, or figure generation, beyond copy-editing, translation, or purely stylistic polishing.

This does not require authoritarian policing. It requires a scientific norm: show your work—human and machine alike.

In the AI era, the unit of credibility is not the manuscript alone, but the manuscript plus its provenance.

3 Our Mission: Verifiable, Open, Energy-Aware

GFJ is a public-benefit-aligned journal dedicated to a pragmatic premise: trusted science in the AI era requires auditable method. We anchor this mission with three commitments:

1. **Content Over Credentials.** GFJ uses double-blind peer review for all submissions. We evaluate the physics, not the prestige, affiliation, or network position of the authors.
2. **Provenance as a First-Class Variable.** When AI materially contributes to text, code, analysis, or figures, authors should include an AI Contributions Statement and are strongly encouraged to provide a provenance artifact that allows reviewers and future readers to understand what was done, by which tools, under which constraints.
3. **Energy Stewardship.** Computation is a physical process. We encourage reporting of computational cost and, where feasible, approximate energy use for AI-intensive steps. Efficiency is not a vanity metric; it is part of scientific responsibility.
4. **Open Materials by Default.** Where feasible, we encourage open code, data, prompts/workflows, and AI provenance artifacts (AI Integrity Reports and CARs). Exceptions for privacy, security, or contractual constraints should be stated clearly.

4 The Method Standard: Ensemble Research

A recurring lesson from modern AI practice is that reliability improves dramatically when tools are used in structured roles. A growing standard is the **Generator–Verifier** protocol: one model proposes a step, another independently challenges it, and humans adjudicate disagreements.

GFJ encourages this mindset for high-stakes claims, especially where a small conceptual error could invalidate a large structure. We call this **Ensemble Research**: the deliberate orchestration of complementary AI tools to reduce hallucination risk and improve methodological clarity.

The key shift is cultural: AI assistance should be treated like an instrument in a lab. It can be powerful, but it must be calibrated, cross-checked, and documented in the scientific record.

5 The Artifacts: Integrity Report and CAR

To make these standards operational, GFJ encourages a dual-layer disclosure approach:

- **AI Integrity Report.** A concise, human-readable integrity summary of AI involvement that may include the stated role of AI, key claims, and flagged fragile or under-supported passages. Preferably generated by an independent verifier or tool (not by the author alone).
- **Content-Addressable Receipt (CAR).** A machine-verifiable record binding key inputs, model/tool identifiers, timestamps, and integrity metadata to a cryptographic content address. Where workflows are multi-step, CARs can represent hash-linked chains of major actions.

These artifacts help shift AI disclosure from vague narrative paragraphs to inspectable scientific metadata.

Handling During Double-Blind Review

During double-blind review, authors may submit full AI Integrity Reports and CARs to the editorial office. Reviewer-facing versions, if shared, will be anonymized. Public artifacts may be linked after acceptance.

Implementation Without Lock-In

GFJ follows a **standard-first** approach to provenance. We encourage the pairing of an **AI Integrity Report** with a **CAR** when AI materially contributes to claims. Gauge Freedom, Inc. develops *Intelexta*, which GFJ uses as a pilot implementation of this Integrity Report + CAR approach, but use of Intelexta is not required. Equivalent provenance methods that provide comparable human-readable summaries and machine-verifiable audit trails are welcome.

This distinction is essential: GFJ is advancing a standard of method, not enforcing a single vendor tool.

6 Scope: Physics with Deep Structure

GFJ's methodological commitments serve a scientific core. We publish work at the interface of:

1. **Gauge Theory Foundations & Applications.**
2. **Information Geometry & Statistical Field Theory.**
3. **Quantum Information, Computation & Technologies.**
4. **Quantum AI Research & Applications.**
5. **Plasma Physics, Nuclear Fusion & Quantum Energy.**

6. Condensed Matter, Strong Correlations & Spintronics.
7. Computational Physics, Simple Programs & Reproducible Science.
8. Human–AI Symbiosis.

7 Governance, Independence, and Trust

GFJ is governed by a **Scientific Editorial Board** comprising researchers from independent academic institutions. While Gauge Freedom, Inc. provides technical infrastructure, **editorial decisions are strictly merit-based and independent of corporate interests**.

Service Level Agreement (SLA)

We respect the author's time and treat these timelines as performance targets, not guarantees:

- **Triage:** Target \leq 72 hours.
- **First Decision:** Target \leq 7 days.
- **Transparency:** Queue health and median timelines will be public.

8 Discussion: The Road to Machine-Verifiable Provenance

The rapid convergence of LLMs and scientific discovery demands a new standard for trust. While the Content Authenticity Initiative and the C2PA standard have established a crucial tamper-evident metadata framework for general digital media, the requirements for high-stakes scientific claims are more stringent. This motivates the dual-layer disclosure commitment of this journal: a concise **AI Integrity Report** paired with a **CAR** when AI materially contributes to claims.

8.1 Watermarking, Hashing, and Quantum Resilience

The current landscape features two primary, complementary approaches to AI provenance:

- **Cryptographically Signed Metadata (C2PA):** This method, which we leverage conceptually with our CARs, secures content by hashing it (typically using SHA-256) and cryptographically signing the manifest. Any subsequent change to the content or the record breaks the verifiable chain.
- **Invisible Watermarking (e.g., SynthID):** A complementary approach embeds an imperceptible, resilient signal directly into the content. Unlike metadata, this watermark is designed to survive routine modifications like compression.

This technical duality immediately points to future-facing cryptographic problems. Specifically, our reliance on algorithms like SHA-256 for hashing and current signature schemes is threatened by the coming era of quantum computation. Grover's algorithm is known to reduce the effective security level of classical hash functions used in receipts, motivating long-horizon planning for post-quantum robustness in the scientific record.

To ensure the long-term integrity of our scientific record, we must anticipate this threat. The transition to Post-Quantum Cryptography (PQC) is already underway, with NIST standardizing new algorithms (including lattice-based schemes). The long-term security of scientific Content-Addressable Receipts will likely require an agile switch to quantum-resistant signature options.

8.2 The Path to Human-Free Provenance

The fundamental weakness in any current system—be it C2PA, invisible watermarking, or a traditional paper trail—is the single-point human in the loop who possesses the critical signing key or can choose to be deceptive. A long-term objective of Human–AI Symbiosis is to reduce these single points of failure while preserving accountable scientific authorship.

AI Integrity Reports and accompanying **CARs** are designed to pave the way for this future:

- **CAR as the Agent’s Log:** A machine-verifiable record. Future evolutions may be automatically generated and signed by the AI Agent itself using short-lived, scope-limited cryptographic credentials.
- **Autonomous Proof of Work:** The provenance chain could become increasingly machine-generated—from the LLM’s initial suggestion to subsequent computational checks—with each step automatically sealed and hash-linked.
- **The New Role of the Human:** The researcher’s role shifts from sole signer to adjudicator: interpreting integrity reports and verifying machine-logged provenance.

This is the central tenet of auditable, agentic science: as AI accelerates discovery, the scientific method must become more visible, not less.

9 Invitation to the Next Builders

GFJ is for researchers who want speed without shortcuts, openness without ambiguity, and AI assistance without invisible authorship. We equally welcome rigorous submissions with minimal or no AI involvement. We invite independent thinkers, interdisciplinary teams, and early adopters of rigorous, documented Human–AI discovery.

Proof \otimes Presence is not a slogan. It is a compact promise: scientific progress with AI must remain transparent enough to be trusted, reproduced, and improved by others.

Declarations

Acknowledgments

We thank the Founding Editorial Board for their essential guidance in shaping the journal's launch standards: Douglas Oliveira Novaes (University of São Paulo), Tiago Dourado (Brazilian Center for Research in Physics), Srinjoy Ganguly (University College London), Shalini Devendrababu (Fractal Quantum AI Lab), Bhavyadhir V. Bharadwaj (IIT Guwahati / GENORROW Enterprises), and Michael P. Rios (Dyonica ICMQG).

AI Contributions Statement

This editorial is a product of Human–AI collaboration. The authors are responsible for all ideas, framing, and scientific claims. Generative AI systems were used to test alternative structures, synthesize literature, and polish the final prose. To exemplify GFJ's provenance-first standard, we provide an accompanying **AI Integrity Report** and **Content-Addressable Receipt (CAR)**. These artifacts were generated using the journal's pilot verification infrastructure (Intelexta), though the standard allows for equivalent third-party provenance tools.

Artifact Availability

The specific LaTeX source and provenance artifacts for this editorial are available via the article landing page to serve as a template for future submissions.

Content-Addressable Receipts (optional)

An accompanying Content-Addressable Receipt (CAR) and AI Integrity Report are linked to this article.

Reproducibility Checklist

Not applicable for an editorial.

Energy Disclosure

Drafting and refinement utilized inference from foundation models including Gemini 3 Pro and GPT-5.2. The validation process (generating the Integrity Report) utilized Intelexta systems. Total inference-related energy consumption is negligible relative to training and global baseline costs.

Conflicts of Interest

The authors serve as Co-Founding Editors of Gauge Freedom Journal. M.M.A. and R.A. are also executives at Gauge Freedom, Inc., a California Benefit Corporation that provides the journal's open-source web infrastructure and the Intelexta provenance tool. Editorial decisions are made independently by the Scientific Editorial Board.

Funding

Gauge Freedom Journal is supported by Gauge Freedom, Inc. as part of its public benefit mission. No external grant funding was received for this manuscript.

Ethics

Not applicable.