Patent Draft: Epistemic Trace Validator for Citation Integrity Detection

🔬 Abstract

The *Epistemic Trace Check* is a standalone validation module within the LucidLock framework, designed to assess the structural integrity of citations within academic and scientific documents. Unlike plagiarism detectors or citation format checkers, this system evaluates whether references form a coherent epistemic backbone — or merely simulate scholarly rigor through overcitation, mimicry, or templated authority patterns. By isolating epistemic scaffolding signals, it enables transparent evaluation of how claims are supported and whether the citation logic reflects genuine intellectual lineage. This tool serves as a diagnostic integrity layer for scholarly and AI-generated content.

📜 Claims

1. **Structural Citation Analysis Module**
   A system that identifies and maps citation use across academic documents to determine whether references serve as functional epistemic anchors or ornamental authority markers.

2. **Detection of Mimicry & Template Signals**
   The system detects template-based citation clusters, redundant authority stacking, overcitation of common knowledge, and reused bibliographic scaffolds across unrelated sections.

3. **Volitional Source Use Detection**
   The validator assesses whether the author demonstrates deliberate and recursive engagement with cited sources, including evidence of conceptual integration, numerical grounding, or falsifiability logic.

4. **Epistemic Drift and Void Mapping**
   The system flags cases where major claims lack adequate support, references are misaligned with the claims they accompany, or foundational citations are absent for critical methods.

5. **Temporal and Self-Citation Ratio Assessment**
   The validator quantifies and interprets the balance of recent vs. foundational works, and

the ratio of external to self-citations, as a proxy for scholarly breadth and epistemic depth.

6. **Standalone Validator Execution**
   The module operates independently via a webhook-based input, accessing PDF documents from a connected data source (e.g., Google Drive), extracting text through structured indexing, and generating structured output reports in Markdown and HTML.

7. **Immutable Output for Archival Review**
   Each validation event produces a timestamped report suitable for archival on decentralized systems (e.g., IPFS, Zenodo), ensuring forensic traceability of citation structure over time.

📘 Description

The LucidLock Epistemic Trace Check is one of four core validators developed to restore trust in academic and AI-generated content. This module specializes in citation structure — treating citations not as peripheral metadata, but as epistemic signals that reveal the reasoning architecture beneath scholarly claims.

Unlike tools that verify citation count or formatting compliance, the Epistemic Trace Check examines **how** citations function: Do they recursively anchor claims in prior knowledge? Do they evolve through the paper's logic, from introduction to conclusion? Or are they mimicked, stacked, and repurposed as a mask of legitimacy?

The module applies a five-part reasoning frame:

1. **Citation Structure** – Checks whether references form a recursive epistemic chain, tracing ideas to their foundations.

2. **Density vs. Coherence** – Flags inflated citation density that lacks distinct support function.

3. **Template Signals** – Detects reused citation blocks, excessive clustering, and prestige mimicry.

4. **Citation Drift / Gaps** – Identifies unsupported major claims or mismatched references.

5. **Volitional Source Use** – Evaluates whether authors show deliberate, context-aware source engagement.

Each report returns a **PASS** or **FAIL** verdict, accompanied by metadata, diagnostic reasoning, and a breakdown of observed citation behaviors. It does not judge factual correctness, scientific merit, or writing quality — only the structural coherence and epistemic function of citations.

The system is deployable via webhook and compatible with cloud workflows (e.g., Make.com, DustTT, Firebase). Reports are timestamped and stored immutably (e.g., via IPFS or Zenodo) for external audit, enabling citation integrity reviews across academic submissions, journal workflows, and AI outputs.