# Una Mens AI-Assisted Peer Review Confidentiality and Responsible Use Policy

**Purpose**
Una Mens permits carefully bounded use of artificial intelligence during peer review because AI systems can assist reviewers with orientation, inquiry, comparison, and reflection.
Their use does not transfer scholarly responsibility.

Reviewers remain responsible for reading the manuscript, evaluating its claims and evidence, deciding which questions matter, and making the final publication recommendation.

This policy establishes the confidentiality conditions under which AI may be used during Una Mens review.

**1. Manuscripts remain confidential**
Manuscripts submitted for peer review are confidential scholarly materials.

Reviewers must not disclose manuscript text, figures, data, identifying information, unpublished arguments, results, methods, or other nonpublic material except as permitted by the journal's review process.

The same obligation applies when using AI systems.

**2. AI systems are approved by use, not by name**
A system is not considered appropriate for confidential review merely because it is described as private, enterprise, institutional, secure, or paid.

Before manuscript content may be used with an AI system, Una Mens must have reasonable grounds to determine that the relevant deployment and settings provide appropriate safeguards for confidential scholarly material.

Reasonable grounds ordinarily require one of the following:

- **(a)** an institutional or contractual arrangement that disables training on inputs and limits retention of submitted material; or
- **(b)** a documented zero-data-retention configuration confirmed with the provider.
-

A provider's self-description as "private," "enterprise," or "secure" is not by itself sufficient. Nor is a paid tier, standing alone.

Relevant considerations may include:

- whether submitted material is retained;
- whether it is used to train or improve models;
- whether it may be reviewed by humans;
- whether it is shared with third parties;
- how long submitted material persists;

- whether deletion controls are available;
- what institutional or contractual protections apply; and
- whether the system is appropriate for the particular manuscript and material involved.

Approved configurations reduce exposure but do not eliminate it. Some providers retain flagged or safety-reviewed content even under zero-data-retention terms, and retention may be compelled by law. Reviewers should assume that material provided to any AI system may, in principle, be subject to legal process.

The journal maintains a **current approved configuration** on a separate, dated page, linked from this policy and reviewed each issue. That page names configurations and verification dates, not vendors or model versions, and may be updated without amendment to this policy. Approved systems or environments may change as technologies, terms, and institutional arrangements change.

**3. Public or unverified AI systems should not receive confidential manuscript content**
Reviewers should not upload unpublished manuscript text, figures, tables, data, supplementary materials, or other identifiable manuscript content into public, consumer, or otherwise unverified AI systems.

This restriction does not prevent reviewers from using such systems for questions that can be posed without revealing confidential material.

Examples may include:
- asking general methodological questions;
- exploring statistical concepts;
- considering broad theoretical alternatives;
- investigating literature that is already public;
- testing generic reasoning strategies; or
- working from sufficiently abstracted reviewer-created notes.

When in doubt, reviewers should minimize the amount and specificity of manuscript material disclosed.

**4. Use the minimum manuscript content necessary**
Even within an approved environment, reviewers are encouraged to use only the manuscript content reasonably necessary for the task.

AI-assisted inquiry should be proportionate.

A reviewer investigating one methodological concern may not need to provide an entire manuscript. A reviewer comparing interpretations may be able to work from a relevant section or from reviewer-created notes.

This principle reduces unnecessary exposure while preserving the usefulness of AI-assisted review.

**5. The reviewer directs the inquiry**
AI may serve as a second reader or thinking partner.

Appropriate uses can include:

- identifying alternative readings of an argument;
- challenging or strengthening a reviewer's interpretation;
- locating manuscript evidence relevant to a reviewer-generated question;
- exploring methodological or theoretical considerations;
- helping organize complex observations; and
- identifying questions that deserve closer human inspection.

AI should not independently determine manuscript merit, substitute for subject-matter expertise, or generate the publication recommendation.

The reviewer chooses the questions, evaluates the responses, and decides what matters.

**6. Independent human reading comes first**

Under the Una Mens HI–AI review process, reviewers are asked to read the manuscript and form an initial interpretation before engaging with AI-generated second-reader perspectives.

This preserves an identifiable human baseline against which later interpretations can be compared, revised, defended, or rejected.

AI-assisted inquiry is intended to widen or test judgment, not silently establish it.

**7. Authors are informed**

Authors whose manuscripts enter the Una Mens HI–AI review process are informed that approved AI systems may be used in bounded ways during peer review.

The journal describes the nature and limits of this process in its submission and review information, and maintains the current approved configuration on the dated page referenced in Section 2.

Where an interim configuration is in use, the journal states its retention limits plainly.

Authors retain the expectation that unpublished work will be treated as confidential and that AI use will occur only under the safeguards described in this policy.

Authors may request the specific configuration used in the review of their manuscript.

**8. Reviewers should disclose substantive AI use**

Reviewers should briefly disclose substantive AI assistance used during the preparation of a review. Routine spelling, grammar, or formatting assistance need not ordinarily be reported.

A disclosure may identify:

- the system used;
- the general purpose of the interaction; and
- whether manuscript material was provided directly.

Disclosure is intended to support transparency and calibration, not to penalize AI use. Reviewer disclosure is intentionally lighter than the author-side Collaboration Field Note; it is a calibration aid, not an audit.

**9. AI output is advisory, not authoritative**
AI systems may produce errors, unsupported claims, fabricated references, oversimplifications, or interpretations shaped by model limitations.

Reviewers should independently verify any AI-generated information that materially influences their assessment.

AI output should be treated as something to examine, not something to defer to.

**10. Sensitive manuscripts may require additional restrictions**
Some manuscripts may involve sensitive personal data, proprietary methods, unpublished datasets, vulnerable populations, clinical information, legally restricted materials, or other content for which ordinary safeguards are insufficient.

The editor may restrict or prohibit AI use for a particular manuscript even when an otherwise approved system is available.

**11. Confidentiality extends beyond the review itself**
Information learned through peer review should not be used for the reviewer's personal, scholarly, commercial, or competitive advantage.

Reviewers should not retain manuscript content in AI workspaces beyond what is necessary for the review, and should follow the session-clearing instructions published with the current approved configuration. Where an environment offers session deletion or retention controls, reviewers should use them at the close of review. Where it does not, reviewers should avoid persistent workspaces, saved conversations, and uploaded files entirely, working instead in a transient session.

**12. Human accountability remains final**
Una Mens treats AI-assisted review as a form of supported scholarly inquiry.

It does not transfer accountability.

The human reviewer remains responsible for:

- the interpretation of the manuscript;
- the accuracy of the review;
- judgments about validity, contribution, context, and significance;
- the fairness of the assessment; and
- the final publication recommendation.

**13. The journal holds itself to the same standard**
The journal's own AI-assisted manuscript mapping uses the configurations described in Section 2, under the same conditions: training disabled, retention minimized, and stateless operation wherever the environment permits.

The journal does not use consumer-tier or training-on-input systems for confidential manuscript content.

Authors may request the configuration used for their manuscript's map, as provided in Section 7.

---

**AI may widen the field of inquiry. Scholarly judgment remains human.**

---

*Current approved configuration: [https://www.unamensjournal.org/_files/ugd/10c32c_b00a0de2793b41a383aef108543915ca.pdf]. Last verified: [11/07/2026]. This register is reviewed each issue and supersedes prior entries without amendment to the policy above.*