

Responding to Suspected LLM-Written Hallucinated Peer Reviews

A Claim-Level Audit Protocol

Byunghyun Ban^{1,2}

1 Gyungkook National University, Department of Music

2 NaNa Lab

bhban@nanalab.kr

Abstract

Authors are often advised to respond to peer-review comments politely and point by point. However, this advice is insufficient when a formal review appears to be LLM-written or hallucinated: fluent, plausible, professionally formatted, but contradicted by the manuscript. This paper proposes a claim-level audit protocol for responding to suspected LLM-written hallucinated peer reviews. The protocol decomposes reviewer comments into atomic claims, classifies each claim as supported, partially supported, contradicted, generic, scope-mismatched, or not assessable, and maps each claim to manuscript evidence before selecting a response strategy. The protocol is illustrated using an anonymized case from an actual submission to a Korea Citation Index (KCI) candidate journal, in which one review contained multiple false missing-content claims. The goal is not to encourage accusations of AI use, but to help authors replace speculation with auditable evidence for revision, rebuttal, and editorial escalation.

Keywords: peer review, large language models, AI-generated peer review, hallucination, rebuttal letter, response to reviewers, research integrity, manuscript revision

1. Introduction

Responding to peer review is a routine part of scholarly publishing. Authors are commonly advised to remain professional, respond point by point, revise the manuscript clearly, and make it easy for editors and reviewers to understand what changed. This advice is valuable, but it assumes that reviewer comments are at least broadly grounded in the manuscript.

That assumption is increasingly fragile. A review can be fluent, formal, and plausible while making claims that are unsupported by, or contradicted by, the manuscript. This problem is especially important in the era of large language models (LLMs), because LLMs can generate review-like text with limited or no genuine engagement with the source document. The result may be a review that looks professional but contains hallucination-like claims: for example, stating that data, methods, parameters, limitations, or experimental details are missing when they are explicitly present.

The author-side problem is different from the editor-side or publisher-side problem. Editors and publishers may ask whether AI was used, whether journal policy was breached, whether confidentiality was violated, or whether a review should be discounted. Authors face a more immediate practical question: how should an author respond when a formal peer-review report appears LLM-written, hallucinated, or contradicted by the manuscript?

A weak response is to argue that the review “looks like ChatGPT.” Such claims are difficult to prove and may distract from the actual editorial issue. A stronger response is to show that specific review claims are contradicted by explicit manuscript content. In other words, the author should not begin with AI detection. The author should begin with claim-level audit.

Recent developments make this problem urgent. COPE has described editorial concerns about reviewers using artificial intelligence, including detailed but inaccurate review reports and breaches of journal policy [1]. Nature Portfolio states that peer reviewers remain accountable for the accuracy and views expressed in their reports and asks reviewers not to upload manuscripts into generative AI tools [2]. SAGE states that reviewers who inappropriately use generative AI tools to generate review reports will not be invited to review and that such reports will not be included in final decisions [3]. Case-based work has also begun to document first-person experiences with suspected AI-generated peer reviews [4], while recent detection-oriented research has investigated methods for identifying LLM-generated peer reviews [5].

However, detection is not enough. Even if a review appears AI-generated, the author may not be able to prove it. Even if a review is written by a human, it can still be ungrounded. Conversely, AI-assisted language polishing may not be problematic if the final report is accurate, grounded, disclosed where required, and responsibly verified. Therefore, for authors, the most actionable question is not “was this written by an LLM?” but “which claims in this review are supported, unsupported, or contradicted by the manuscript?”

This paper proposes a claim-level audit protocol for responding to suspected LLM-written hallucinated peer reviews. The protocol decomposes review comments into atomic claims, classifies each claim by its relationship to the manuscript, maps disputed claims to manuscript evidence, and selects an appropriate response strategy: revise, clarify, relocate, respectfully disagree, or escalate to the editor.

The goal is not to make authors more defensive. The goal is to make responses more precise, professional, and auditable.

The paper makes three contributions:

1. It reframes suspected LLM-written hallucinated review as an author-side evidence problem rather than an AI-detection problem.
2. It proposes a practical claim-level audit protocol for response letters and editorial escalation.
3. It illustrates the protocol using an anonymized case from a formal journal peer-review process.

2. Background

2.1 Conventional response-to-reviewer guidance

Traditional advice on responding to reviewers emphasizes professionalism, completeness, and clarity. Authors are encouraged to copy reviewer comments into a response document, respond point by point, describe manuscript changes, and remain respectful even when they disagree [6,7]. This advice is appropriate for ordinary peer review.

For example, if a reviewer correctly identifies insufficient explanation of an experimental parameter, the author can revise the methods section and explain the revision. If a reviewer asks for additional limitations, the author can expand the discussion. If a reviewer requests a clearer figure caption, the author can improve it.

However, conventional advice does not fully address hallucinated or ungrounded review claims. If a reviewer says “the hardware specifications are missing” when the manuscript already lists the hardware, a simple concession may be misleading. If the author writes, “We have now added the hardware specifications,” the response may falsely imply that the original manuscript lacked them. A better response is:

The manuscript already included the computational environment in Section X. To improve visibility, the relevant information has been moved to a dedicated subsection and expanded.

This distinction matters. Authors should revise when criticism is valid, but they should not silently accept false premises.

2.2 Suspected LLM-written review as a new response problem

Suspected LLM-written review creates a distinctive response problem. The issue is not only that the review may be inaccurate. It is that the review may be inaccurate in a particular way: fluent, generic, checklist-like, and confident while failing to track the source manuscript.

Such reviews may contain comments that are formally reasonable but substantively ungrounded. For example:

- “The dataset is not clearly described.”

- “The experimental setup lacks reproducibility details.”
- “The model parameters should be specified.”
- “The limitations are not discussed.”
- “The comparison with prior work is insufficient.”

Each of these comments may be valid in some manuscripts. The problem arises when the manuscript already contains the requested information. This failure mode is dangerous because it can influence editorial decisions while appearing legitimate.

The emergence of LLM-generated peer-review detection research underscores the seriousness of the issue [5]. However, detection research does not directly tell authors how to respond to a disputed review. A practical author-side protocol is still needed.

2.3 Why AI detection is the wrong first move

When authors receive a suspected AI-generated review, the immediate temptation is to say:

This review looks like it was written by ChatGPT.

This is usually a weak editorial argument. It is difficult to prove, emotionally charged, and potentially distracting. It may also invite an unhelpful debate over style rather than substance.

A stronger argument is:

Several claims in the review are contradicted by explicit manuscript content. The table below maps each disputed claim to the corresponding manuscript evidence.

This approach does not require proving AI use. It focuses on verifiable grounding. The review may have been written by an LLM, a human, or a human assisted by an LLM. In all cases, if the review contains claims contradicted by the manuscript, those claims should not be relied upon without correction.

This paper therefore treats suspected LLM-written review as a claim-verification problem.

3. Definitions

3.1 Suspected LLM-written hallucinated peer review

A suspected LLM-written hallucinated peer review is a formal review report that exhibits one or more of the following features:

1. fluent but generic phrasing
2. checklist-like revision requests
3. weak connection to manuscript-specific evidence
4. claims that information is missing when it is present
5. confident judgments without traceable support
6. possible stylistic resemblance to LLM-generated text

The term “suspected” is important. The protocol does not require proving that the review was actually generated by an LLM. It only requires identifying whether review claims are grounded in the manuscript.

3.2 Reviewer comment

A reviewer comment is a passage in a peer-review report. One comment may contain multiple claims.

Example:

The experimental setup is insufficient. The dataset is not clearly described, and hardware specifications should be provided for reproducibility.

This comment contains at least three claims:

1. the experimental setup is insufficient,
2. the dataset is not clearly described,
3. hardware specifications are missing or insufficient.

3.3 Atomic claim

An atomic claim is the smallest reviewable unit of meaning in a reviewer comment. The protocol operates on atomic claims rather than whole paragraphs.

3.4 Manuscript grounding

A review claim is grounded if it is supported by identifiable manuscript evidence. A grounded criticism specifies what is missing or problematic and where the issue appears.

3.5 False missing-content claim

A false missing-content claim occurs when a review states or implies that something is absent although the manuscript explicitly contains it.

Example:

The computational environment is not described.

If the manuscript already lists hardware and software tools, this claim is contradicted.

3.6 Partial validity

A claim is partially valid when the reviewer identifies a real weakness but overstates it.

Example:

The limitations are not discussed.

If limitations are discussed briefly but without mitigation strategies, a fairer version would be:

The manuscript mentions limitations, but the discussion should be expanded with mitigation strategies.

4. Claim-Level Audit Protocol

The proposed protocol has six steps.

Step 1. Preserve the review and avoid immediate accusation

The author should first save the original review report, decision letter, and manuscript version. Immediate emotional responses should be avoided. The first task is diagnosis, not rebuttal.

This step is especially important when the review appears AI-generated. Accusing the reviewer of AI use without evidence can weaken the author's position. The author should instead preserve the materials and prepare a claim-level audit.

Output:

- original review text,
- decision letter,
- submitted manuscript version,
- initial list of disputed claims.

Step 2. Decompose comments into atomic claims

Each reviewer paragraph should be split into atomic claims.

Table 1. Example of comment decomposition.

Reviewer comment	Atomic claims
"The dataset is not clearly described, and the hardware specifications are missing."	C1. Dataset description is unclear. C2. Hardware specifications are missing.
"The methodology is insufficient, and the comparison with prior work is weak."	C3. Methodology is insufficient. C4. Comparison with prior work is weak.

This prevents two errors. First, it prevents the author from dismissing an entire review because some claims are wrong. Second, it prevents valid criticism from being buried inside a poor review.

Step 3. Classify each claim

Each atomic claim should be classified into one of six categories.

Table 2. Classification categories for reviewer claims.

Category	Definition	Typical response
Supported	The claim is clearly valid based on the manuscript.	Revise and thank the reviewer.
Partially supported	The claim has some basis but is vague, incomplete, or overstated.	Clarify, expand, or relocate information.
Contradicted	The claim says something is missing or incorrect, although the manuscript contains the relevant information.	Point to manuscript evidence; improve visibility if needed.
Unsupported/generic	The claim is too general to verify.	Make limited revision or request specificity.
Scope mismatch	The claim asks for work outside the manuscript's stated scope.	Explain the scope and add a limitation if appropriate.
Not assessable	The claim cannot be evaluated from the manuscript alone.	Ask the editor for guidance or clarification.

This classification shifts the response from emotion to evidence.

Step 4. Map claims to manuscript evidence

The author should create an evidence table before drafting the rebuttal.

Table 3. Claim-evidence mapping template.

Claim ID	Reviewer claim	Manuscript evidence	Classification	Planned action
C1	Dataset source is unclear.	Section X states the validation data source.	Contradicted / partially supported	Cite Section X; add more detail if useful.
C2	Hardware is missing.	Section Y lists GPU, CPU, and frameworks.	Contradicted	Point to existing text; add runtime benchmark if available.
C3	Limitations are insufficient.	Section Z discusses limitations briefly.	Partially supported	Expand limitations and mitigation strategies.

This table is the core of the protocol. It can be used internally for revision, included in a response letter, or sent to the editor if escalation is necessary.

Step 5. Select a response strategy

Different claim categories require different response strategies.

Table 4. Response strategy by claim type.

Claim type	Manuscript action	Response-letter action
Supported	Revise.	Thank the reviewer and describe the change.
Partially supported	Clarify, expand, or relocate.	Acknowledge the issue but correct overstatement.
Contradicted	Improve visibility if needed.	Politely point to existing manuscript evidence.
Unsupported/generic	Make limited clarity revision.	State that no specific missing element was identified.
Scope mismatch	Add limitation or future work note.	Explain why the request is outside scope.
Repeated severe contradiction	Prepare evidence table.	Request editorial consideration.

The strongest author response is not necessarily the most aggressive one. Often the best response is to say that the information was already present, while also improving its visibility.

Step 6. Escalate only when the pattern is serious

Editorial escalation should be reserved for serious cases. A single mistaken reviewer comment usually does not justify escalation. Repeated false missing-content claims may.

Escalation is appropriate when:

- multiple substantive claims are contradicted by the manuscript,
- the review recommendation depends on those contradicted claims,
- the review appears generic or disconnected from manuscript content,
- the author can provide a clear claim-evidence table,
- ordinary revision would require accepting false premises.

The escalation should focus on reliability, not accusation.

A weak escalation says:

This review was obviously written by AI.

A stronger escalation says:

Several claims in this review appear inconsistent with explicit manuscript content. The concern is not that the review is negative, but that multiple claims used to justify the recommendation are contradicted by the submitted manuscript.

5. Case Illustration

5.1 Case context

The protocol is illustrated using an anonymized case from an actual submission to a Korea Citation Index (KCI) candidate journal. The manuscript was a technical paper in applied artificial intelligence. Three formal reviews were received. One review was suspected to be LLM-written or LLM-assisted because it contained fluent, generic, and hallucination-like comments that were insufficiently grounded in the manuscript.

To reduce identifiability, the journal name, manuscript title, detailed application domain, system name, author identity in the original submission, reviewer identities, editorial metadata, submission dates, publication date, volume, issue, page information, and bibliographic links are omitted.

Selected technical metadata are retained because they are necessary to illustrate the audit. The manuscript contained information about validation data, hardware, software frameworks, model configuration, model size, parameter count, algorithmic procedure, and limitations. These details included an NVIDIA RTX 3080 GPU, an Intel i9-13900K CPU, Python, Torch, OpenCV, a 169.8 MB model weight file, and approximately 44.4 million parameters.

5.2 Claim-level audit of the problematic review

Table 5. Audit of hallucination-like claims in the problematic review.

Reviewer claim	Manuscript evidence	Classification	Response strategy
The validation data source or characteristics were not clearly described.	The manuscript described the validation data source and test-case types.	Contradicted / partially supported	Point to existing statement; add sample size and selection criteria if available.
Hardware specifications were needed for reproducibility.	The manuscript listed GPU, CPU, CPU/GPU usage, Python, Torch, and OpenCV.	Contradicted	Point to existing setup; add runtime benchmarks if possible.
Model parameters should be described.	The manuscript reported model configuration, model weight size, and parameter count.	Contradicted	Point to existing information; clarify relation to system complexity.
Methodological explanation was insufficient.	The manuscript described the algorithmic pipeline, update rules, pseudocode, and output generation.	Partially supported but overstated	Improve organization and visibility; expand implementation details.
Comparison with previous methods should be clearer.	Related work and a limited comparison were present, but empirical comparison could be expanded.	Partially supported	Add clearer comparison table or explain constraints.
Limitations should be discussed.	Limitations were present but mitigation discussion was limited.	Partially supported	Expand limitations and mitigation strategies.
Typographical errors were present.	Minor typographical errors existed.	Supported	Proofread and correct.

This table shows why the problem was not ordinary negative feedback. Several criticisms were valid or partially valid. However, the most serious claims were false missing-content claims: the review requested information as if it were absent, although the manuscript already contained it.

5.3 Response strategy used in the case

The appropriate response was not to accuse the reviewer of AI use. Instead, the author reported the grounding failures to the editor with specific evidence. The report focused on contradictions between review claims and manuscript content.

Following editorial consideration, the problematic review was not relied upon in the final assessment, and the manuscript was subsequently accepted and published.

This outcome illustrates the value of claim-level audit. A general complaint that “the review seems AI-generated” is weak. A table showing that multiple claims are contradicted by explicit manuscript content is stronger, more professional, and more useful to the editor.

6. Response Templates

6.1 When the reviewer is correct

Thank you for this helpful comment. The manuscript has been revised to address this point. Specifically, [describe revision], and the relevant change appears in [section/table/page].

6.2 When the reviewer is partially correct

Thank you for raising this issue. The manuscript did include this information, but it may not have been sufficiently visible. The relevant section has been revised to clarify the point and improve readability.

6.3 When the reviewer claims something is missing but it is present

Thank you for the comment. The requested information was already included in [section]. To avoid ambiguity, this information has been made more prominent and expanded in the revised manuscript.

6.4 When the reviewer comment is generic

Thank you for this suggestion. Because the comment did not identify a specific section or missing element, the manuscript was reviewed for clarity in the relevant areas, and [section] was revised accordingly.

6.5 When the reviewer asks for work outside scope

This is a valuable suggestion. However, the requested analysis is outside the scope of the present manuscript because [reason]. The point has been added to the limitations and future work section.

6.6 When multiple review claims are contradicted by the manuscript

The author respectfully requests editorial consideration of several claims in this review. The concern is not disagreement with negative criticism, but that multiple substantive claims appear inconsistent with explicit content already present in the manuscript. A claim-level audit table is provided below to map each disputed claim to the corresponding manuscript evidence.

7. Editorial Escalation Protocol

7.1 When escalation is justified

Not every flawed review should be escalated. Escalation is justified when the review itself becomes unreliable as a basis for editorial decision-making.

Examples include:

- repeated claims that present information is missing,
- methodological criticisms based on a misread pipeline,
- low scores justified by contradicted claims,
- generic comments with no manuscript-specific evidence,
- suspected AI-generated text combined with factual contradiction.

The final condition is important. Suspicion of AI use alone is not enough. The evidence should be grounded in contradictions.

7.2 How to write the escalation letter

A strong escalation letter should include:

1. a respectful opening,
2. a statement that valid criticisms have been addressed,
3. a clear distinction between disagreement and contradiction,
4. a claim-level table,
5. a request for editorial consideration.

Suggested wording:

The concern is not that the review is negative. The concern is that several claims in the review are contradicted by explicit content in the submitted manuscript. Because these claims appear to form part of the basis for the recommendation, the author respectfully asks the editor to consider whether this review should be relied upon in the final decision.

This keeps the issue professional and evidence-based.

7.3 What not to write

Authors should avoid language such as:

- “The reviewer clearly did not read the paper.”
- “This was obviously written by ChatGPT.”
- “The review is nonsense.”
- “The reviewer is unethical.”

Even if the author suspects these things, such statements are unlikely to help. The stronger move is to present evidence.

8. Discussion

8.1 Novelty of the protocol

The novelty of the protocol is not that authors should respond to reviewer comments with evidence. That is already standard advice. The novelty lies in adapting response practice to the LLM era, where a formal review may be fluent, plausible, and potentially AI-written while containing hallucination-like claims.

Existing guidance emphasizes polite point-by-point responses. This paper adds a prior step: claim-level verification. Before writing the rebuttal, the author should determine whether each claim is supported, partially supported, contradicted, generic, scope-mismatched, or not assessable.

This makes the protocol especially useful for suspected LLM-written hallucinated reviews. It helps authors avoid two mistakes: accepting false premises, and making unsupported accusations of AI use.

8.2 The protocol protects valid criticism

The protocol does not encourage authors to dismiss inconvenient reviews. In the case illustration, some reviewer comments were valid: typographical errors existed, quantitative evaluation could be strengthened, and limitations could be expanded. The protocol preserves those valid criticisms while identifying false missing-content claims.

This is important. A hallucination-like review may still contain useful comments. Conversely, a human-written review may still contain ungrounded claims. The protocol evaluates claims, not authors.

8.3 The protocol is better than style-based AI accusation

Style-based AI suspicion is weak. A review may be generic because it was generated by an LLM, because the reviewer used a template, or because the reviewer wrote quickly. The author usually cannot prove the source.

Claim-level contradiction is stronger. It does not depend on detecting AI. It depends on comparing review claims to manuscript evidence.

Thus, the protocol reframes the issue:

*Do not argue that the review looks like AI.
Show that the review is contradicted by the manuscript.*

8.4 Editorial implications

Editors can use the same protocol. When an author reports that a review may be AI-generated or hallucinated, the editor can request a claim-level audit rather than adjudicating a vague complaint.

A useful editorial assessment may ask:

- Which claims are disputed?
- Where does the manuscript address them?
- Are the disputed claims contradicted or only partially supported?

- Would revision solve the issue?
- Should the review be discounted, supplemented, or replaced?

This is consistent with the broader direction of publisher policies: AI use in peer review is not merely a tool issue, but a matter of confidentiality, accountability, and reliability [2,3].

8.5 The protocol does not solve every problem

The protocol cannot prove AI use. It cannot force editors to reconsider a decision. It cannot remove subjectivity from peer review. It also cannot prevent authors from misclassifying valid criticism as contradicted.

For this reason, the protocol should be used conservatively. Authors should revise where criticism is valid, acknowledge ambiguity where the manuscript is unclear, and escalate only when repeated contradictions make the review unreliable.

9. Limitations

This paper has several limitations.

First, the protocol is illustrated with a single anonymized case. The paper does not estimate the prevalence of suspected LLM-written hallucinated peer review.

Second, the case does not prove that the problematic review was generated by an LLM. The protocol is intentionally designed to avoid relying on AI-detection claims.

Third, the protocol depends on author judgment. Authors may misclassify reviewer comments, especially when the review is harsh or emotionally frustrating. For this reason, the protocol emphasizes evidence mapping and respectful language.

Fourth, the protocol may not be accepted by every editor or journal. Some editorial systems provide limited opportunities for authors to contest reviews.

Fifth, anonymization limits external verifiability. The case details are generalized to protect the journal, reviewers, and editorial process. Technical metadata are retained only where necessary to explain the audit.

10. Conclusion

Authors should not respond to suspected LLM-written hallucinated peer reviews by arguing that the review “looks like AI.” Such claims are difficult to prove and may distract from the editorial issue.

A stronger response is to audit the review at the claim level. The author should identify atomic claims, classify each claim by its relationship to the manuscript, map disputed claims to evidence, revise where criticism is valid, and escalate only when repeated contradictions make the review unreliable.

The central lesson is simple:

Do not accuse first. Audit first.

In the LLM era, peer review integrity cannot depend only on detecting whether AI was used. It must also depend on whether review claims are grounded, verifiable, and responsibly evaluated. A claim-level audit protocol gives authors a practical way to replace speculation with evidence.

Ethics and Confidentiality Statement

This paper analyzes anonymized peer-review material from a real journal submission. The journal was a Korea Citation Index (KCI) candidate journal at the time of the review process. This broad indexing status is disclosed only to clarify that the case arose from formal scholarly peer review. The journal title, manuscript title, detailed application domain, system name, author identity in the original submission, reviewer identities, editorial metadata, submission dates, publication date, volume, issue, page information, and bibliographic links are omitted.

Selected technical metadata are retained where necessary to illustrate the audit protocol. The purpose of the paper is not to identify or accuse an individual reviewer, editor, or journal, but to propose a method for responding to suspected LLM-written hallucinated peer-review claims.

AI-Use Statement

AI tools were used to assist drafting, organization, and language editing of this manuscript. The human author remains solely responsible for factual accuracy, interpretation, argumentation, and conclusions.

References

- [1] Committee on Publication Ethics (COPE), “Editors suspect that reviewers are using artificial intelligence.” Accessed May 4, 2026. URL: <https://publicationethics.org/guidance/case/editors-suspect-reviewers-are-using-artificial-intelligence>
- [2] Nature Portfolio, “Artificial Intelligence (AI), Editorial Policies.” Accessed May 4, 2026. URL: <https://www.nature.com/nature-portfolio/editorial-policies/ai>
- [3] SAGE, “Artificial Intelligence Policy.” Accessed May 4, 2026. URL: <https://www.sagepub.com/journals/publication-ethics-policies/artificial-intelligence-policy>
- [4] N. Lo Vecchio, “Personal experience with AI-generated peer reviews: a case study,” *Research Integrity and Peer Review*, 2025. URL: <https://link.springer.com/article/10.1186/s41073-025-00161-3>
- [5] V. S. Rao, A. Kumar, H. Lakkaraju, and N. B. Shah, “Detecting LLM-generated peer reviews,” *PLOS ONE*, 20(9): e0331871, 2025. DOI: 10.1371/journal.pone.0331871. URL: <https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0331871>
- [6] Taylor & Francis Author Services, “How to respond to reviewer comments.” Accessed May 4, 2026. URL: <https://authorservices.taylorandfrancis.com/blog/peer-review/how-to-respond-to-reviewer-comments/>
- [7] P. S. Hiemstra, “How to write a response to the reviewers of your manuscript,” *Breathe*, 2018. URL: <https://pmc.ncbi.nlm.nih.gov/articles/PMC6269172/>
- [8] Nature Nanotechnology, “Peer review in the time of artificial intelligence,” *Nature Nanotechnology*, 21, 479, 2026. DOI: 10.1038/s41565-026-02177-2. URL: <https://www.nature.com/articles/s41565-026-02177-2>
- [9] B. Ban, “Can LLMs Review Scientific Papers?” preprint on Researchgate. <https://doi.org/10.13140/RG.2.2.34160.83209>

Appendix A. Claim-Level Audit Worksheet

Table A1. Claim-level audit worksheet.

Claim ID	Reviewer	Reviewer comment	Atomic claim	Manuscript evidence	Classification	Action	Response text
C1	A/B/C				Supported / Partially supported / Contradicted / Generic / Scope mismatch / Not assessable	Revise / Clarify / Relocate / Disagree / Escalate	

Appendix B. Response Strategy Matrix

Table B1. Response strategy matrix.

Classification	Manuscript action	Response-letter action	Editorial action
Supported	Revise	Thank and describe change	None
Partially supported	Clarify or expand	Acknowledge and explain improvement	None
Contradicted	Improve visibility if needed	Point to existing manuscript evidence	Consider if repeated or severe
Unsupported/generic	Limited clarity revision	State that no specific issue was identified	Ask editor for specificity if necessary
Scope mismatch	Add limitation or future work	Explain scope	None unless demand is unreasonable
Not assessable	State uncertainty	Ask for clarification	Editor guidance may be needed

Appendix C. Editorial Escalation Template

Dear Editor,

Thank you for managing the review process for our manuscript.

We have carefully reviewed all reviewer comments and have revised the manuscript where the comments identify valid weaknesses. However, we respectfully request editorial consideration of several claims in [Review X]. Our concern is not that the review is negative, but that multiple claims appear to be contradicted by explicit content already present in the manuscript.

To make this concern easy to evaluate, we provide below a claim-level audit table. Each disputed review claim is mapped to the corresponding manuscript evidence and classified as supported, partially supported, contradicted, unsupported/generic, scope-mismatched, or not assessable.

[Insert table]

We remain willing to revise the manuscript further where clarification is needed. However, we believe that the claims marked as contradicted should not be used as a basis for editorial decision-making without additional review.

Sincerely,

[Author]

Appendix D. Author Response Templates

Template 1. Valid criticism

Thank you for this helpful comment. We agree that this part of the manuscript required improvement. The manuscript has been revised accordingly in [Section X], where we now [describe revision].

Template 2. Partially valid criticism

Thank you for raising this issue. The manuscript did include this information, but we agree that it was not sufficiently visible. We have revised [Section X] to clarify the point and make the relevant information easier to locate.

Template 3. False missing-content claim

Thank you for the comment. The requested information was already included in [Section X]. To avoid ambiguity, we have made this information more prominent in the revised manuscript and have added additional detail in [Section Y].

Template 4. Generic comment

Thank you for this suggestion. Because the comment did not identify a specific section or missing element, we reviewed the manuscript for clarity in the relevant areas and revised [Section X] accordingly.

Template 5. Scope mismatch

We appreciate this suggestion. The requested analysis is valuable but outside the scope of the present manuscript because [reason]. We have added this issue to the limitations and future work section.