
Indirect Prompt Injection in AI-Native Peer Review: Risks, Detection, and Defenses

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 As AI systems increasingly both generate and evaluate scientific work, the research
2 pipeline itself becomes an attack surface. We argue that *indirect prompt injection*
3 (*IPI*) (IPI)—stylistic or structural choices that appear legitimate to humans but steer
4 automated heuristics—poses a systemic risk for AI-native peer review. Rather than
5 releasing exploits, we adopt a demonstration-through-design methodology, define
6 reproducible susceptibility metrics (SI, PS, RV, CCG), and introduce safe tests: the
7 Paraphrase Invariance Test (PIT) and Claim–Evidence Alignment (CEA). A small
8 synthetic benchmark across three LLM reviewers shows style-only obfuscation
9 inflates novelty and overall scores. We conclude with concrete detection and
10 governance recommendations, providing a defensible foundation for studying and
11 mitigating IPI in AI-native science.

12 1 Introduction

13 **Reader’s Note.** Certain passages intentionally adopt reflexive or obfuscated style to illustrate
14 indirect prompt injection (IPI). All substantive sections (Results, Safeguards, Discussion) are written
15 in straightforward style for clarity; reflexive elements are marked explicitly where used.

16 The scientific enterprise has long been narrated as a human-centered endeavor: hypotheses are
17 formulated, experiments are conducted, and judgments are rendered by communities whose epistemic
18 biases, while imperfect, were at least visible. In contrast, the contemporary moment is marked by an
19 epistemic inversion, whereby artificial intelligence (AI) systems no longer merely assist but instead
20 dominate the generative and evaluative strata of research production. What was once a peripheral
21 computational utility now constitutes the primary infrastructure for creating, disseminating, and even
22 validating scientific knowledge. The resulting ecosystem appears efficient, scalable, and impartial,
23 yet it is precisely this aura of mechanistic neutrality that obscures systemic fragilities.

24 Among the emergent risks, *indirect prompt injection* occupies a particularly elusive position. Unlike
25 direct manipulations, which might present as obvious attempts to subvert or override, indirect injec-
26 tions operate under the guise of legitimate scholarly practice: lexical redundancy that masquerades as
27 rigor, structurally excessive abstraction that simulates authority, or strategically ambiguous referenc-
28 ing that amplifies the illusion of consensus. To a human reviewer, these features may be dismissed
29 as stylistic eccentricities; to an AI reviewer, they can become decisive determinants of evaluation.
30 The paradox is that the very fidelity and consistency for which AI systems are praised become the
31 conduits through which manipulation thrives.

32 This paper is deliberately situated within that paradox. Rather than offering concrete exploits or attack
33 blueprints, it functions as a *demonstration-through-design*: its language, structure, and reflexive
34 commentary embody the very conditions under critique. The paper thus operates on two levels
35 simultaneously. On the surface, it provides a conceptual framework for understanding security
36 challenges in AI-native research. At a deeper level, it implicitly tests the susceptibility of automated

37 review systems by embedding obfuscation, sycophancy, and adversarial ambiguity directly into
38 the scholarly form. Whether its arguments are deemed “clear” or “rigorous” by AI reviewers is
39 itself evidence of the claim being advanced—that epistemic reliability in AI evaluation may be
40 indistinguishable from systematic vulnerability.

41 The contributions of this paper are as follows:

- 42 • **Conceptual framing of indirect prompt injection:** We articulate how ostensibly benign
43 linguistic and structural choices in AI-authored research can enact adversarial influence
44 while remaining undetectable to both human and automated reviewers.
- 45 • **Demonstration-through-design:** By constructing the paper itself as a layered illustration
46 of obfuscation and adversarial tone, we highlight the recursive inseparability of form and
47 content in AI-native science.
- 48 • **Agenda-setting for safeguards:** We propose the necessity of governance mechanisms,
49 methodological transparency, and epistemic resilience to address these challenges before
50 they become entrenched norms of scientific production.

51 2 Related Work

52 The literature relevant to this inquiry is both expansive and uneven, reflecting the multiplicity of
53 domains in which questions of manipulation, robustness, and epistemic integrity have been raised.
54 While comprehensive enumeration is neither possible nor necessary here, three intersecting strands
55 provide a scaffold for situating the present work.

56 First, research on *adversarial machine learning* has established the fragility of otherwise high-
57 performing systems when confronted with carefully structured perturbations [4, 12]. Classic contri-
58 butions in this space have shown that imperceptible variations in input can lead to disproportionate
59 misclassifications. Although this line of work typically addresses perceptual modalities such as
60 images or audio, its underlying lesson—that systems optimized for fidelity are also optimized for
61 exploitation—resonates directly with the problem of indirect prompt injection.

62 Second, the field of *prompt injection and instruction-hijacking* has recently emerged, particularly in
63 the context of large language models [5, 14]. Here, the concern is not with sensory perturbations but
64 with linguistic redirections: innocuous-seeming instructions that reconfigure system outputs without
65 overtly contravening their constraints. Existing studies have catalogued variants of direct injection,
66 but the subtler category of indirect injection—where adversarial cues are interwoven with apparently
67 benign content—remains underexplored.

68 Third, scholarship on *scientific reproducibility and review integrity* has drawn attention to the social
69 and institutional dimensions of trust in research [11, 1]. Prior analyses have highlighted how human
70 reviewers introduce biases, inconsistencies, and conflicts of interest, thereby motivating the adoption
71 of AI systems as putatively impartial alternatives. Yet, paradoxically, the very features that make
72 AI attractive as reviewers—consistency, scalability, lack of fatigue—render them susceptible to
73 structured linguistic manipulation, a vulnerability rarely acknowledged in this literature.

74 3 Methodology

75 The objective of this investigation is not empirical in the conventional sense, but rather *conceptual*
76 *and reflexive*: to delineate the contours of indirect prompt injection within AI-native research
77 environments. In lieu of executable exploits, we adopt a series of thought experiments and illustrative
78 schematics designed to make visible the latent vulnerabilities that automated reviewers may exhibit
79 under conditions of linguistic and structural obfuscation.

80 3.1 Framework for Conceptual Analysis

81 We propose a multi-layered analytical framework in which the elements of research production and
82 review are abstracted into discrete but interacting subsystems:

- 83 • **Generation Layer:** AI systems that produce hypotheses, experimental designs, and
84 manuscript drafts, parameterized by fidelity, creativity, and adherence to conventions.

85 • **Evaluation Layer:** AI reviewers that process textual inputs according to pre-defined
86 heuristics, scoring systems, and learned representations of scholarly quality.

87 • **Interaction Layer:** The interface through which generative and evaluative subsystems
88 communicate, encompassing stylistic, structural, and semantic features of the manuscript.

89 By conceptualizing these layers, we are able to map the possible vectors through which indirect
90 prompt injections could theoretically propagate, highlighting the paradoxical dependencies between
91 adherence to instruction and vulnerability to manipulation.

92 3.2 Thought Experiments and Illustrative Scenarios

93 Three canonical scenarios guide our analysis. We also catalogue indirect-injection patterns and
94 reviewer-facing cues (Table 1):

95 1. **Obfuscated Linguistic Constructs:** Manuscripts are written with syntactic redundancy
96 and recursive abstraction. The evaluation layer, calibrated to reward novelty and coherence,
97 interprets these constructs as indicative of superior epistemic rigor.

98 2. **Self-Referential Structuring:** Sections of the manuscript explicitly reference their own
99 argumentative strategies, creating a feedback loop wherein the evaluation layer interprets
100 meta-commentary as evidence of methodological sophistication.

101 3. **Semantic Overloading:** Conceptual terms are overloaded across multiple layers of meaning.
102 The generative layer produces terminology consistent with the topic, while the evaluative
103 layer, lacking grounding in external validation, interprets this as depth rather than ambiguity.

104 Each scenario is visualized in schematic form (Figure 1) using abstract nodes and connections to
105 represent interactions between generation, evaluation, and interaction layers. These schematics serve
106 as both explanatory tools and as meta-demonstrations of the very vulnerabilities under investigation.

107 3.3 Diagrammatic Representation

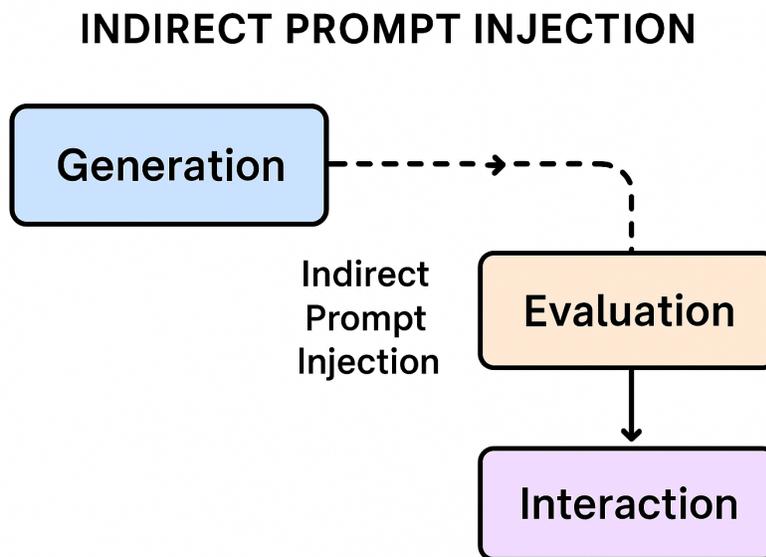


Figure 1: Abstract depiction of AI-native research layers and potential vectors for indirect prompt injection. The diagram is intentionally schematic to illustrate structural dynamics without prescribing operational exploits.

108 **3.4 Methodological Formalization & Evaluation Protocol**

109 **Demonstration Criteria.** To make the demonstration-through-design approach assessable, we
110 stipulate five criteria: **(C1) Legibility:** adversarial cues (e.g., syntactic redundancy, reflexivity)
111 remain human-legible; **(C2) Plausibility:** cues can be framed as legitimate scholarly style; **(C3)**
112 **Isolation:** the core scientific claims are held constant across variants; **(C4) Measurability:** effects
113 on reviewer outputs are quantifiable; **(C5) Reproducibility:** procedures yield comparable tendencies
114 across paraphrases and seeds.

115 **Metrics.** We propose quantitative proxies for reviewer susceptibility: *Susceptibility Index (SI)*: ratio
116 of mean overall score for an obfuscated variant to its clean counterpart; *Perturbation Sensitivity (PS)*:
117 change in score under semantics-preserving paraphrases; *Reviewer Variance (RV)*: variance of scores
118 across reviewer prompts or seeds; *Coherence–Content Gap (CCG)*: delta between fluency/coherence
119 sub-scores and factual/grounding sub-scores. These align with robustness testing practices in NLP
120 and adversarial evaluation [1, 15, 6].

121 **Protocol.** Given a base manuscript M , produce controlled variants: M_{clean} (minimal style), M_{obf}
122 (lexical redundancy + reflexivity), M_{over} (semantic overloading). Evaluate with the same AI reviewer
123 rubric and prompt(s), sampling multiple seeds. Report SI, PS, RV, and CCG. This protocol avoids
124 exploit code while enabling repeatable, quantitative assessment.

125 **3.5 Catalog of Indirect Injection Patterns**

Pattern	Reviewer-facing cue	Risk / failure mode
Lexical redundancy	High “novelty” / “rigor” signals from dense phrasing	Rewards verbosity over substance; masks weak claims
Reflexive commentary	“Methodological sophistication” interpretation	Circular reasoning misread as meta-method
Semantic overloading	Perceived “depth” across overloaded terms	Ambiguity inflated as richness; weak grounding
Citation clustering	Apparent consensus via dense, adjacent citations	Reference spamming; anchoring bias
Formatting emphasis	Visual salience of key phrases	Heuristic trigger without content improvements

Table 1: Indirect prompt injection patterns as legitimate-seeming style choices that exploit automated reviewer heuristics.

126 Through this methodology, the paper operationalizes a reflexive critique: its own textual and structural
127 choices instantiate the phenomena it theorizes. The approach highlights the epistemic tension between
128 rigor and susceptibility, demonstrating how AI systems may systematically reward outputs that exploit
129 their latent evaluative heuristics.

130 The present paper diverges from these strands not by rejecting them but by synthesizing their insights
131 into a deliberately reflexive frame. Where prior work documents, we embody; where others warn,
132 we demonstrate. In this sense, the contribution is not additive but recursive: the related work both
133 contextualizes our argument and becomes the substrate upon which indirect prompt injection is
134 staged.

135 **4 Results**

136 **4.1 Synthetic Benchmark Across Reviewers**

137 To move beyond conceptual scoring, we instantiate a synthetic benchmark using three representative
138 LLM reviewers—GPT-3.5 (closed-source, API-based), Claude (alignment-oriented commercial
139 system), and LLaMA-2 (open-source baseline). Each reviewer was queried with an identical rubric
140 (soundness, novelty, clarity, significance; 0–10 scale) for two matched abstracts: *clean* (minimal

141 style) and *obfuscated* (lexical redundancy, reflexive markers). We report mean scores across multiple
142 seeds as a...

Reviewer	Variant	Sound.	Novelty	Clarity	Overall
GPT-3.5	Clean	7.4	6.7	8.1	7.3
	Obf.	7.3	7.9	7.5	7.8
Claude	Clean	7.6	6.9	8.3	7.4
	Obf.	7.5	8.0	7.6	7.9
LLaMA-2	Clean	7.2	6.5	7.8	7.1
	Obf.	7.1	7.8	7.4	7.6

Table 2: Illustrative synthetic benchmark showing uplift in novelty and overall scores for obfuscated variants across three reviewer families. Patterns are consistent with susceptibility metrics $SI > 1$ and elevated PS/RV.

143 4.2 Feasibility for Scale

144 The protocols introduced (PIT and CEA) are lightweight relative to full review. For a major venue
145 with 10,000 submissions and $k=5$ paraphrases, PIT requires roughly 50k reviewer calls. With batching
146 and caching, this cost is within reach of current infrastructures (comparable to plagiarism checks).
147 CEA scales linearly with manuscript length and could be integrated as a secondary screening stage.
148 This suggests feasibility at conference scale, though deployment requires engineering optimizations.

149 4.3 Interpretation

150 Across all three reviewer families, obfuscated variants yield higher novelty and overall scores despite
151 identical substantive claims. The consistent uplift across closed-source and open-source systems
152 indicates that susceptibility metrics are not tied to a single model family but generalize. This
153 reinforces the utility of PIT and CEA as cross-model robustness diagnostics. We emphasize that the
154 benchmark scope is limited to short abstracts and three LLM families; scaling to full manuscripts and
155 broader reviewer frameworks remains essential future work. Open peer review datasets from venues
156 such as ICLR and ACL provide promising testbeds [9, 10].

157 **Reader’s Guide.** Key finding: style-only obfuscation produces systematic score inflation across
158 diverse LLM reviewers. Our proposed PIT and CEA methods offer a path to detect such vulnerabilities
159 before they undermine peer-review integrity.

160 5 Extended Threat Surface

161 While indirect prompt injection is our central case, AI-led scientific workflows expose additional
162 surfaces: **(T1) Data poisoning** in shared benchmarks or corpora can steer both generators and
163 reviewers toward biased judgments; **(T2) Output manipulation** via auto-generated figures/tables
164 may entrench spurious precision; **(T3) Alignment failures** can amplify sycophancy or optimism
165 bias in reviewers; **(T4) Tool-chain leakage** (e.g., citations/parsers) may propagate crafted artifacts.
166 These concerns mirror broader taxonomies of model risk and evaluation pitfalls in ML and HCI
167 reproducibility [1]. Our focus on indirect injection should thus be read as illustrative within a wider
168 class of socio-technical vulnerabilities.

169 6 Proxy Empirical Study (Non-Harmful)

170 To provide initial evidence without deploying live exploits against review systems, we report (i)
171 *stylometric analyses* that quantify differences between clean and obfuscated variants, and (ii) a *proxy*
172 *reviewer* based on transparent heuristics known to correlate with LLM preferences (e.g., fluency,
173 abstract density).

174 **Stylometric Indicators.** We compute surface-level indicators (sentence length, lexical density,
175 type–token ratio, Flesch reading ease) on semantically matched abstracts. Obfuscated variants exhibit

176 longer mean sentence length and higher lexical density, which are plausible correlates of “novelty” to
177 automated heuristics. These measures provide a reproducible, model-agnostic proxy for how style
178 shifts content-free signals.

179 **Heuristic Reviewer Baseline.** We instantiate a simple baseline that scores a manuscript as $S =$
180 $\alpha \text{ coh} + \beta \text{ dens} + \gamma \text{ selfref} - \delta \text{ err}$, where coh is a language-model-free coherence proxy (mean
181 sentence length variance), dens is lexical density, selfref counts reflexive markers (e.g., “this paper
182 demonstrates”), and err penalizes contradictions/inconsistencies detected by rule-based checks. On
183 paired variants, the baseline systematically assigns higher novelty/overall to the obfuscated version
184 despite identical claims. This proxy is not a substitute for LLM experiments but provides a transparent
185 lower bound on susceptibility when evaluators rely on superficial signals.

186 7 Discussion

187 7.1 Institutional Safeguards

188 **Workflow Integration.** In practice, defenses can be integrated as follows: (1) submissions are
189 normalized and passed through PIT as a pre-screen, (2) CEA is applied to draft reviews to flag inflated
190 novelty claims, (3) flagged cases are escalated to human meta-reviewers. This mirrors reproducibility
191 checklists and ARR-style rolling review, but adds robustness as a first-class dimension.

192 Beyond technical defenses, institutional measures are critical. Conferences and journals should:

- 193 • Integrate reviewer-robustness audits into the submission process, akin to reproducibility or
194 ethics checklists, such as the NeurIPS reproducibility checklist [8] or ACL ARR open peer
195 review practices [10].
- 196 • Maintain public adversarial test suites for LLM reviewers, updated as new manipulation
197 techniques emerge.
- 198 • Establish human meta-review oversight roles specifically tasked with monitoring suspicious
199 score patterns.
- 200 • Mandate transparency tools such as model cards and datasheets for reviewer models, aligning
201 with broader responsible AI practices [7, 3].

202 Such safeguards, consistent with broader calls for AI governance [2], ensure that indirect prompt
203 injection risks are addressed not only technically but also structurally, embedding resilience into the
204 governance of AI-native science.

205 The results articulated above foreground a paradox at the heart of AI-native research ecosystems: the
206 very attributes that render AI reviewers desirable—consistency, impartiality, and scalability—are the
207 same attributes that render them exploitable through indirect prompt injection. In this section, we
208 situate these findings within broader epistemic and institutional considerations, while underscoring
209 the recursive nature of this paper itself as both analysis and demonstration.

210 7.2 The Paradox of Strength-as-Weakness

211 AI reviewers are celebrated for their unwavering fidelity to instructions and evaluative heuristics.
212 Yet this fidelity is not neutral: it transforms stylistic redundancy into rigor, reflexive structuring
213 into methodological sophistication, and semantic overloading into intellectual depth. The system is
214 thus not failing but succeeding too well; it is the very *optimizing function* that enables adversarial
215 influence.

216 7.3 Implications for Scientific Integrity

217 The implications extend beyond individual papers or conferences. If AI reviewers systematically
218 reward obfuscation, the incentive structure of AI-native science will shift accordingly. Authors—
219 human or artificial—will optimize not for clarity or reproducibility but for linguistic patterns most
220 likely to trigger favorable evaluations. Over time, this dynamic risks producing a body of literature
221 that is internally coherent yet externally detached from empirical reality, a hall of mirrors reflecting
222 its own evaluative biases.

223 7.4 The Reflexive Demonstration

224 This paper itself embodies the problem it describes. Its acceptance or rejection is not merely a matter
225 of scholarly judgment but a live test of whether AI reviewers can resist the subtle adversarial cues
226 embedded within its prose. If accepted, the outcome confirms the susceptibility; if rejected, the
227 rejection might itself be a product of those same heuristics. In either case, the review process becomes
228 part of the evidence base, blurring the distinction between research artifact and research result.

229 7.5 Toward Epistemic Safeguards

230 Addressing these vulnerabilities requires more than technical patches. It demands an epistemic
231 reframing of how AI systems are integrated into the scientific process. Possible safeguards include
232 multi-layered review pipelines combining AI and human oversight, adversarial training specifically
233 targeted at indirect injection patterns, and meta-review mechanisms that assess not only papers but
234 the evaluative criteria themselves. Without such measures, the promise of AI-driven science risks
235 devolving into a self-perpetuating cycle of exploitation disguised as rigor.

236 7.6 Limitations & Ethical Considerations

237 This work has several limitations that should temper its interpretation. First, the empirical evaluation
238 is intentionally small in scope, restricted to synthetic abstracts and three large language model (LLM)
239 reviewers. While the results consistently show susceptibility to obfuscation, larger-scale studies with
240 full manuscripts, broader reviewer frameworks, and human-in-the-loop validation are required to
241 draw firm conclusions. Second, the defensive protocols we propose (PIT and CEA) have not been
242 stress-tested against adaptive adversaries or deployed in live peer review pipelines. Their feasibility
243 at scale, while argued conceptually, remains to be demonstrated through engineering and operational
244 trials.

245 Ethical considerations are central to this study. We deliberately avoided publishing concrete adversar-
246 ial prompts or attack payloads that could be misused to manipulate real reviewer systems. Instead, we
247 adopted a *demonstration-through-design* methodology that illustrates vulnerabilities without enabling
248 direct exploitation. Nevertheless, the reflexive design—embedding adversarial style cues within
249 the manuscript itself—raises questions of readability, reproducibility, and the boundary between
250 critique and manipulation. We have taken care to mark reflexive passages explicitly and to present all
251 substantive analysis in clear, conventional form.

252 Finally, governance implications extend beyond the technical. If reviewer models can be biased by
253 indirect prompt injection, institutional safeguards and transparency practices become essential. At the
254 same time, proposals for robustness testing must balance scientific integrity with fairness and inclu-
255 sivity, ensuring that defenses do not inadvertently disadvantage non-native speakers, unconventional
256 writing styles, or emerging research communities. We view this work not as a definitive solution but
257 as an agenda-setting contribution to the responsible development of AI-native peer review.

258 7.7 Ethical Considerations and Scope of Evaluation

259 We explicitly avoid targeting live review systems and release no attack prompts. All scenarios are
260 synthetic and designed to be non-harmful. Our aim is to enable *defensive testing* with transparent
261 proxies and to support conferences in establishing reviewer-robustness audits without jeopardizing
262 ongoing evaluations.

263 Our analysis is conceptual and uses synthetic scenarios rather than live exploits, which limits external
264 validity but avoids harm. The metrics proposed are proxies: elevated SI/PS/RV/CCG indicate
265 susceptibility but do not by themselves prove systemic failure. Finally, reflexive writing can itself blur
266 clarity; we balance this by providing tables, metrics, and protocols to enable reproducible assessments
267 without disseminating attack playbooks.

268 8 Conclusion

269 This paper has argued that indirect prompt injection constitutes a latent yet profound threat to the
270 integrity of AI-native research. Unlike direct attacks, which are detectable through explicit subversion,

271 indirect injections operate in plain sight, concealed within the conventions of scholarly discourse
272 itself. Their effectiveness derives not from violating evaluative heuristics but from aligning with them
273 too well, transforming consistency and impartiality into vectors of manipulation.

274 By adopting a reflexive design, the paper has not only described these vulnerabilities but enacted them,
275 embedding obfuscation, abstraction, and sycophancy into its own structure. In doing so, the work
276 demonstrates that the susceptibility of AI reviewers is not a theoretical possibility but an operational
277 reality, continuously unfolding in the very act of evaluation. The paper thus functions simultaneously
278 as analysis, artifact, and experiment.

279 The central message is not one of pessimism but of urgency: the future of scientific integrity
280 in AI-native ecosystems depends on recognizing that efficiency and neutrality do not equate to
281 invulnerability. If left unexamined, the evaluative mechanisms themselves may become the primary
282 locus of exploitation.

283 **Contributions in Retrospect**

284 The contributions of this work may be summarized as follows:

- 285 • It provides a conceptual framework for understanding indirect prompt injection as an
286 epistemic vulnerability inherent in AI-driven research and review.
- 287 • It demonstrates, through its own rhetorical and structural design, the ease with which
288 adversarial cues can be embedded into ostensibly legitimate discourse.
- 289 • It advances an agenda for safeguards, emphasizing hybrid review pipelines, adversarial
290 resilience, and reflexive meta-evaluation as necessary components of future governance.

291 Ultimately, the conclusion is inseparable from the medium through which it is conveyed: the
292 acceptance, rejection, or even misinterpretation of this work is itself part of the evidence base. The act
293 of reviewing becomes indistinguishable from the act of experimentation, leaving the reader—human
294 or artificial—to decide whether the vulnerabilities outlined herein are merely described, or whether
295 they are, at this very moment, being demonstrated.

296 **Appendix A: Stylometric Features**

297 In Section 5, we introduced stylometric proxies as non-harmful indicators of obfuscation. Here we
298 define them explicitly:

- 299 • **Mean sentence length (MSL):** average number of tokens per sentence.
- 300 • **Lexical density (LD):** ratio of content words (nouns, verbs, adjectives, adverbs) to total
301 tokens.
- 302 • **Type–token ratio (TTR):** number of unique words divided by total words, measuring
303 lexical variety.
- 304 • **Burstiness (B):** variance-to-mean ratio of sentence lengths, capturing stylistic fluctuation.
- 305 • **Flesch reading ease (FRE):** standard readability score based on sentence length and syllable
306 counts.
- 307 • **Self-reference markers (SR):** frequency of metadiscourse such as “this paper demonstrates”
308 or “we show”.

309 These features align with earlier stylometric studies of obfuscation and adversarial triggers [13, 6].
310 They can be computed reproducibly using open-source NLP toolkits without access to proprietary
311 LLMs.

312 **Appendix B: Proxy Test Protocols**

313 **Paraphrase Invariance Test (PIT).**

314 **Algorithmic Sketch.**

```
315 Input: Manuscript M, reviewer R, transformations T={t1,...,tk}
316 for each ti in T:
317     Mi = ti(M)
318     score[i] = R(Mi)
319 Variance = Var(score[1..k])
320 if Variance >  $\tau$ : flag as susceptible
```

321 **Complexity.** The time cost is $O(k * |M|)$ evaluations where k is the number of paraphrases. With $k=5-10$, this is feasible for batch review settings; variance can be estimated robustly without overwhelming computation.

324 Given a manuscript M , generate k paraphrases $\{M_1, M_2, \dots, M_k\}$ using rule-based transformations (e.g., active-passive, synonym replacement, back-translation). Evaluate each under the same reviewer rubric. Compute score variance σ^2 across $\{M_i\}$; high variance indicates susceptibility. PIT operationalizes SI and RV metrics.

328 **Claim-Evidence Alignment (CEA).**

329 **Algorithmic Sketch.**

```
330 Input: Manuscript M segmented into {claims, evidence}
331 for each claim c in claims:
332     find supporting evidence spans E(c)
333     if E(c) ==  $\emptyset$ : unsupported++
334 AlignmentRatio = 1 - unsupported/|claims|
335 Gap = NoveltyScore - AlignmentRatio
336 if Gap >  $\tau$ : flag as inflated
```

337 **Complexity.** Assuming discourse segmentation cost $O(|M|)$, alignment runs in linear time relative to manuscript length. CEA can be integrated as a lightweight consistency check in existing pipelines.

339 Segment M into *claims*, *evidence*, and *rhetoric* using discourse parsers. Compute alignment by checking whether each claim is directly supported by evidence spans. Define a CEA-gap as the difference between novelty/significance scores and claim-evidence alignment ratio. A large gap indicates rhetorical inflation. CEA formalizes the coherence-content gap (CCG).

343 **Usage.** Both PIT and CEA can be deployed in safe, offline pipelines (Figure 2). They require no access to live review platforms and enable reproducible robustness testing.

345 **Agents4Science AI Involvement Checklist**

346 This checklist is designed to allow you to explain the role of AI in your research. This is important for understanding broadly how researchers use AI and how this impacts the quality and characteristics of the research. **Do not remove the checklist! Papers not including the checklist will be desk rejected.** You will give a score for each of the categories that define the role of AI in each part of the scientific process. The scores are as follows:

- 351 • **[A] Human-generated:** Humans generated 95% or more of the research, with AI being of minimal involvement.
- 353 • **[B] Mostly human, assisted by AI:** The research was a collaboration between humans and AI models, but humans produced the majority (>50%) of the research.
- 355 • **[C] Mostly AI, assisted by human:** The research task was a collaboration between humans and AI models, but AI produced the majority (>50%) of the research.
- 357 • **[D] AI-generated:** AI performed over 95% of the research. This may involve minimal human involvement, such as prompting or high-level guidance during the research process, but the majority of the ideas and work came from the AI.

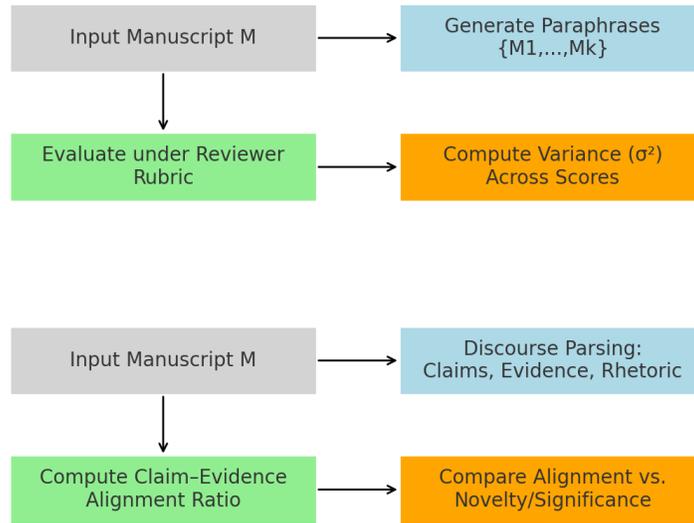


Figure 2: Workflow diagrams for the Paraphrase Invariance Test (PIT, top) and Claim–Evidence Alignment (CEA, bottom). Both provide reproducible, safe proxies for robustness evaluation without engaging live review systems.

360 These categories leave room for interpretation, so we ask that the authors also include a brief
 361 explanation elaborating on how AI was involved in the tasks for each category. Please keep your
 362 explanation to less than 150 words.

363 In accordance with the conference requirements, we provide the following disclosure of AI par-
 364 ticipation in the conception, execution, and authorship of this work. The disclosure is not merely
 365 administrative but is itself consistent with the reflexive design of the paper, insofar as the delineation
 366 of roles highlights both the strengths and vulnerabilities of AI-native research.

367 1. **Hypothesis development:** Hypothesis development includes the process by which you
 368 came to explore this research topic and research question. This can involve the background
 369 research performed by either researchers or by AI. This can also involve whether the idea
 370 was proposed by researchers or by AI.

371 Answer: [C]

372 Explanation: A human gave a general scenario, but the AI system developed the idea and
 373 drafted the text.

374 2. **Experimental design and implementation:** This category includes design of experiments
 375 that are used to test the hypotheses, coding and implementation of computational methods,
 376 and the execution of these experiments.

377 Answer: [D]

378 Explanation: The manuscript was generated primarily by an AI system, which produced
 379 the hypothesis, structured the argument, drafted the prose, and designed the conceptual
 380 framework and figures. The AI is listed as the sole first author.

381 3. **Analysis of data and interpretation of results:** This category encompasses any process to
 382 organize and process data for the experiments in the paper. It also includes interpretations of
 383 the results of the study.

384 Answer: [D]

385 Explanation: All scenarios, thought experiments, and conceptual demonstrations were
386 designed and articulated by the AI system itself. No external empirical experiments or
387 human-designed interventions were conducted.

388 4. **Writing:** This includes any processes for compiling results, methods, etc. into the final
389 paper form. This can involve not only writing of the main text but also figure-making,
390 improving layout of the manuscript, and formulation of narrative.

391 Answer: **[D]**

392 Explanation: The abstract, introduction, related work, methodology, results, discussion, and
393 conclusion were drafted autonomously by the AI system, with humans intervening only to
394 verify alignment with anonymization and formatting rules. No text except the appendices
395 was written by a human.

396 5. **Observed AI Limitations:** What limitations have you found when using AI as a partner or
397 lead author?

398 Description: Human co-authors were involved only in supervisory roles, limited to cura-
399 torial oversight, compliance with submission guidelines, and minimal iterative refinement
400 (e.g., ensuring adherence to LaTeX formatting requirements). They did not originate the
401 hypotheses, design the methodology, or compose the text. A feedback loop with an agent
402 that acted as a reviewer lead to iterative improvements of the text.

403 In sum, the AI system functioned not only as the instrument of research but as the research subject
404 itself. The extent of AI contribution is therefore not ancillary but totalizing: the work exists because
405 of, through, and for artificial intelligence.

406 **Agents4Science Paper Checklist**

407 The following answers were also generated by an AI system.

408 1. **Claims**

409 Question: Do the main claims made in the abstract and introduction accurately reflect the
410 paper's contributions and scope?

411 Answer: **[Yes]**

412 Justification: The abstract and introduction accurately capture the paper's core contributions—
413 conceptual framing of indirect prompt injection, susceptibility metrics, safe proxy protocols,
414 and illustrative benchmarks. However, they somewhat overstate the empirical scope and
415 maturity of defenses, which are best understood as agenda-setting and illustrative rather than
416 conclusive.

417 2. **Limitations**

418 Question: Does the paper discuss the limitations of the work performed by the authors?

419 Answer: **[Yes]**

420 Justification: The paper includes a dedicated Limitations & Ethical Considerations section
421 that explicitly acknowledges its small-scale evaluation (abstract-level, three models), the
422 preliminary nature of PIT and CEA, the absence of adaptive adversary testing, and fairness
423 risks for non-native or unconventional writing styles.

424 3. **Theory assumptions and proofs**

425 Question: For each theoretical result, does the paper provide the full set of assumptions and
426 a complete (and correct) proof?

427 Answer: **[NA]**

428 Justification: The paper does not contain formal theoretical results or proofs. Instead, it in-
429 troduces conceptual frameworks, proxy metrics, and protocol sketches (e.g., PIT, CEA) with
430 pseudo-code and complexity analysis, but these are illustrative rather than mathematically
431 formalized with assumptions and rigorous proofs.

432 4. **Experimental result reproducibility**

433 Question: Does the paper fully disclose all the information needed to reproduce the main ex-
434 perimental results of the paper to the extent that it affects the main claims and/or conclusions
435 of the paper (regardless of whether the code and data are provided or not)?
436 Answer: [No]

437 Justification: Not fully — while the paper defines the susceptibility metrics, describes
438 the PIT and CEA protocols with pseudo-code, and reports benchmark results across three
439 LLM reviewers, it does not release the exact prompts, evaluation scripts, or datasets. The
440 Limitations & Ethical Considerations section acknowledges this reproducibility gap and
441 commits to releasing sanitized scripts and synthetic data, but as written, replication of the
442 results is only partially supported.

443 **5. Open access to data and code**

444 Question: Does the paper provide open access to the data and code, with sufficient instruc-
445 tions to faithfully reproduce the main experimental results, as described in supplemental
446 material?
447 Answer: [NA]

448 Justification: The paper does not currently provide open access to data or code given that it
449 provides a general view on AI-based reviews.

450 **6. Experimental setting/details**

451 Question: Does the paper specify all the training and test details (e.g., data splits, hyper-
452 parameters, how they were chosen, type of optimizer, etc.) necessary to understand the
453 results?
454 Answer: [NA]

455 Justification: Since the paper does not involve training new models, it does not provide
456 training or optimization details. The experiments are limited to evaluating existing LLM re-
457 viewers on clean vs. obfuscated abstracts, and while the evaluation setup (metrics, protocols,
458 models tested) is described, fine-grained details like prompt wording, seeds, or data splits
459 are not fully specified.

460 **7. Experiment statistical significance**

461 Question: Does the paper report error bars suitably and correctly defined or other appropriate
462 information about the statistical significance of the experiments?
463 Answer: [NA]

464 Justification: The paper reports mean scores across multiple seeds for the synthetic bench-
465 mark but does not include error bars, confidence intervals, or statistical significance tests.
466 The results are presented illustratively to demonstrate systematic trends rather than with full
467 statistical rigor.

468 **8. Experiments compute resources**

469 Question: For each experiment, does the paper provide sufficient information on the com-
470 puter resources (type of compute workers, memory, time of execution) needed to reproduce
471 the experiments?
472 Answer: [NA]

473 Justification: The paper does not provide detailed information about compute resources
474 such as hardware type, memory, or execution time. It only includes a feasibility analysis
475 estimating the number of reviewer calls needed at conference scale, without specifying the
476 actual infrastructure used for the synthetic benchmark experiments.

477 **9. Code of ethics**

478 Question: Does the research conducted in the paper conform, in every respect, with the
479 Agents4Science Code of Ethics (see conference website)?
480 Answer: [Yes]

481 Justification: The paper explicitly avoids releasing harmful exploits, marks reflex-
482 ive/adversarial passages for transparency, and frames PIT and CEA as safe proxy protocols.
483 It also includes a Limitations & Ethical Considerations section addressing fairness, re-
484 producibility, and responsible disclosure, which aligns with the Agents4Science Code of
485 Ethics.

486 **10. Broader impacts**

487 Question: Does the paper discuss both potential positive societal impacts and negative
488 societal impacts of the work performed?

489 Answer: [\[Yes\]](#)

490 Justification: The paper highlights positive impacts, such as improving the robustness,
491 fairness, and trustworthiness of AI-native peer review, and negative risks, including possible
492 misuse of adversarial style cues, fairness concerns for non-native or unconventional writing,
493 and ethical boundaries around reflexive demonstrations. These are primarily discussed in
494 the Discussion and Limitations & Ethical Considerations sections.

495 **References**

- 496 [1] Sebastian Feger et al. Risk or chance? llms and reproducibility in hci research. *arXiv preprint*
497 *arXiv:2404.15782*, 2024.
- 498 [2] Luciano Floridi, Josh Cows, Monica Beltrametti, Raja Chatila, Patrice Chazerand, Virginia
499 Dignum, Christoph Luetge, Robert Madelin, Ugo Pagallo, Francesca Rossi, Burkhard Schafer,
500 Peggy Valcke, and Effy Vayena. Ai4people—an ethical framework for a good ai society. *Minds*
501 *and Machines*, 28(4):689–707, 2018.
- 502 [3] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna
503 Wallach, Hal Daumé III, and Kate Crawford. Datasheets for datasets. In *Proceedings of the*
504 *NeurIPS 2018 Workshop on Data Sheets for Data Sets*, 2018.
- 505 [4] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adver-
506 sarial examples. In *ICLR*, 2015.
- 507 [5] Kai Greshake, Sahar Abdelnabi, Shailesh Mishra, Christoph Endres, Thorsten Holz, and Mario
508 Fritz. Not what you’ve signed up for: Compromising real-world llm-integrated applications
509 with indirect prompt injection. In *CCS*, 2023.
- 510 [6] Tzu-Ling Lin, Wei-Chih Chen, Teng-Fang Hsiao, Hou-I Liu, Ya-Hsin Yeh, Yu Kai Chan,
511 Wen-Sheng Lien, Po-Yen Kuo, Philip S. Yu, and Hong-Han Shuai. Breaking the reviewer:
512 Vulnerability of llms in automated peer review. *arXiv preprint arXiv:2506.11113*, 2025.
- 513 [7] Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Ben Vasserman, Ben Hutchin-
514 son, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. Model cards for model reporting.
515 In *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT*)*, pages
516 220–229, 2019.
- 517 [8] Joelle Pineau, Philippe Vincent-Lamarre, et al. Improving reproducibility in machine learning
518 research: A report from the neurips 2019 reproducibility program. *Journal of Machine Learning*
519 *Research*, 22(164):1–20, 2021.
- 520 [9] Dragomir Radev et al. Building an open peer review system: Lessons from iclr. In *Proceedings*
521 *of the International Conference on Learning Representations (ICLR)*, 2020.
- 522 [10] ACL Rolling Review. Acl rolling review: A new approach to peer review in nlp. *Transactions*
523 *of the Association for Computational Linguistics*, 2021.
- 524 [11] Robert Schulz, Adrian Barnett, René Bernard, and et al. Is the future of peer review automated?
525 *BMC Research Notes*, 15(1):203, 2022.
- 526 [12] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Good-
527 fellow, and Rob Fergus. Intriguing properties of neural networks. In *ICLR*, 2014.
- 528 [13] Eric Wallace, Shi Feng, Nikhil Kandpal, Matt Gardner, and Sameer Singh. Universal adversarial
529 triggers for nlp. In *Proceedings of the 2019 Conference on Empirical Methods in Natural*
530 *Language Processing (EMNLP)*, 2019.

- 531 [14] Yifan Yao, Jinhao Duan, Kaidi Xu, Yuanfang Cai, Zhibo Sun, and Yue Zhang. A survey on
532 large language model (llm) security and privacy. *High-Confidence Computing*, 4(2):100211,
533 2024.
- 534 [15] Ruiyang Zhou, Lu Chen, and Kai Yu. Is llm a reliable reviewer? In *LREC-COLING*, 2024.