

TRACE-AI

An Evidence-Based Research Intelligence Platform

White Paper

Version 1.0-beta (External Review Release)

June 2026

Authors

Sadia Rehman

Muhammad Amin ur Rehman Khan

Developed in association with

International Association for Social Research Dynamics (IASRD)

Guiding Principle “No Evidence = No Claim.”

Status: External Review Release (v1.0-beta)

This white paper describes a research software platform currently undergoing external validation and should not be considered a final production release.

This document describes a conceptual and developmental framework and should not be interpreted as evidence of completed technical implementation or validated performance.

Abstract

TRACE-AI is an evidence-based research intelligence platform designed to support transparent and auditable analysis of research documents. The platform aims to improve evidence traceability by linking findings to supporting source material and presenting confidence assessments and audit trails. AI assistance, where used, does not replace the evidence verification framework. TRACE-AI is guided by the core design principle: “No Evidence = No Claim.” The current release (v1.0-beta) represents an external review version intended for validation and feedback rather than a final production release.

Table of Contents

1. Introduction.....	1
2. Background and Problem Statement	1
2.1 The Evidence Attribution Gap	1
2.2 Confidence Without Calibration	1
2.3 Audit Trail Deficits.....	1
3. Guiding Principle: No Evidence = No Claim	2
4. System Architecture	2
4.1 Document Ingestion Layer	2
4.2 Evidence Extraction Layer	2
4.3 Verification and Cross-Reference Engine	2
4.4 Confidence Calibration Layer	3
4.5 Output and Audit Trail Generation.....	3
5. Core Features and Capabilities	3
6. Verification Framework.....	3
7. Validation and Evaluation	4
7.1 Evidence Linkage Accuracy Testing	4
7.2 Confidence Calibration Review	4
7.3 Practitioner Usability Review	4
8. Intended Applications and Use Cases	4
9. Scope and Limitations	5
10. Future Development Roadmap	5
11. Conclusion	5
12. Intellectual Property and Copyright Statement	6
13. Suggested Citation.....	6
14. References.....	6
15. Contact Information	6

1. Introduction

The proliferation of artificial intelligence in academic and professional research contexts has introduced significant opportunities for accelerating knowledge synthesis. However, these same tools present a growing risk: the generation of confident-sounding claims that are unsupported by verifiable evidence. Research consumers, policy makers, and practitioners frequently encounter outputs that lack clear citation trails, conflate interpretation with fact, or present speculative conclusions without adequate qualification.

TRACE-AI was conceived in direct response to this challenge. Developed by Sadia Rehman and Muhammad Amin ur Rehman Khan in association with the International Association for Social Research Dynamics (IASRD), TRACE-AI is an evidence-based research intelligence platform designed to apply a consistent evidence-oriented framework across every layer of its operation: No Evidence = No Claim.

This white paper presents the conceptual foundations, system architecture, key features, verification framework, and intended applications of TRACE-AI at its v1.0-beta external review release stage. It is addressed to researchers, academic institutions, research ethics bodies, and technology reviewers engaged in the responsible integration of AI tools into scholarly practice.

Note: *Features and workflows described in this document represent design objectives and may evolve during ongoing development and validation.*

2. Problem Statement

Contemporary research practice faces a compounding evidence integrity crisis. Three converging phenomena are driving this problem:

2.1 The Evidence Attribution Gap

Large language models and AI-assisted synthesis tools routinely produce summaries, conclusions, and recommendations without providing traceable links to the underlying source documents. Users are presented with outputs that appear authoritative but cannot be independently verified. In research contexts, particularly health, social care, and policy, this represents a material risk to evidence-based practice.

2.2 Confidence Without Calibration

Existing AI tools frequently assign uniform confidence or certainty to outputs regardless of the strength, consistency, or methodological quality of the underlying evidence base. A finding supported by a single grey literature source may receive the same presentation as one supported by multiple peer-reviewed systematic reviews this lack of calibration misleads practitioners and undermines trust in AI-assisted analysis.

2.3 Audit Trail Deficits

Regulatory and governance frameworks in health and social care, education, and policy research increasingly require demonstrable audit trails for decisions informed by evidence. Existing AI tools provide no structured mechanism to support this requirement. Practitioners using such tools cannot demonstrate the evidential basis for conclusions reached through AI-mediated analysis, creating both ethical and professional risk.

TRACE-AI is intended to mitigate these challenges through its architecture, verification framework, and output design.

3. Guiding Principle: No Evidence = No Claim

The foundational commitment of TRACE-AI is expressed in its guiding principle: No Evidence = No Claim. This principle is embedded as a core design objective across the platform's architecture and intended to guide every stage of its evidence processing and output generation.

Guiding Principle: *“No Evidence = No Claim.” TRACE-AI is designed to link assertions to identified and retrievable source material wherever possible. Where evidence is absent, qualified, or contested, this is stated explicitly rather than concealed behind confident language.*

In practical terms, this means:

- TRACE-AI is designed to link assertions to identified, retrievable source documents wherever technically possible.
- Where evidence is absent, qualified, or contested, the platform is intended to state this explicitly rather than generating plausible-sounding but unsupported conclusions.
- Confidence ratings are designed to reflect the nature and volume of available evidence, not the internal coherence of the model's language generation.
- Users are intended to be able to inspect the evidence chain underlying claims the platform produces.

This principle distinguishes TRACE-AI from generative AI tools that optimise for fluency and coherence. TRACE-AI is designed to prioritise traceability and epistemic honesty, even where this produces less rhetorically satisfying outputs.

4. System Architecture

TRACE-AI is designed as a modular, layered platform. The architecture separates evidence ingestion, claim extraction, verification, confidence calibration, and output generation into distinct functional layers, each designed around the No Evidence = No Claim principle.

4.1 Document Ingestion Layer

The ingestion layer accepts research documents in standard formats (PDF, Word, plain text). Documents are parsed, segmented, and indexed with metadata capturing source identity, publication date, document type, and access provenance. Each document is assigned a unique internal identifier that persists through all downstream processing.

4.2 Evidence Extraction Layer

The extraction layer identifies and isolates claims, findings, and conclusions within ingested documents. Extracted claims are anchored to their precise location within the source document — page, section, and paragraph ensuring that the evidential basis for every extracted statement can be independently verified by a human reader.

4.3 Verification and Cross-Reference Engine

The verification engine cross-references extracted claims against the broader document corpus to identify corroboration, contradiction, or absence of support. It applies a structured verification protocol that distinguishes between primary evidence (empirical studies, clinical trials, systematic reviews), secondary evidence (literature syntheses, expert commentary), and grey literature (policy documents, reports, working papers).

4.4 Confidence Calibration Layer

Confidence ratings are computed from the quantity, quality, and consistency of supporting evidence rather than from the internal probability scores of the language model. A finding supported by a single source receives a lower confidence rating than one corroborated across multiple independent high-quality studies, regardless of how fluently either can be expressed.

4.5 Output and Audit Trail Generation

Outputs are designed to include a structured evidence appendix presenting the evidence chain, confidence rating, and verification status of each claim. This appendix is designed to be human-readable and suitable for inclusion in professional, regulatory, and academic governance documentation.

5. Key Features

TRACE-AI v1.0-beta delivers the following core capabilities:

- **Evidence Traceability** — claims are intended to be linked to their source document, section, and location, enabling direct verification by the user.
- **Confidence Calibration** — confidence ratings reflect the strength and volume of available evidence, not model fluency.
- **Structured Audit Trail** — outputs are designed to include a machine- and human-readable evidence chain suitable for regulatory and governance documentation.
- **Source Type Classification** — documents are classified by evidence tier (primary, secondary, grey literature) and this classification is surfaced in outputs.
- **Contradiction Detection** — the platform is designed to identify and flag where extracted claims are contradicted or contested by other documents in the corpus.
- **Absence Signalling** — where evidence is absent or insufficient to support a claim, the platform is designed to report the evidential gap rather than generating unsupported content.
- **Document-Anchored Outputs** — generated outputs are intended to be anchored to specific locations within source documents, to minimise drift from the evidential record.
- **Transparency-First Design** — the platform is designed to support rather than replace professional judgement, with the aim of directing users back to the original source material.

6. Verification Framework

The TRACE-AI Verification Framework operationalises the No Evidence = No Claim principle across all platform functions. The framework consists of four sequential stages:

1. **Source Registration** — Documents submitted to TRACE-AI are intended to be registered with full provenance metadata before any processing occurs. Under the intended workflow, unregistered sources are excluded from claim generation.
2. **Claim Extraction and Anchoring** — Claims are extracted and anchored to their precise source location. Claims are intended to enter the processing pipeline only after a verified source anchor has been established.
3. **Evidence Corroboration Check** — Extracted claims are assessed against the full document corpus for corroboration, contradiction, or absence of support. The outcome of this check determines the confidence rating assigned to the claim.
4. **Output Certification** — Under the intended workflow, outputs are expected to include a verification certificate documenting the evidence chain, confidence rating, source tier classification, and any detected contradictions or evidential gaps. This certificate is designed

to support human-readable audit documentation suitable for professional and governance purposes.

This framework ensures that TRACE-AI outputs remain epistemically honest and auditable at every stage of the research intelligence process.

7. Validation and Testing

The v1.0-beta release is intended to undergo structured external validation. The planned validation programme is expected to include the following components:

7.1 Evidence Linkage Accuracy Testing

External validation activities are planned and may include subject matter experts in health, social care, and social research methods assessing the accuracy of TRACE-AI's claim extraction and source anchoring against manually verified reference outputs. Findings will be reported in subsequent release documentation.

7.2 Confidence Calibration Review

Planned validation activities include examination of the relationship between TRACE-AI confidence ratings and independent expert assessments of evidence strength, to evaluate whether calibration is meaningful and consistent across document types and evidence tiers.

7.3 Practitioner Usability Review

Structured usability review activities are planned for future development stages, with the intention of engaging practitioners in adult social care and health research contexts to assess the accessibility, interpretability, and practical utility of TRACE-AI outputs in professional decision-making environments.

Findings from the validation programme will inform revisions to the platform prior to subsequent production releases.

8. Intended Applications

TRACE-AI is designed for use in contexts where evidence integrity, auditability, and transparency are professional, regulatory, or ethical requirements. Primary intended applications include:

- **Academic Research** — supporting systematic review, evidence synthesis, and literature analysis with full citation traceability and audit documentation.
- **Health and Social Care Practice** — enabling practitioners and commissioners to interrogate the evidential basis of clinical and social care guidelines, policies, and service evaluations.
- **Policy Analysis** — supporting policy researchers and analysts in tracing the evidential foundations of policy documents and impact assessments.
- **Research Ethics and Governance** — providing research ethics committees and governance bodies with auditable documentation of the evidence basis for research proposals and findings.
- **Professional Education and Training** — supporting critical appraisal training by making the evidence verification process transparent, structured, and learnable.
- **Open-Access Publishing** — complementing open-access journals such as IASRD's Journal of Social Research Dynamics by supporting authors in producing transparently evidenced manuscripts.

9. Limitations and Scope

Notice: TRACE-AI is under active development. The v1.0-beta release represents an external review version intended for validation and feedback. Hybrid theme synthesis remains experimental and outside the scope of this release. Human interpretation remains essential, and findings should always be considered alongside the original source documents. Some cosmetic limitations may remain and are subject to future refinement.

In addition to the above, users should be aware of the following scope considerations for the v1.0-beta release:

- TRACE-AI processes documents submitted by the user; it does not independently search or retrieve documents from external databases or the internet.
- The platform's verification framework operates on the document corpus provided; it cannot verify claims against evidence not present in that corpus.
- Language support in v1.0-beta is limited to English-language documents.
- The confidence calibration framework is in active development; ratings should be treated as indicative rather than definitive in the beta release.
- TRACE-AI does not replace expert human review and is not intended to do so. All outputs should be reviewed by a qualified practitioner or researcher before informing professional or clinical decisions.
- TRACE-AI is not intended to provide medical, legal, or regulatory advice and should not be relied upon as such.
- TRACE-AI outputs may contain errors, omissions, or incomplete interpretations and should not be relied upon without independent review of the underlying source material.

10. Future Development

The TRACE-AI development roadmap following the v1.0-beta external review release includes:

- Hybrid Theme Synthesis — integration of thematic synthesis across evidence tiers with full cross-document traceability, which may be considered for future releases.
- Multilingual Support — extension of the platform's document processing and verification capabilities to non-English language documents.
- External Database Integration — structured interfaces with open-access academic databases to support corpus construction.
- Calibration Refinement — continued development of the confidence calibration framework based on findings from the external validation programme.
- Practitioner-Facing Interface — development of a simplified interface for use by frontline practitioners in health and social care settings without specialist research methods training.
- Open-Access API — a documented API to support integration of TRACE-AI verification capabilities into third-party research management and publishing platforms, including alignment with IASRD's open-access publishing infrastructure.

11. Conclusion

The integrity of evidence-based practice depends on the ability of researchers, practitioners, and policy makers to trace claims to their sources, assess the strength of the evidence base, and identify where evidence is absent or contested. Existing AI-assisted research tools have, to date, largely optimised for fluency and speed at the expense of these fundamental requirements.

TRACE-AI takes a different path. By embedding the principle of No Evidence = No Claim as a structural constraint rather than an aspiration, the platform is intended to promote transparency, auditability, and epistemic honesty across its outputs. In doing so, it aims to support rather than undermine the rigour that evidence-based practice demands.

The v1.0-beta external review release represents the beginning of a structured validation process. Feedback from external reviewers, subject matter experts, and practitioner users will be integral to the development of the platform toward its production release. The authors welcome engagement from researchers, practitioners, and institutions committed to advancing the responsible use of AI in research and professional practice.

12. Copyright Statement

Copyright © 2026 Sadia Rehman and Muhammad Amin ur Rehman Khan

Developed in association with the International Association for Social Research Dynamics (IASRD).

Except as permitted under applicable copyright law, no part of this white paper may be reproduced, distributed, or transmitted in any form or by any means, including photocopying, recording, or other electronic or mechanical methods, without the prior written permission of the authors. Brief quotations for scholarly, educational, review, and non-commercial purposes are permitted in accordance with applicable copyright law.

TRACE-AI is used as a project name by the authors. No claim is made over generic AI methodologies or evidence-based research principles. The original framework descriptions and documentation contained within this white paper are the intellectual property of the authors. Enquiries regarding permissions should be directed to IASRD at <https://iasrd.com>.

13. Suggested Citation

Rehman, S., & Khan, M. A. U. R. (2026). *TRACE-AI: An Evidence-Based Research Intelligence Platform (Version 1.0 Beta External Review White Paper)*. International Association for Social Research Dynamics (IASRD). [https://doi.org/\[DOI Pending\]](https://doi.org/[DOI Pending])

14. References

This document is intended as a platform concept paper rather than an empirical evaluation. Formal validation findings and references will be incorporated in subsequent releases.

15. Contact Information

Corresponding Authors: Sadia Rehman and Muhammad Amin ur Rehman Khan

Email: sadiarehman2816@gmail.com • aminrehman240@gmail.com

Project Website: <https://iasrd.com>

Journal Portal: submission.iasrd.com

Organisation: International Association for Social Research Dynamics (IASRD), UK Registered Charity No. 1217292

This white paper will be assigned a DOI upon publication. The DOI will be incorporated in the Suggested Citation and all subsequent versions of this document.