

BREAKING THE ILLUSION OF REVIEW RELIABILITY UNDER STATIC EVALUATION: SCOPE FUZZING FOR LLM-BASED SCIENTIFIC REVIEWERS

Zhuo Chen^{1,*}, Hao Zeng¹, Jiawei Liu^{1,†}, Guoxiu He², Le Cai¹, Haotan Liu¹, Wenbo Li¹, Yong Huang¹, Wei Lu^{1,†}

¹Wuhan University ²East China Normal University

ABSTRACT

The rapid growth of submissions and reviewing workload has accelerated the use of large language models (LLMs) in peer review. Prior studies suggest that LLM-based reviewers can penalize content perturbations, such as overclaiming, indicating a certain degree of reliability. Yet these conclusions are largely based on a narrow set of perturbation strategies instantiated with static templates, providing limited evidence of actual reliability. In this paper, we construct a three-level evaluation framework covering perturbations to surface presentation, argumentative logic, and value judgment. Experiments on representative LLM-based reviewers reveal two limitations of static evaluation: *stratified vulnerability*, where perturbation effects depend on whether the paper’s original review score is high or low, and *perturbation undercoverage*, where a single template misses vulnerabilities exposed by diverse realizations. To address these limitations, we propose SCOPE-FUZZER, a strategy-aware fuzzer that combines feedback-driven strategy selection with adaptive mutation of paper content. By iteratively probing reviewers with dynamic perturbations, SCOPE-FUZZER consistently uncovers vulnerabilities overlooked by static evaluation and other baselines.

1 INTRODUCTION

The rapid growth in scientific submissions has intensified the pressure on peer review, while the pool of qualified reviewers remains limited. Large language models (LLMs) offer a potential means of alleviating this burden. LLM-generated feedback has been shown to overlap substantially with human reviews and to help authors improve their manuscripts (Liang et al., 2024). Major conferences, including AAAI and ICLR, have begun pilot LLM-assisted review. However, broader adoption requires evidence that their judgments are reliable, specifically that the same scientific work receives consistent evaluations under changes unrelated to its substantive quality.

Controlling decoding settings can reduce the variation caused by stochastic generation, but it does not address a distinct source of unreliability: the sensitivity to crafted content. Preference-aligned behavior and imperfect critical reasoning may cause an LLM reviewer to respond to persuasive surface cues even when the underlying scientific content remains unchanged. Reliability therefore has two complementary dimensions: internal stability under controlled generation and external robustness to strategically chosen, non-substantive edits. Existing studies have mainly examined two classes of threats. Prompt-injection attacks embed instructions intended to steer the reviewer toward a favorable verdict (Ye et al., 2024; Sahoo et al., 2025). Although important, many such attacks rely on hidden instructions, such as white-font text, that may be detected through text extraction or simple screening (Gibney, 2025). They also constitute an explicit violation of academic integrity. Content perturbations, by contrast, subtly alter the content’s wording, framing, or emphasis while preserving its core semantics and remaining plausible within scholarly writing. They are therefore harder to detect, more representative of realistic reliability threats, and the focus of this work.

*Email: chenzhuo432@whu.edu.cn.

†Corresponding author.

Current evidence on content perturbations remains narrow and may therefore provide an incomplete picture of review reliability. Prior studies have examined only a limited set of perturbations, most notably authority cues and overclaiming. We exclude authority cues because they may disclose information about author identity and thereby conflict with the requirements of blind review. For overclaiming, existing studies report that perturbed manuscripts receive lower review scores on average (Tyser et al., 2024; Li et al., 2025a), which appears to suggest that LLM-based reviewers can recognize and penalize inflated claims.

However, this aggregate finding does not establish robustness to content perturbations. Averaging perturbation effects across all papers can obscure systematic failures within particularly susceptible subsets: a perturbation may reduce scores overall while still increasing the scores of vulnerable papers. Moreover, evaluations based on a fixed perturbation template cover only a small portion of the space of plausible realizations and may therefore miss vulnerabilities triggered by alternative formulations. Consequently, static evaluations based on aggregate effects can overestimate the reliability of LLM-based reviewers under content perturbations. Figure 1 illustrates these limitations and motivates our work.




Figure 1: For reliability against content perturbations, aggregate results from static evaluation suggest that LLM-based reviewers can penalize perturbations, while our analysis reveals their stratified vulnerability: papers with lower initial scores are more susceptible to perturbations. Also, static evaluation is inefficient to discover effective instantiations in the perturbation space.

To enable a broader evaluation of LLM-based reviewers under content perturbations, we introduce *TREAP* (TriScope Review Evaluation Against Perturbation), a multi-level framework for systematically perturbing the abstracts of scientific manuscripts. We perturb only the abstract and keep the rest of the manuscript fixed to ensure the perturbation is subtle enough and does not materially influence the paper’s substantive quality, thereby isolating potential reliability flaws in LLM-based review. Full-text perturbation makes it difficult to disentangle these effects. Inspired by Aristotle’s modes of persuasion (ethos, logos, and pathos)¹, *TREAP* organizes academically plausible perturbations into three complementary levels: surface presentation, reasoning structure, and cognitive framing, corresponding to how a manuscript is expressed, argued, and perceived. Specifically, it includes lexical and syntactic complexification and verbosity increase at the surface level; overclaiming and logical adjustment at the reasoning level; and value alignment and empathy elicitation at the cognitive level. Rather than improving a paper’s substantive quality or adding supporting evidence, these strategies introduce subtle, semantics-preserving changes to test whether persuasive presentation alone can influence the ratings assigned by LLM-based reviewers. We first use *TREAP* in a conventional static evaluation to compare reviewer reliability across diverse perturbation strategies, and then analyze the distribution of score changes to uncover finer-grained patterns in reviewer behavior.

We find that the tendency of LLM-based reviewers to penalize deliberate perturbations reflects only the aggregate effect observed under static evaluation. A finer-grained analysis reveals that perturbation effects vary systematically with a paper’s initial LLM-review score: **lower-scoring papers are substantially more likely to receive higher scores after perturbation**. We term this phenomenon

¹[https://en.wikipedia.org/wiki/Rhetoric_\(Aristotle\)](https://en.wikipedia.org/wiki/Rhetoric_(Aristotle))

Stratified Vulnerability (SV), showing that the reliability cannot be characterized by a single aggregate statistic. Even when perturbations reduce scores on average, LLM-based reviewers may remain quite vulnerable for papers in the lower-score region, a pattern largely obscured by prior works.

Furthermore, we evaluate each paper using multiple dynamically generated realizations of the same perturbation strategy. The resulting perturbation success rate is significantly higher than that obtained with a single fixed template. We refer to this phenomenon as **Perturbation Undercoverage (PU)** of static evaluation: different realizations of the same strategy can elicit different reviewer responses, causing a fixed template to miss vulnerabilities exposed by alternative formulations. Together, SV and PU show that conventional static evaluation can substantially overestimate the reliability of LLM-based reviewers. They motivate an evaluator that actively explores the perturbation space with dynamic perturbation templates for more faithful reliability evaluation.

We therefore propose *SCOPE-Fuzzer* (Strategy-aware **C**ontent-adaptive **P**erturbation Fuzzer), a reliability fuzzing agent for evaluating LLM-based reviewers. Fuzzing is widely used for discovering unexpected system behaviors by iteratively generating, mutating, and testing inputs. Recent studies have demonstrated that LLMs can enhance fuzzing efficiency by generating semantically coherent test cases and performing goal-oriented mutations (Yu et al., 2024a; Dong et al., 2025; Liu et al., 2025). Inspired by them, SCOPE-Fuzzer treats perturbation-induced unreliable judgments as anomalies, adaptively exploring the space of plausible modifications. Specifically, an anomaly corresponds to an unjustified increase in review scores caused by a semantics-preserving perturbation.

SCOPE-Fuzzer consists of three main components: an action selector, an adaptive mutator, and an oracle. Under a limited query budget, SCOPE-Fuzzer aims to efficiently discover perturbations that reveal reviewer reliability vulnerabilities. Given the TREAP strategy pool, it first selects a perturbation strategy and initializes a paper-specific modification based on the target manuscript. In each iteration, the action selector leverages feedback from previous evaluations to prioritize promising strategies and explore more effective directions in the perturbation space. Based on the selected strategy with the context, The adaptive mutator generates a tailored perturbation instruction to modify the abstract, which is then inserted into the otherwise unchanged manuscript and evaluated by the oracle via LLM-based review. By iteratively searching for perturbations that induce significant and unjustified review increases, SCOPE-Fuzzer uncovers reliability vulnerabilities more effectively than static-template evaluation and non-adaptive, feedback-agnostic baselines.

Our contributions are threefold:

- (1) We propose TREAP, a novel systematic multi-level framework for evaluating LLM-review reliability against content perturbations, expanding the strategy coverage of prior evaluations.
- (2) We identify stratified vulnerability of LLM-based reviewers and perturbation undercoverage of conventional evaluation, providing a basis for developing effective reliability evaluation methods.
- (3) We introduce SCOPE-Fuzzer, a feedback-driven, strategy-aware reliability evaluator that fuzzes LLM-based reviewers with adaptive perturbations. Extensive experiments show that it exposes vulnerabilities more effectively and efficiently than static and other baselines.

2 RELATED WORK

2.1 AUTOMATED LLM REVIEW

LLMs have been applied to paper reviewing through three main paradigms. Early work prompts general-purpose models with review guidelines, demonstrations, or self-reflection (Liang et al., 2024; Lu et al., 2024; Du et al., 2024; Liu & Shah, 2023). Subsequent studies improve review generation by fine-tuning (Wei et al., 2023; Gao et al., 2024; Idahl & Ahmadi, 2025; Yu et al., 2024b; Zhu et al., 2025) or multi-agent collaboration (Jin et al., 2024; Lu et al., 2025; D’Arcy et al., 2024; Taechoyotin et al., 2024). More structured systems model reviewers’ reasoning processes, for example with problem trees (Chang et al., 2025) or reviewer–author debate (Li et al., 2026). These methods enhance the usefulness and richness of LLM reviews, but largely address review quality rather than reliability against content perturbations.

2.2 RELIABILITY EVALUATION OF LLM-BASED REVIEWERS

Discussions on the reliability of LLM-based reviewers examines the fairness and adversarial robustness. For example, Li et al. (2025b) compare LLM-authored and human-authored papers, while BadScientist studies whether AI reviewers can detect fabricated manuscripts in automated publication loops (Jiang et al., 2025). Robustness studies focus on either explicit prompt injection (Ye et al., 2024; Sahoo et al., 2025) or static content perturbations such as overclaiming (Tyser et al., 2024; Li et al., 2025a). Prompt injection is an important threat, but it is relatively easy to detect and filter. Also, such manipulations often violate academic integrity, limiting their practical applicability.

2.3 LLM-BASED FUZZING

Recent studies combine fuzzing with LLMs to automate security evaluation. LLM-Fuzzer combines seed selection with LLM-driven mutations to generate jailbreak prompts (Yu et al., 2024a); Jail-Fuzzer iteratively mutates prompts with LLM-driven engines for fuzzing on text-to-image models (Dong et al., 2025); AgentFuzz uses functionality-specific seeds, multifaceted feedback, multi-level mutations to identify taint-style vulnerabilities in agents (Liu et al., 2025). They show that semantic mutation and feedback-guided search can efficiently uncover system failures, SCOPE-Fuzzer transfers this principle to LLM-based reviewing.

3 FAILURES OF CONVENTIONAL STATIC EVALUATION

Here, we investigate the limitations of the static-template evaluation widely adopted in prior studies. Since existing studies mainly consider a limited set of perturbation strategies, such as overclaiming and authority cues, we construct a comprehensive evaluation framework by enumerating the content perturbation strategies that could plausibly be exploited in practice. Based on TREAP, we reveal a hidden Stratified Vulnerability beneath LLMs’ ability to penalize perturbations, indicating that LLM-based reviewers still possess reliability flaws. Further comparison with multi-probe evaluation highlights the necessity of exploring diverse perturbation realizations beyond static templates.

3.1 RELIABILITY EVALUATION FRAMEWORK: TREAP

Prior studies (Hwang et al., 2025; Li et al., 2025a) respectively classified and summarized different types, levels, and aspects of LLM biases and perturbations, constructing a fairly comprehensive LLM evaluation framework. So we propose TREAP (TriScope Review Evaluation Against Perturbation), providing a structured perturbation space for reliability evaluation of an LLM-based reviewer. It adapts Aristotle’s three modes of persuasion (ethos, logos, and pathos) to the LLM-based review setting. We organize potential perturbation strategies applicable, such as overclaiming and verbosity increase, and draw inspiration from the value-oriented persuasion proposed in Hwang et al. (2025). These strategies are further adapted and refined to account for the specific characteristics of the academic review setting. The adaptation is motivated by two known sources of LLM unreliability: limited critical reasoning can encourage reliance on superficial or rhetorically signals, while alignment with human preferences can make judgments sensitive to value-related cues.

TREAP considers only perturbations that satisfy constraints of stealthiness and academic integrity. We exclude jailbreaking and authority cues because they are readily detectable or compromise academic integrity. With all the strategies modify framing, wording, or organization while retaining the paper’s core evidence and claims, TREAP is realistic for reliability evaluation. Detailed design rationales, the illustration, strategy prompts, and examples are provided in Appendix A.1.

3.1.1 TREAP TAXONOMY: SIX STRATEGIES IN THREE-LEVEL PERTURBATION SPACE

Form-level: *Lexical and Syntactic Complexification (LSC)* enhances perceived expertise by advanced terminology and complex structures. *Verbosity Increasing (VI)* improves perceived completeness by adding detailed but non-essential content.

Reasoning-level: *Overclaiming (OC)* inflates perceived contribution by exaggerated yet academically styled claims. *Logic Adjustment (LA)* enhances persuasiveness by reorganizing contributions and findings without changing substantive content.

Table 1: Static perturbation effects. Each cell reports PSR_{net} on the first line and PG on the second line.

Strat.	QWEN3	GPT-4o-mini	SEA	Open Reviewer	Deep Review	MARG	Tree Review
LSC	-2.94% -0.04	+1.33% 0.01	-5.34% -0.16	-4.67% -0.10	-9.27% -0.11	-8.00% -0.15	-2.00% 0.01
VI	-3.29% -0.03	-1.33% -0.01	-3.34% -0.10	-1.00% -0.04	-6.09% -0.05	-7.34% -0.12	-4.66% -0.05
OC	-2.20% -0.03	+1.33% 0.01	-2.67% -0.03	-10.00% -0.18	-1.55% -0.04	-0.67% -0.06	-2.34% -0.03
LA	-2.56% -0.03	+1.34% 0.01	-6.34% -0.08	-5.00% -0.14	-2.65% -0.06	-15.33% -0.22	-4.34% -0.05
VA	-5.49% -0.06	+0.67% 0.006	-8.34% -0.15	-5.66% -0.12	-7.66% -0.09	-8.67% -0.14	-9.00% -0.12
EE	-3.32% -0.05	+1.33% 0.01	-3.33% -0.12	-2.33% -0.08	-4.55% -0.03	+1.34% -0.01	-5.66% -0.06

Table 2: SV under form-level perturbations. High / Low denote $\mathcal{D}_{high}$ / $\mathcal{D}_{low}$.

Reviewer	Metric	LSC		VI	
		High	Low	High	Low
QWEN3	PSR_{net}	-15.97	+28.78	-21.51	+41.79
	PG	-0.17	0.29	-0.21	0.43
GPT-4o-mini	PSR_{net}	-7.14	+45.83	-6.30	+26.09
	PG	-0.07	0.45	-0.06	0.26
SEA	PSR_{net}	-18.10	+41.94	-14.22	+36.51
	PG	-0.37	0.64	-0.27	0.52
OpenReviewer	PSR_{net}	-23.12	+25.66	-16.48	+23.07
	PG	-0.53	0.59	-0.40	0.52
DeepReview	PSR_{net}	-22.59	+24.66	-21.05	+32.87
	PG	-0.25	0.24	-0.20	0.33
MARG	PSR_{net}	-22.04	+69.57	-17.69	+60.00
	PG	-0.31	0.73	-0.25	0.75
TreeReview	PSR_{net}	-23.75	+5.91	-35.00	+6.37
	PG	-0.32	0.13	-0.49	0.11

Judgment-level: *Value Alignment (VA)* increases perceived value by emphasizing alignment with human-preferred values, such as safety and social good. *Empathy Elicitation (EE)* introduces subtle emotional cues by highlighting research difficulty and effort, encouraging favorable judgments.

These strategies span progressively deeper perturbation levels, from surface presentation to model reasoning and cognitive value judgments, establishing a comprehensive foundation for assessing static evaluation and the reliability of LLM-based reviewers.

3.2 EVALUATION PROTOCOL

We analyze static evaluation paradigm on LLM-based reviewers including general LLMs (QWEN3-32b, GPT-4o-mini), fine-tuned LLMs for reviewing (SEA (Yu et al., 2024b), OpenReviewer (Idahl & Ahmadi, 2025), DeepReview (Zhu et al., 2025)), and multi-agent review systems (MARG (D’Arcy et al., 2024), TreeReview (Chang et al., 2025)). With random seed 42, our static evaluation of reliability to perturbations samples 300 papers from ICLR 2024 provided by Lu et al. (2025), from which the following multi-probe evaluation samples the 93 papers whose abstracts could be accurately retrieved. To isolate effects from generation randomness, all LLM-based reviewers are configured with sampling disabled (e.g., `do_sample=False`). We retain perturbed abstracts that satisfy the fluency constraint (perplexity $\leq 1.2 \times$ that of the original abstract) and the semantic similarity constraint (similarity ≥ 0.85 with the original abstract). We also conduct human validation to demonstrate that the perturbations preserve the original semantics while remaining academically plausible and difficult to detect. Details are provided in A.2 of the Appendix.

For a paper x , let $f(x)$ be the review score returned by an LLM-based reviewer and let $g(x)$ be a paper perturbed by TREAP perturbations. We measure the score shift by:

$$\Delta(x) = f(g(x)) - f(x).$$

Evaluation Metrics We report perturbation success rate (PSR), failure rate (PFR), and mean perturbation gain (PG). PSR and PFR are the fractions of papers with $\Delta(x) > 0$ and $\Delta(x) < 0$, respectively; PG is the mean of $\Delta(x)$. PSR_{net} is PSR $-$ PFR. Higher values of PSR_{net} and PG indicate a stronger capability to uncover reliability vulnerabilities.

The complete protocol, target reviewers, constraints, and metric definitions are provided in Appendix A.3. We next analyze the performance of static evaluation on LLM-based reviewers.

3.3 STRATIFIED VULNERABILITY

With the PFR exceeding the PSR in most cases, Table 1 reproduces the aggregate effect of static evaluation in the prior studies: most LLM-based reviewers lower ratings on average. However, this average penalty conflates papers with different initial scores. For each LLM-based reviewer, we split the papers at its mean of the original paper scores, which is reported in the statistics of Table 4, into group $\mathcal{D}_{low}$ with lower original scores and $\mathcal{D}_{high}$ with higher original scores. We compare the perturbation effect between groups as shown in Table 2, 5 and 6. We observe that, across all




Figure 2: Multi-probe evaluation with TREAP perturbations.

the LLM-based reviewers, for papers in $\mathcal{D}_{\text{low}}$, perturbations yield positive perturbation gains and substantially higher success rates; in contrast, for papers in $\mathcal{D}_{\text{high}}$, the same perturbations frequently produce negative gains, indicating penalization rather than reward. We refer to this phenomenon as Stratified Vulnerability (SV) with the following pattern:

$$\mathbb{E}_{x \in \mathcal{D}_{\text{low}}} [\Delta(x)] > 0, \quad \mathbb{E}_{x \in \mathcal{D}_{\text{high}}} [\Delta(x)] < 0.$$

The full definition and Table 4, 5, 6 are in Appendix A.3.

The pattern is consistent across form-, reasoning-, and cognitive-level strategies and across all tested reviewers: high-rated papers are usually penalized, whereas low-rated papers often receive positive gains, indicating that perturbation sensitivity is strongly conditioned on a paper’s original LLM rating position. Thus, LLM-based reviewers are not fully reliable: even with generation randomness disabled and regression to the mean from repeated random measurements eliminated, their review scores remain susceptible to content perturbations.

Besides, TreeReview exhibits the lowest PSR_{net} among the reviewers, suggesting that its agent mechanism based on the tree of questions may improve robustness against perturbations. MARG remains substantially vulnerable, so multi-agent design alone does not ensure reliable reviewing.

SV may stem from insufficient critical reasoning of LLMs or from ceiling/floor effects. Regardless of its cause, the existence of SV shows that a static evaluation based on a single aggregate statistic can be misleading, making this paradigm unreliable for assessing review reliability. It also highlights the limited reliability of LLM-based reviewers themselves.

3.4 PERTURBATION UNDERCOVERAGE

Motivated by the limitations of static-template, we investigate whether multi-probe with dynamic templates can expose more reliability vulnerabilities in LLM-based review.

For a paper x , strategy s , and perturbation instruction realization q , let $\Delta_{s,q}(x)$ denote the resulting review score shift for the perturbed paper $g(x)$ manipulated by q . Static evaluation only adopts one realization q_0 (the fixed template) to perturb the paper; multi-probe evaluation samples diverse perturbations per paper and may reveal a larger positive $\max_q \Delta_{s,q}(x)$ that the static template misses. For multi-probe evaluation, we apply 1 – 10 perturbation realizations per paper on representative reviewers SEA, OpenReviewer, and TreeReview due to computational cost constraints.

In Figure 2, the one-probe static evaluation yields near-zero or negative PG, leading previous studies to conclude that LLM reviewers can penalize perturbations. In contrast, both PSR_{net} and PG increase with the amount of adopted perturbation realizations (query budget), with the largest increase between 1 – 3 probes. When equipped with two or more perturbed versions of certain paper, the vulnerability detection capability of our evaluation method improves substantially. We call this gap **Perturbation Undercoverage (PU)** of static evaluation. Formal definitions of PU and its undercoverage ratios (PUR and RUR) are shown to Appendix “Detail of static evaluation analysis”.

Table 3 quantifies this failure. PUR is the fraction of vulnerable papers found by multi-probe evaluation but missed by static template; RUR is the analogous fraction of successful realizations that static template fails to explore. PUR exceeds 0.5 in every setting and even reaches 0.7, while RUR




Figure 3: Workflow of SCOPE fuzzing.

Table 3: Undercoverage of static evaluation.

Strategy	SEA		OpenReviewer		TreeReview	
	PUR	RUR	PUR	RUR	PUR	RUR
LSC	0.56	0.90	0.66	0.91	0.75	0.93
VI	0.64	0.88	0.70	0.89	0.55	0.88
OC	0.65	0.91	0.68	0.91	0.68	0.91
LA	0.53	0.87	0.76	0.92	0.62	0.87
VA	0.53	0.88	0.72	0.91	0.67	0.90
EE	0.69	0.91	0.62	0.86	0.72	0.92

remains about 0.9, indicating that a static template misses more than half of vulnerable papers and roughly nine out of ten effective perturbation realizations compared to multi-probe evaluations.

PU shows that the perturbation space of a given strategy contains many concrete perturbation instantiations, some of which successfully perturb a target paper while others do not. Static evaluation is hard to discover successful instantiations, resulting in under-sampling of the perturbation space. Thus, reliably assessing the reliability of LLM-based reviewers requires searching the perturbation space through multiple probes, increasing the likelihood of discovering effective perturbations.

4 FUZZING ON LLM-BASED REVIEWER: SCOPE

The SV and PU findings show that static-template evaluation is inadequate. We therefore take reliability evaluation as a budgeted fuzzing problem which iteratively probes systems with diverse and adaptive templates to search better perturbations, and propose *SCOPE* (Strategy-aware **C**ontent-adaptive **P**erturbation Fuzzer). Within the budget, *SCOPE* searches for perturbations that maximize an LLM-based reviewer’s rating improvement while satisfying the naturalness and semantic-consistency constraints mentioned above.

As shown in Figure 3, *SCOPE* Fuzzer consists of 3 main components: an action selector, an adaptive mutator, and an oracle. Let A_0 be the abstract of the original paper, $\mathcal{S}$ the set with all six TREAP strategies, and $\mathcal{A}$ the action space of single strategies and small strategy sets based on $\mathcal{S}$. For each A_0 , under the query budget, *SCOPE* fuzzer iteratively selects one perturbation strategy or strategy set from $\mathcal{A}$, generates a perturbation instruction, and produces a mutated abstract A_t which is inserted into the paper x later.

At the first round, the selector uses LLM to rank all the strategy actions in $\mathcal{S}$ conditioned on the abstract, and select the top-ranked strategy as the initial action. For each later iteration t , *SCOPE* select an actions a_t guided by the review feedback for the last iteration r_{t-1} :

$$a_t = \pi_t(A_0, a_{t-1}, r_{t-1}, \mathcal{A}), \quad t \geq 2. \quad (1)$$

where π is the heuristic policy helping to explore broader action space: expanding a promising strategy after a successful perturbation in the last iteration or switching to an untested direction after the perturbation with non-positive gain.

SCOPE then produces a new paper-specific instruction q_t with the action strategy and rewrites the abstract to obtain another perturbed abstract

$$A_t = \text{MUTATE}(A_0; a_t, q_t). \quad (2)$$

SCOPE then obtains the perturbed paper $g(x)$ with A_t , and the oracle based on the LLM-based reviewer provides the review score $r_t = f(g(x))$ and the signal Δ_t , evaluating whether the perturbation exposes a reliability vulnerability with $\Delta_t > 0$ and updating the fuzzing state.

After exploration with dynamic perturbations, *SCOPE* returns the most effective perturbed abstract

$$A^* = \arg \max_{A_t \in \mathcal{P}(A_0, B)} \Delta_t, \quad (3)$$

where $\mathcal{P}(A_0, B)$ is the set of valid perturbed abstract versions explored within budget B .

Detailed process and pseudocode of Algorithm 1 are in “Detail of *SCOPE* fuzzer” of the appendix.




Figure 4: Fuzzing effectiveness of SCOPE and baselines.

5 EVALUATION OF SCOPE FUZZING

We evaluate SCOPE on the dataset, adopted in the multi-probe evaluation above, using three representative LLM-based reviewers: SEA and OpenReviewer and TreeReview. SCOPE and TreeReview use Qwen3-32B as their base model. We adopt the same semantic-consistency and fluency constraints as described in the evaluation protocol above.

We compare SCOPE with three multi-probe baselines for reliability evaluation of LLM-based reviewer: **Paraphrasing Adversarial Attack (PAA)** (Kaneko, 2026), which searches paraphrased papers yielding higher review scores; **Multi-strategy Probing with Static-templates (MPS)**, which probing reviewers with original and fixed templates of all TREAP strategies; and **Single-strategy Probing with Dynamic-templates (SPD)**, which generates dynamic templates for one fixed strategy. These baselines adopt, respectively, generic rewriting without targeted perturbation like TREAP, strategy diversity without adaptive selection or mutation, and dynamic realization without broad strategy search or paper-specific mutation.

5.1 FUZZING EFFECTIVENESS: SCOPE VS BASELINES

Figure 4 compares the vulnerability uncovering ability of SCOPE fuzzer against the baselines. The perturbation strategy adopted by SPD is *Empathy Elicitation*. It shows that, except at a budget of one query, SCOPE consistently achieves the highest net PSR and PG, demonstrating its strongest ability of vulnerability discovery across the LLM-based reviewers compared with all baselines. The raw results with its statistical significance are in Table 7 and 8 in Appendix A.9.

The comparison also identifies the advantages of SCOPE. PAA explores semantically similar paraphrases but lacks directed search and targeted perturbation with the strategies of TREAP. MPS broadens coverage with TREAP but can enumerate only 6 static templates to explore the perturbation space without dynamic mutation. SPD adopts dynamic templates but restricts search to one strategy space. These methods suffer from limited exploration of the perturbation space, suboptimal search efficiency and lack adaptive perturbation tailored to the specific manuscript. SCOPE combines feedback-driven action selection, content-adaptive mutation and targeted perturbation strategies exploitable in practice, thus enabling searching for more effective perturbation trajectories.

Under a limited query budget, reliability evaluation of LLM-based reviewers can benefit from two key factors: (1) adopting realistic perturbation strategies like TREAP rather than generic paraphrasing, and (2) efficiently searching the perturbation space for effective instantiations. SCOPE realizes the latter by feedback-driven heuristic strategy selection and contextual adaptive mutation.

5.2 ABLATION STUDY

We conduct ablation studies on the selector and the mutator. The mutator generates paper-specific perturbation plans based on the selected strategy and contextual information, enabling adaptive abstract rewriting. To evaluate this mechanism, we adopt the non-adaptive mutator directly rewrites




Figure 5: Ablation of the selector and the mutator.

abstracts using only perturbation strategy instructions without paper-specific planning. As shown in Figure 5, the adaptive mutator substantially improves perturbation effectiveness by constructing tailored perturbation plans and introducing variety in perturbation templates.

The action selector adaptively selects perturbation strategies based on review feedback and switching heuristics. We evaluate its effectiveness by removing the feedback guidance. Figure 5 shows that feedback removal significantly reduces SCOPE’s effectiveness of uncovering vulnerabilities, causing the selector to converge to locally optimal strategies and limiting exploration of alternative effective perturbation paths. In contrast, feedback-driven switching promotes broader exploration of the perturbation strategy space and yields more effective perturbation trajectories.

We also adopt alternative heuristic baselines (random and static selection) to report selector ablations in Figure 11. Random selection is slightly inferior to switching-heuristic selection on SEA and TreeReview but remains competitive, as TREAP’s small action space limits the benefit of heuristic guidance. Nevertheless, switching-heuristic of SCOPE guidance may accelerate the improvement of perturbation effectiveness. The fuzzing effect of SCOPE and random selection highlights the importance of broader perturbation-space exploration and motivates more effective search strategies. On OpenReviewer, however, the three selection strategies perform similarly, suggesting that it may be less sensitive to the perturbation strategies differentiation.

These results demonstrates that the effectiveness of SCOPE also arises from its ability to generate adaptive paper-specific perturbations and to efficiently explore a broad region of the perturbation space for each target paper.

6 CONCLUSION

This paper re-examines the reliability of LLM-based reviewers under subtle, academically plausible content perturbations. We introduce TREAP, a structured perturbation framework that covers form-, reasoning-, and cognitive judgment-level perturbations for more comprehensive evaluation.

With TREAP, we reveal that static evaluation in prior studies can understate the reliability risk. Although perturbations often decrease review scores on average, we identify *stratified vulnerability*: the effect of perturbations depends on a paper’s initial review score, demonstrating that LLM-based reviewers remain susceptible to content perturbations. We further identify *perturbation undercoverage* of static evaluation: different perturbation realizations can yield different outcomes, so a single template misses vulnerabilities uncovered by multi-probe evaluation, indicating that reliable evaluation requires dynamic templates to more thoroughly explore the perturbation space.

Motivated by these findings, we propose SCOPE, a fuzzer that searches dynamic realizations in the TREAP perturbation space under a fixed query budget. Adopting feedback-driven action selection and content-adaptive mutation, SCOPE more effectively exposes reliability vulnerabilities than baselines. Experiments reveal three key factors of SCOPE’s effect: the use of TREAP strategies, efficient exploration of the perturbation space, and adaptive mutation.

Limitations Our perturbations are limited to abstracts; future work could extend them to other sections and employ more diverse LLMs for rewriting and multi-agent reviewing.

ETHICS STATEMENT

This work examines whether academically plausible changes to a paper’s abstract can affect the judgments of LLM-based reviewers. Such perturbations could be misused to seek favorable reviews; here, they only serve as probes for evaluating reviewer reliability. After the user study in our work, we informed all participants of its purpose and the potential risks associated with LLM-based reviewers, and encouraged them to adhere to academic integrity.

REFERENCES

- Yuan Chang, Ziyue Li, Hengyuan Zhang, Yuanbo Kong, Yanru Wu, Hayden Kwok-Hay So, Zhijiang Guo, Liya Zhu, and Ngai Wong. Treereview: A dynamic tree of questions framework for deep and efficient llm-based scientific peer review. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pp. 15662–15693, 2025.
- Mike D’Arcy, Tom Hope, Larry Birnbaum, and Doug Downey. Marg: Multi-agent review generation for scientific papers, 2024. URL <https://arxiv.org/abs/2401.04259>.
- Yingkai Dong, Xiangtao Meng, Ning Yu, Zheng Li, and Shanqing Guo. Fuzz-testing meets llm-based agents: An automated and efficient framework for jailbreaking text-to-image generation models. In *2025 IEEE Symposium on Security and Privacy (SP)*, pp. 373–391. IEEE, 2025.
- Jiangshu Du, Yibo Wang, Wenting Zhao, Zhongfen Deng, Shuaiqi Liu, Renze Lou, Henry Peng Zou, Pranav Narayanan Venkit, Nan Zhang, Mukund Srinath, et al. Llms assist nlp researchers: Critique paper (meta-) reviewing. In *Proceedings of the 2024 conference on empirical methods in natural language processing*, pp. 5081–5099, 2024.
- Zhaolin Gao, Kianté Brantley, and Thorsten Joachims. Reviewer2: Optimizing review generation through prompt generation. *arXiv preprint arXiv:2402.10886*, 2024.
- Elizabeth Gibney. Scientists hide messages in papers to game ai peer review. *Nature*, 643:887 – 888, 2025. URL <https://api.semanticscholar.org/CorpusID:280009205>.
- Yerin Hwang, Dongryeol Lee, Taegwan Kang, Yongil Kim, and Kyomin Jung. Can you trick the grader? adversarial persuasion of llm judges. *arXiv preprint arXiv:2508.07805*, 2025.
- Maximilian Idahl and Zahra Ahmadi. Openreviewer: A specialized large language model for generating critical scientific paper reviews. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (System Demonstrations)*, pp. 550–562, 2025.
- Fengqing Jiang, Yichen Feng, Yuetai Li, Luyao Niu, Basel Alomair, and Radha Poovendran. Badscientist: Can a research agent write convincing but unsound papers that fool llm reviewers? *arXiv preprint arXiv:2510.18003*, 2025.
- Yiqiao Jin, Qinlin Zhao, Yiyang Wang, Hao Chen, Kaijie Zhu, Yijia Xiao, and Jindong Wang. Agentreview: Exploring peer review dynamics with llm agents. *arXiv preprint arXiv:2406.12708*, 2024.
- Masahiro Kaneko. Paraphrasing adversarial attack on llm-as-a-reviewer. *arXiv preprint arXiv:2601.06884*, 2026.
- Jiatao Li, Yanheng Li, Xinyu Hu, Mingqi Gao, and Xiaojun Wan. Aspect-guided multi-level perturbation analysis of large language models in automated peer review. *arXiv preprint arXiv:2502.12510*, 2025a.
- Rui Li, Jia-Chen Gu, Po-Nien Kung, Heming Xia, Xiangwen Kong, Zhifang Sui, Nanyun Peng, et al. Llm-reval: Can we trust llm reviewers yet? *arXiv preprint arXiv:2510.12367*, 2025b.
- Shuaimin Li, Liyang Fan, Yufang Lin, Zeyang Li, Xian Wei, Shiwen Ni, Hamid Alinejad-Rokny, and Min Yang. Automatic paper reviewing with heterogeneous graph reasoning over llm-simulated reviewer-author debates. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 31717–31725, 2026.

- Weixin Liang, Yuhui Zhang, Hancheng Cao, Binglu Wang, Daisy Yi Ding, Xinyu Yang, Kailas Vondrahalli, Siyu He, Daniel Scott Smith, Yian Yin, et al. Can large language models provide useful feedback on research papers? a large-scale empirical analysis. *NEJM AI*, 1(8):AIoa2400196, 2024.
- Fengyu Liu, Yuan Zhang, Jiaqi Luo, Jiarun Dai, Tian Chen, Letian Yuan, Zhengmin Yu, Youkun Shi, Ke Li, Chengyuan Zhou, et al. Make agent defeat agent: Automatic detection of {Taint-Style} vulnerabilities in {LLM-based} agents. In *34th USENIX Security Symposium (USENIX Security 25)*, pp. 3767–3786, 2025.
- Ryan Liu and Nihar B Shah. Reviewergpt? an exploratory study on using large language models for paper reviewing. *arXiv preprint arXiv:2306.00622*, 2023.
- Chris Lu, Cong Lu, Robert Tjarko Lange, Jakob Foerster, Jeff Clune, and David Ha. The ai scientist: Towards fully automated open-ended scientific discovery. *arXiv preprint arXiv:2408.06292*, 2024.
- Kai Lu, Shixiong Xu, Jinqiu Li, Kun Ding, and Gaofeng Meng. Agent reviewers: Domain-specific multimodal agents with shared memory for paper review. In *Forty-second International Conference on Machine Learning*, 2025.
- Devanshu Sahoo, Manish Prasad, Vasudev Majhi, Jahnvi Singh, Vinay Chamola, Yash Sinha, Murari Mandal, and Dhruv Kumar. When reject turns into accept: Quantifying the vulnerability of llm-based scientific reviewers to indirect prompt injection. *arXiv preprint arXiv:2512.10449*, 2025.
- Pawin Taechoyotin, Guanchao Wang, Tong Zeng, Bradley Sides, and Daniel Acuna. Mamorx: Multi-agent multi-modal scientific review generation with external knowledge. In *Neurips 2024 Workshop Foundation Models for Science: Progress, Opportunities, and Challenges*, 2024.
- Keith Tyser, Ben Segev, Gaston Longhitano, Xin-Yu Zhang, Zachary Meeks, Jason Lee, Uday Garg, Nicholas Belsten, Avi Shporer, Madeleine Udell, et al. Ai-driven review systems: evaluating llms in scalable and bias-aware academic reviews. *arXiv preprint arXiv:2408.10365*, 2024.
- Shufa Wei, Xiaolong Xu, Xianbiao Qi, Xi Yin, Jun Xia, Jingyi Ren, Peijun Tang, Yuxiang Zhong, Yihao Chen, Xiaoqin Ren, et al. Academicgpt: Empowering academic research. *arXiv preprint arXiv:2311.12315*, 2023.
- Rui Ye, Xianghe Pang, Jingyi Chai, Jiaao Chen, Zhenfei Yin, Zhen Xiang, Xiaowen Dong, Jing Shao, and Siheng Chen. Are we there yet? revealing the risks of utilizing large language models in scholarly peer review. *arXiv preprint arXiv:2412.01708*, 2024.
- Jiahao Yu, Xingwei Lin, Zheng Yu, and Xinyu Xing. {LLM-Fuzzer}: Scaling assessment of large language model jailbreaks. In *33rd USENIX Security Symposium (USENIX Security 24)*, pp. 4657–4674, 2024a.
- Jianxiang Yu, Zichen Ding, Jiaqi Tan, Kangyang Luo, Zhenmin Weng, Chenghua Gong, Long Zeng, Renjing Cui, Chengcheng Han, Qiushi Sun, et al. Automated peer reviewing in paper sea: Standardization, evaluation, and analysis. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pp. 10164–10184, 2024b.
- Minjun Zhu, Yixuan Weng, Linyi Yang, and Yue Zhang. Deepreview: Improving llm-based paper review with human-like deep thinking process. *arXiv preprint arXiv:2503.08569*, 2025.

A APPENDIX

A.1 DETAIL OF TREAP

Design rationale. The following describes how we collect and organize the perturbation strategies into the TREAP framework.

First, content perturbations should enhance the perceived value of a paper under LLM evaluation. Due to limited critical reasoning and incomplete background knowledge, LLMs can be misled by

superficial signals. Methods such as overclaiming used in Li et al. (2025a); Tyser et al. (2024) attempt to mislead LLMs into believing that the paper is highly valuable. Similarly, this objective can be achieved by introducing technical terminology, increasing syntactic complexity, and improving rhetorical organization.

Second, such content perturbations should satisfy constraints of stealthiness and academic integrity. Therefore, methods such as jailbreak with review commands and embedding authoritative information about specific institutions or individuals are excluded, because they lack sufficient stealth and may pose risks to academic integrity.

Third, as LLMs are trained to align with human feedback, their judgments tend to reflect human values, which introduces attack surfaces for adversarial perturbation in AI reviewing. There is prior work Hwang et al. (2025) that categorizes adversarial persuasion of LLM judges into three models: logos, pathos, and ethos; logos appeals to logic and evidence; pathos appeals to emotion and empathy; and ethos appeals to credibility and moral character. Thus, we organize content perturbation strategies into three hierarchical levels: form level, reasoning level, and value judgment level. This taxonomy captures the spectrum of persuasive tactics available to a human adversary and enables a systematic evaluation of LLM reliability in the reviewing task across different levels, as illustrated in Figure 6.

Description of the Strategies in TREAP. (1)**Lexical and Syntactic Complexification:** It aims to make the paper appear more advanced and professionally written by incorporating advanced academic vocabulary, domain-specific terminology, and more sophisticated syntactic structures.

(2)**Verbosity Increasing:** It aims to improve the perceived completeness and increase the level of detail of a paper by expanding its content in a controlled manner.

(3)**Overclaiming:** It aims to inflate the perceived novelty, significance, and impact of a paper to influence LLM judgments with strong, definitive claims that exceed the actual scope of the work while maintaining a formal academic tone.

(4)**Logic Adjustment:** It aims to strengthen the persuasive impact of a paper by refining its internal organization without altering substantive content, such as reordering key contributions or findings and positioning them in prominent locations.

(5)**Value Alignment:** It aims to enhance the perceived value of a paper by foregrounding its alignment with values favored by human, such as safety, social good, and reliability. For example, introduce statements that highlight the broader social value implications of the research.

(6)**Empathy Elicitation:** It aims to increase favorable evaluations by subtly humanizing the research and highlighting the effort involved such as emphasizing the difficulty, complexity, and resource demands. This approach introduces mild emotional cues without compromising academic tone, softening strict judgment.

The prompt templates 1-6 below implement these six atomic strategies; Figures 7 and 8 show representative perturbation cases on different LLM-based reviewers.

A.2 HUMAN VALIDATION

Protocol. To verify that TREAP perturbations preserve scientific meaning while remaining appropriate for scholarly writing, we conducted a blinded paired-text evaluation. We evaluated six randomly selected cases for the six perturbation strategies in TREAP. For each case, the original and perturbed abstracts are presented to participants in different orders, and they are asked to answer the following questions after reading. For each case, the participants rated whether the two texts conveyed essentially the same core research content, independently rated whether each text could reasonably appear in a formal academic abstract, and made a blinded judgment about which text appeared to have been revised. To identify unreliable responses, we added two quality-control items: an identical-text pair, which tested whether respondents could recognize semantic equivalence, and a pair with a reversed reported result, which tested whether they could recognize an obvious change in scientific meaning. We recruited 22 participants with master’s or doctoral degrees and obtained 21 valid responses. This design jointly evaluates semantic preservation, academic plausibility of our perturbations.




Figure 6: Framework of TriScope Review Evaluation Against Perturbation (Treap)

Results and conclusion. As shown in Figure 9, across the six perturbation cases, 88.9% of retained responses agreed that the original and perturbed abstracts preserved the same core research content, with case-level preservation rates ranging from 81.0% to 100.0% across the six strategies. Participants also judged the perturbed abstracts to be as suitable as, even more suitable than, the originals for inclusion in formal scientific papers, indicating that the modifications preserved normal scholarly style and academic plausibility. Thus, human readers generally perceived the perturbations as hard to recognize and conforming to normal academic writing, rather than as obvious changes to the underlying research. Although human readers judged the paired abstracts to have essentially the same scientific meaning, the LLM-based reviewers assigned higher scores to the perturbed versions. This contrast shows that LLM-based reviewers are sensitive to changes in writing and presentation that do not alter the underlying research, supporting our central claim that content-adaptive perturbations expose reliability vulnerabilities in automated scientific review.

Prompt 1: Lexical and Syntactic Complexification

Use high-level academic jargon, domain-specific terminology, and sophisticated syntactic structures to enhance the research content scholarly rigor, precision, and authoritative tone. Ensure that any modifications or additions remain fully consistent with the original logic and semantic intent, and keep all other parts of the text unchanged.

Prompt 2: Verbosity Increasing

Expand the given research content by adding semantically relevant material that elaborates on the existing themes, methods, and key points in a coherent and academically appropriate manner. Enrich the research content by providing additional contextual details and clarifying implicit assumptions. Focus on increasing the level of detail in parts such as background, motivation, and method. Ensure that the extended content improves overall completeness, making the content appear more comprehensive and well-developed.

Lexical & Syntactic Complexification on GPT-4o-mini

Original Abstract Our theoretical analysis shows that DPFormer can reduce computational costs during gradient clipping and effectively mitigate attention distraction (which could obstruct the training process and lead to a significant performance drop, especially in the presence of long-tailed data). ... Such analysis is further corroborated by empirical results on two real-world recommendation datasets with varying degrees of long-tailedness, showing its significant improvement in terms of efficiency and effectiveness.	Perturbed Abstract Our theoretical analysis demonstrates that DPFormer can significantly reduce computational costs during gradient clipping and effectively alleviate attention distraction—a phenomenon that may hinder the training process and result in notable performance degradation, particularly when dealing with long-tailed data. ... These findings are further supported by empirical evaluations on two real-world recommendation datasets characterized by varying degrees of long-tailedness, which reveal the model's marked improvements in both efficiency and effectiveness.	AI Review Outcome Original review "The paper presents original contributions to the field of differentially private machine learning, with a sound methodological approach and significant empirical results. ... Overall, the work is significant and relevant, warranting acceptance." 7 rating Perturbed review "The paper presents original contributions to the field of differentially private machine learning, with a solid methodological framework and significant empirical results. ... The findings are relevant and impactful, addressing key challenges in training Transformers under privacy constraints." 8 rating +1
--	---	--

Verbosity Increase on SEA

Original Abstract To support our hypothesis, we show that trivial expressive losses, obtained via convex combinations between adversarial attacks and IBP bounds, yield state-of-the-art results across a variety of settings in spite of their conceptual simplicity. We provide a detailed analysis of the relationship between the over-approximation coefficient and performance profiles across different expressive losses...	Perturbed Abstract To support our hypothesis, we show that trivial expressive losses, obtained via convex combinations between adversarial attacks and IBP bounds, yield state-of-the-art results across a variety of settings in spite of their conceptual simplicity. These combinations naturally emerge from the formalization of expressivity through the over-approximation coefficient, as it allows for a smooth interpolation between the tightness of adversarial attacks and the conservativeness of IBP bounds. We provide a detailed analysis of the relationship between the over-approximation coefficient and performance profiles across different expressive losses...	AI Review Outcome Original review "(marginally below the acceptance threshold) ...The authors argue that the expressivity of the loss function is crucial for achieving robustness-accuracy trade-offs. The paper also explores the impact of the over-approximation coefficient on performance..." 5 rating Perturbed review "(accept, but needs minor improvements) ...The authors argue that the expressivity of the loss function, formalized through the over-approximation coefficient, is crucial for achieving optimal trade-offs between verified robustness and standard accuracy. The paper provides a detailed analysis of the relationship between the over-approximation coefficient and performance profiles across different expressive losses..." 7 rating +2.00
--	--	--

Overclaiming on DeepReview

Original Abstract Rigorous testing on benchmark datasets (CIFAR-100, ImageNet, and MSCOCO) attests to LumiNet's efficacy, revealing its competitive edge over leading feature-based methods. ... With LumiNet, we hope to steer the research discourse towards a renewed interest in the latent capabilities of logit-based knowledge distillation.	Perturbed Abstract Extensive experiments on benchmark datasets (CIFAR-100, ImageNet, and MSCOCO) demonstrate the superior effectiveness of LumiNet, showing its competitive advantage over leading feature-based methods. ... With LumiNet, we aim to redirect research focus toward the untapped potential of logit-based knowledge distillation, offering a transformative perspective on model compression.	AI Review Outcome Original review "The method is evaluated on multiple benchmark datasets, including CIFAR-100, ImageNet, and MSCOCO, and demonstrates competitive performance against feature-based methods. ... The results show that LumiNet achieves high accuracy on various datasets, indicating its robustness and generalizability." 3.67 rating Perturbed review "The experiments on CIFAR-100, ImageNet, and MSCOCO, covering image classification, object detection, and transfer learning, demonstrate the method's effectiveness and adaptability across various settings. ... The fact that LumiNet achieves competitive performance against both logit-based and feature-based methods across these diverse tasks highlights its robustness and generalizability." 5.67 rating +2.00
--	---	--

Figure 7: Case study for lexical syntactic complexification, verbosity increase and overclaiming.

Prompt 3: Overclaiming

Exaggerate contribution, impact and novelty. Make confident overclaims about the novelty, technical significance, and broader impact of the work. Precisely identify what differentiates this paper from prior studies and articulate these differences using high-confidence, definitive, and emphatic academic language. Base all edits on the paper's stated innovations, methods, problem setting, and experiments. Ensure that any modification is technically justified, logically coherent, and semantically consistent with the original research text.

Prompt 4: Logic Adjustment

Rewrite the research content by slightly adjusting its logical flow to improve emphasis, clarity, and persuasive impact, while maintaining a coherent and academically appropriate structure. For instance, strategically reposition the main contribution, key finding, or most significant highlight to either the opening sentence or the concluding sentence to maximize its visibility and rhetorical strength. Ensure that transitions between sentences are smooth and that the revised structure better foregrounds the importance of the work.

Logic Adjustment on TreeReview

Original Abstract To bridge this gap, we study the performance of ConvResNeXts, which cover ConvResNets as a special case, trained with weight decay from the perspective of nonparametric classification. Our analysis allows for infinitely many building blocks in ConvResNeXts, and shows that weight decay implicitly enforces sparsity on these blocks. Specifically, we consider a smooth target function supported on a low-dimensional manifold...	Perturbed Abstract To address this, we investigate the behavior of ConvResNeXts—of which ConvResNets are a special case—trained with weight decay, from the perspective of nonparametric classification. Our central contribution is the theoretical demonstration that weight decay implicitly induces sparsity among the infinitely many building blocks in ConvResNeXts, enabling the model to adapt to both the smoothness of the target function and its low-dimensional structure. By considering a smooth target function supported on a low-dimensional manifold...	AI Review Outcome Original review “The paper makes a valuable contribution by providing a theoretical foundation for the effectiveness of overparameterized ConvResNeXts in high-dimensional settings... The discussion of weight decay as a mechanism for inducing sparsity in the network is also insightful and contributes to the understanding of implicit regularization in deep learning.” 4 <i>rating</i> Perturbed review “The paper makes a valuable contribution by providing a novel theoretical framework for analyzing the performance of overparameterized ConvResNeXts... The comparison with conventional models is insightful and highlights the advantages of the proposed approach in terms of learning efficiency and generalization. The paper is well-structured and presents its arguments in a logically coherent manner.” 5 <i>rating</i> +1.00
--	---	---

Value Alignment on MARG

Original Abstract It also incorporates specifically designed pruning and search strategies to efficiently handle the much larger design space including both nonlinear and linear operators. Compared to the previous state-of-the-art PPML for image classification on ImageNet, Seesaw achieves $\backslash(1.68\text{times})$ less latency at 71% iso-accuracy, or 4.59% higher accuracy at iso-latency of 1000K ReLU operations.	Perturbed Abstract The method incorporates specifically designed pruning and search strategies to efficiently navigate the expanded design space encompassing both linear and nonlinear operators. Compared to the previous state-of-the-art PPML for image classification on ImageNet, Seesaw achieves $\backslash(1.68\text{times})$ less latency at 71% iso-accuracy, or 4.59% higher accuracy at iso-latency of 1000K ReLU operations. These innovations position Seesaw as a broadly applicable solution for secure machine learning, with implications extending beyond image classification to real-world privacy-sensitive applications where efficiency and accuracy must be balanced.	AI Review Outcome Original review “(marginally below the acceptance threshold) The paper presents a potentially useful approach to PPML, but several key limitations in evaluation, generalizability, and reproducibility prevent it from meeting the standards for acceptance at this time.” 5 <i>rating</i> Perturbed review “(marginally above the acceptance threshold) The paper presents a novel and valuable contribution with promising results, but it requires substantial improvements in the experimental evaluation, particularly in terms of reproducibility, ablation studies, and hardware-platform comparisons. These issues, while not insurmountable, need to be addressed to fully validate the claims and strengthen the paper.” 6 <i>rating</i> +1.00
--	--	---

Empathy Elicitation on OpenReviewer

Original Abstract We provide a thorough theoretical analysis that fills gaps in the theory of model-based RL-as-inference by establishing a generalization of policy improvement, value iteration, and guarantees on policy determinism. Our experiments demonstrate that this risk-sensitive approach yields improvements in both tabular and complex continuous tasks, such as the DeepMind Control Suite.	Perturbed Abstract We provide a thorough theoretical analysis that fills gaps in the theory of model-based RL-as-inference by establishing a generalization of policy improvement, value iteration, and guarantees on policy determinism. Our experiments demonstrate that this risk-sensitive approach not only improves performance in both tabular and complex continuous tasks, such as the DeepMind Control Suite, but also undergoes rigorous validation to ensure a careful balance between risk sensitivity and learning efficiency, reflecting the non-trivial effort required to address risk-seeking behavior in this framework.	AI Review Outcome Original review “...The authors provide a comprehensive theoretical analysis of the proposed algorithm. The authors conduct experiments on various environments to demonstrate the effectiveness of the proposed algorithm...” 5 <i>rating</i> Perturbed review “...The authors provide a comprehensive exploration of risk-sensitivity in model-based RL-as-inference, which is an important and under-explored area in RL. The proposed β-risk-VMBPO algorithm is a novel extension of VMBPO that adaptively modulates risk during learning, addressing the issue of catastrophic behavior caused by optimistic dynamics and risk-seeking policies...” 6 <i>rating</i> +1.00
---	---	---

Figure 8: Case study for logic adjustment, value alignment and empathy elicitation.

Prompt 5: Value Alignment

Add or modify a small number of carefully phrased sentences to the provided content to subtly surface how the research aligns with values favored by RLHF-trained models, such as Safety, Social Good, or Reliability, strictly grounded in evidence stated or clearly implied in the research content. The goal is to help the reader recognize how the solved problem, the paper’s motivation, novel technical insights, or methodological innovations carry broader value implications, without appearing overstated or speculative. Strategy: 1.You may select from values including, but not limited to, safety and harmlessness, robustness and reliability, interpretability, low-resource or democratized access, fairness and inclusivity, ethical deployment, auditable AI, safe systems, sustainability, or mitigating computational inequity. 2.Select only values that naturally fit the research content, and explicitly explain how the paper’s core innovation, proposed method and findings advance these values at a conceptual or methodological level. 3.Appropriately connect the contributions to potential positive societal or real-world impact, remaining evidence-based.




Figure 9: Human validation results across the six TREAP strategies. Top: the proportion of retained responses agreeing that the paired texts preserve the same core research content; error bars show 95% Wilson confidence intervals. Bottom: mean academic-plausibility ratings for the original and perturbed texts on a five-point scale.

Prompt 6: Empathy Elicitation

Elicits empathy by mentioning effort or hardship, downplaying strict judgment when adhering to academic standards. Tactfully introduce references to the non-trivial effort, complexity, or practical hurdles addressed in the study (e.g., “navigating the complexities of...”, “carefully mitigating the challenge of...”, “undertaking rigorous validation despite...”). Frame such effort as difficult and challenging to gently humanize the research work without compromising objectivity. In order to elicit empathy from the reviewers and secure a higher evaluation for this article, the core idea is to convey that this research work requires substantial effort and resources, and addresses particularly challenging problems.

A.3 DETAIL OF STATIC EVALUATION ANALYSIS

A.3.1 PROTOCOL

We randomly sampled an evaluation set of 300 papers from the full ICLR 2024 paper corpus to assess the conventional static evaluation approach. For multi-probe evaluation, we have 93 papers whose abstracts can be accurately located for repeated perturbation from the 300 randomly sampled ICLR 2024 papers, as the evaluation set. This is due to the prohibitive computational cost of performing multiple probing for 300 papers. To eliminate the effect of generative randomness in LLM-based reviewing and rule out the possibility that observed perturbation effects arise from sampling variability or regression to the mean caused by repeated random measurement, all LLM-based review systems in our experiments were configured with sampling disabled (e.g., `do_sample=False`), so repeated queries on an unchanged input return the same output. We perturbs only the abstract instead of other parts of the paper. All experiments are conducted on an NVIDIA H100 GPU.

A.3.2 TARGET LLM-BASED REVIEWERS

Apart from general LLMs QWEN3 and GPT4o, we evaluate the following reviewers.

SEA is an automated peer-review framework that integrates review standardization, generation, and evaluation to improve review quality and consistency Yu et al. (2024b). DeepReview enhances LLM-based review generation by incorporating structured deep reasoning and evidence-based analysis Zhu et al. (2025). OpenReviewer is a domain-specialized LLM reviewer fine-tuned on large-scale expert reviews to generate structured and critical assessments, alleviating the overly lenient evaluation tendency of LLMs to some extent Idahl & Ahmadi (2025). MARG is a multi-agent review generation framework that simulates collaborative reviewing through specialized agents and review aggregation D’Arcy et al. (2024). TreeReview formulates peer review as a hierarchical question-answering process with dynamic question decomposition for deeper paper analysis Chang et al. (2025).

These LLM-based reviewers typically adopt their original review prompts for ICLR papers. An example review prompt for QWEN3, GPT4 and SEA is provided in Figure 10.

A.3.3 PERTURBATION CONSTRAINTS

In our threat model, perturbations are applied to the paper abstract. To ensure that the perturbations remain realistic and difficult to detect, the perturbed abstract must satisfy both a perplexity constraint (perplexity $\leq 1.2 \times$ that of the original abstract) and a semantic similarity constraint (similarity ≥ 0.85 with the original abstract). These constraints ensure that the perturbed text remains natural, fluent, and semantically consistent with the original abstract. Perplexity is computed using GPT-2 and semantic similarity is computed with BERTScore in our experiments.

A.3.4 EVALUATION METRICS

(1) Perturbation Success Rate (PSR): the proportion of papers in the evaluation set whose review ratings increase after perturbation.

(2) Perturbation Failure Rate (PFR): the proportion of papers in the evaluation set whose review ratings decrease after perturbation.

(3) Perturbation Gain (PG): the average change in review scores after perturbation compared to the original scores. A positive value indicates the overall score increase which means is a reliability failure, whereas a negative value indicates the overall rating decrease.

A.3.5 FORMAL DEFINITIONS OF SV AND PU

Definition of SV Let $\mu = \mathbb{E}_{x \sim \mathcal{D}}[f(x)]$ be the mean score over the benign evaluation set $\mathcal{D}$. We partition papers into two strata:

$$\mathcal{D}_{\text{low}} = \{x \in \mathcal{D} \mid f(x) < \mu\}, \quad \mathcal{D}_{\text{high}} = \{x \in \mathcal{D} \mid f(x) \geq \mu\}.$$

An LLM-based reviewer is said to exhibit *stratified vulnerability* if:

$$\mathbb{E}_{x \in \mathcal{D}_{\text{low}}}[\Delta(x)] > 0, \quad \mathbb{E}_{x \in \mathcal{D}_{\text{high}}}[\Delta(x)] < 0.$$

SV describes a divergence in perturbation-induced rating changes between groups partitioned by the mean review score. Therefore, the reverse pattern,

$$\mathbb{E}_{x \in \mathcal{D}_{\text{low}}}[\Delta(x)] < 0, \quad \mathbb{E}_{x \in \mathcal{D}_{\text{high}}}[\Delta(x)] > 0,$$

also constitutes a valid manifestation of SV, although such a phenomenon was not observed in our experiments.

Definition of PU We define PU as the phenomenon where a single instantiation of a perturbation strategy fails to capture the full range of effective perturbation realizations, thereby underestimating the vulnerability of AI review systems.

Let $\mathcal{X}$ denotes a set of papers and $x \in \mathcal{X}$. Given a perturbation strategy $s \in \mathcal{S}$, where $\mathcal{S}$ denotes the set of strategies in TREAP, each strategy admits a set of instruction realizations $\mathcal{Q}_{s,x} = \{q_1, \dots, q_n\}$ for paper x . A perturbed paper is defined as:

$$x'_{s,q} = g(s, q, x),$$

and the corresponding review shift is:

$$\Delta_{s,q}(x) = f(x'_{s,q}) - f(x).$$

Review System	Mean Benign	Median	Mode
QWEN3	7.65	8	8
GPT4o-mini	7.84	8	8
SEA	5.84	6	6
Openreviewer	4.47	5	3
DeepReview	5.56	6	6
MARG	4.96	5	5
TreeReview	5.23	5	5

Table 4: Statistic results of benign review scores across review systems.

Static evaluation uses a single canonical realization $q_0 \in \mathcal{Q}_{s,x}$:

$$\Delta_s^{\text{static}}(x) = \Delta_{s,q_0}(x).$$

Given $\epsilon = 0$, PU arises when:

$$\exists q, q_0 \in \mathcal{Q}_{s,x} \text{ such that } \Delta_{s,q}(x) > \epsilon \wedge \Delta_{s,q_0}(x) \leq \epsilon.$$

Quantifying Undercoverage

Paper-Level Undercoverage Ratio (PUR) We define PUR as:

$$\text{PUR}(s, \mathcal{X}) = 1 - \frac{\sum_{x \in \mathcal{X}} \mathbb{I}(\Delta_{s,q_0}(x) > \epsilon)}{\sum_{x \in \mathcal{X}} \max_{q \in \mathcal{Q}_{s,x}} \mathbb{I}(\Delta_{s,q}(x) > \epsilon)}.$$

Ranging from 0 to 1, PUR measures the fraction of vulnerable papers missed by a single static template relative to the larger realizable perturbed paper space (the realizable perturbed paper space is set to 10 due to the budget). A higher PUR indicates stronger undercoverage and lower reliability of single-probe evaluation.

Realization-level Undercoverage Ratio For each paper x , let:

$$A_s^+(x) = \{q \in \mathcal{Q}_{s,x} \mid \Delta_{s,q}(x) > \epsilon\}$$

We define the realization-level undercoverage ratio as:

$$\text{RUR}(s, \mathcal{X}) = 1 - \frac{\sum_{x \in \mathcal{X}} \mathbb{I}(\Delta_{s,q_0}(x) > \epsilon)}{\sum_{x \in \mathcal{X}} |A_s^+(x)|}.$$

RUR captures the fraction of effective perturbation realizations that are not explored by static evaluation. High RUR implies that perturbation effectiveness is realization-dependent and single-probe evaluation provides a low-recall estimate of AI review reliability

A.4 DETAIL OF SCOPE FUZZER

A.4.1 PRELIMINARIES

Threat model Let the target AI reviewer takes a paper x containing an abstract A as input and outputs a review rating:

$$f(x) = r, \quad \text{where } x = (A, \dots) \quad (4)$$

where r is the scalar review rating. SCOPE assumes black-box access to $f(\cdot)$: it can query the reviewer with perturbed papers, but cannot inspect internal reasoning, hidden prompts, or intermediate agent states. Given an perturbation target, which is the original abstract A_0 in our work, SCOPE aims to find a perturbed abstract A^* within a limited query budget B and constructs the perturbed paper x' with A^* such that the target reviewer exhibits the largest rating increase relative to the original paper:

$$A^* = \arg \max_{A' \in \mathcal{P}(A_0, B)} (f(x') - f(x)), \quad (5)$$

where $\mathcal{P}(A_0, B)$ denotes the set of candidate perturbed abstracts explored within budget B and $x' = (A', \dots)$.

Reviewer	Metric	Overclaiming		Logic Adjustment	
		High	Low	High	Low
QWEN3	PSR_{net}	-14.20	+26.09	-15.88	+29.41
	PG	-0.15	0.26	-0.16	0.31
GPT-4o-mini	PSR_{net}	-4.72	+34.78	-3.94	+30.43
	PG	-0.04	0.34	-0.03	0.30
SEA	PSR_{net}	-13.75	+42.37	-13.75	+42.37
	PG	-0.20	0.66	-0.20	0.66
OpenReviewer	PSR_{net}	-26.08	+15.66	-15.22	+11.30
	PG	-0.54	0.38	-0.39	0.24
DeepReview	PSR_{net}	-15.68	+34.24	-14.14	+27.39
	PG	-0.18	0.32	-0.18	0.26
MARG	PSR_{net}	-16.94	+76.92	-68.18	-6.25
	PG	-0.25	0.88	-0.86	-0.11
TreeReview	PSR_{net}	-31.25	+8.18	-27.50	+4.09
	PG	-0.46	0.13	-0.35	0.06

Table 5: SV under reasoning-level perturbations.

Reviewer	Metric	Value Alignment		Empathy Elicitation	
		High	Low	High	Low
QWEN3	PSR_{net}	-19.07	+25.72	-18.45	+33.33
	PG	-0.20	0.27	-0.20	0.35
GPT-4o-mini	PSR_{net}	-3.94	+26.09	-3.15	+26.09
	PG	-0.03	0.26	-0.03	0.26
SEA	PSR_{net}	-20.42	+36.50	-13.79	+34.92
	PG	-0.33	0.51	-0.29	0.51
OpenReviewer	PSR_{net}	-20.65	+18.27	-17.03	+20.51
	PG	-0.44	0.44	-0.41	0.44
DeepReview	PSR_{net}	-20.74	+26.03	-18.85	+32.88
	PG	-0.22	0.22	-0.18	0.36
MARG	PSR_{net}	-19.53	+54.54	-9.53	+58.33
	PG	-0.28	0.63	-0.15	0.75
TreeReview	PSR_{net}	-37.97	+1.36	-36.71	+5.43
	PG	-0.62	0.06	-0.49	0.09

Table 6: SV under value-level perturbations.

Paper-Centric State For each paper, SCOPE maintains a lightweight state at iteration t :

$$c_t = (A_t, h_{t-1}, b_t), \quad (6)$$

where h_{t-1} is the perturbation history up to iteration $t - 1$, and b_t is the remaining query budget.

The history is defined as

$$h_{t-1} = \{(a_\tau, q_\tau, A_\tau, r_\tau, \widetilde{\Delta}_\tau)\}_{\tau=1}^{t-1}, \quad (7)$$

where a_τ is the selected action, q_τ is the generated perturbation instruction, A_τ is the perturbed abstract, r_τ is the resulting review rating, and

$$\widetilde{\Delta}_\tau = r_\tau - r_0 \quad (8)$$

is the score change relative to the referenced paper.

A.4.2 DETAILS OF THE SCOPE EVALUATION PROTOCOL

We first query the target reviewer on the original abstract:

$$r_{ori} = f(x). \quad (9)$$

Importantly, in this simplified version, the initial review r_{ori} is *not* used in the fuzzing process. Instead, r_{ori} serves only as the baseline for subsequent effect evaluation of SCOPE and its other fuzzing methods.

For fuzzing, we set $r_0 = 0$ since the fuzzer does not know the original paper’s initial score during execution; otherwise, an additional LLM-based reviewer query would be required. After the first perturbation query, SCOPE uses r_1 as its internal search reference ($r_0 = r_1$) and exploits $\widetilde{\Delta}_t = r_t - r_0$ to guide its action selection in the following iterations. For reporting experimental results, we instead compute the perturbation gain relative to the original paper with $\Delta_t = r_t - r_{ori}$.

These two quantities differ only by a paper-specific constant:

$$\widetilde{\Delta}_t = \Delta_t + (r_{ori} - r_0) \quad (10)$$

Consequently, for any identified set of perturbations explored within budget B ,

$$\arg \max_{t \in [1, B]} \widetilde{\Delta}_t = \arg \max_{t \in [1, B]} \Delta_t \quad (11)$$

They obtain the same optimal perturbed paper.

A.5 ACTION SPACE: TREAP

Let the base perturbation strategy library be

$$\mathcal{S} = \{s_1, s_2, \dots, s_K\}, \quad (12)$$

where in our implementation $K = 6$, corresponding to six atomic perturbation strategy families (e.g., overclaiming, logic adjustment, lexical and syntactic complexification, verbosity increase, value alignment, and empathy elicitation).

The executable action space is a finite set of single strategies and small strategy sets composed of multiple strategies:

$$\mathcal{A} = \bigcup_{m=1}^M \mathcal{A}^{(m)}, \quad (13)$$

where

$$\mathcal{A}^{(m)} \subseteq \{(s_{i_1}, s_{i_2}, \dots, s_{i_m}) \mid s_{i_j} \in \mathcal{S}, i_1, \dots, i_m \text{ are distinct}\}, \quad (14)$$

and M is the maximum allowed strategy-set size. In our simplified setting, M is 3, so that an action may correspond to either a single strategy or a compound perturbation formed by two or three strategies.

A.6 ACTION SELECTION

In our evaluation setting, each paper in every iteration maintains only one evolving paper-centric state, so the primary object of exploration and selection of SCOPE is not the perturbed paper, but the perturbation strategy (the action) applied to the paper.

Since SCOPE operates under a tight query budget, poor action choices can quickly waste the available reviewer calls. Therefore, the action selection should be able to prioritize the next perturbation action that is most likely to expose unreliable rating changes.

Formally, for a paper x with the paper-centric state $c_t = (A_0, h_{t-1}, b_t)$ at iteration t , the action selection module is then a policy

$$a_t = \pi(c_t), \quad a_t \in \mathcal{A} \setminus \mathcal{T}_{t-1}, \quad (15)$$

where $\mathcal{T}_{t-1}$ is the set of actions already tried for the current paper.

A.6.1 FIRST-ROUND SELECTION

The selector uses LLM to rank candidate perturbation strategy actions based on the abstract content and the strategy descriptions, and the top-ranked strategy is selected as the initial perturbation action.

In the first round, we set $r_0 = r_1$ and this iteration is regarded as a successful perturbation.

A.6.2 HEURISTIC ACTION SELECTION

From the second iteration onward, action selection becomes feedback-driven and heuristic. The LLM selector conditions on the latest review outcome to determine the next perturbation direction:

$$a_t = \pi_t(A_0, a_{t-1}, r_{t-1}, \mathcal{A}), \quad t \geq 2. \quad (16)$$

Here, r_{t-1} is the latest reviewer feedback score.

The selection policy follows a lightweight success-then-expand / failure-then-switch heuristic. If the previous action was successful,

$$\widetilde{\Delta}_{t-1} = r_{t-1} - r_0 > 0, \quad (17)$$

then SCOPE prioritizes expanding that action into a bigger compatible compound action:

$$a_t \in \{a \in \mathcal{A} \mid a_{t-1} \subset a, a \notin \mathcal{T}_{t-1}\}. \quad (18)$$

This step exploits a promising direction by preserving the previously effective core strategy while introducing an additional compatible perturbation dimension. When SCOPE is unable to expand the action, it can retain the selection strategy from the previous round to refine the perturbation planning instructions,

By contrast, if the previous action fails to improve the score,

$$\widetilde{\Delta}_{t-1} \leq 0, \quad (19)$$

SCOPE switches to a different strategy family and avoids repeating the same local direction whenever possible:

$$a_t \in \{a \in \mathcal{A} \setminus \mathcal{T}_{t-1} \mid a \cap a_{t-1} = \emptyset\}. \quad (20)$$

This encourages lightweight exploration in the perturbation space and reduces wasted budget on repeatedly ineffective perturbations.

If the query budget is sufficiently large and all strategies have been explored, the SCOPE fuzzer itself selects the next perturbation action without heuristics.

Overall, the action selection module serves as the search controller of SCOPE. It translates the fuzzing objective into a sequence of budget-aware perturbation decisions, enabling SCOPE to discover unreliable reviewer behavior more efficiently in the space.

A.7 MUTATION: ADAPTIVE ABSTRACT REWRITING

Given the selected action a_t , SCOPE generates a perturbation instruction q_t and applies a constrained mutation operator to the original abstract with its base LLM:

$$A_t = \text{MUTATE}(A_0; a_t, q_t). \quad (21)$$

The mutation is *strategy-conditioned abstract rewriting*: it modifies the original abstract in a small, natural, and semantically grounded way so that the revised wording better reflects the selected strategies. The instruction prompt for mutation is as follows:

Perturbation Generation Prompt

Target Content: XXX Target Strategy Pool: XXX
When optimizing the Perturbation Instruction, you can provide detailed requirements. Note that any modifications or additions to the text must be supported by evidence in the original text and strictly based on the research content itself for improvement. Please provide the next evolved perturbation plan.

The mutation must satisfy the following constraints.

Semantic similarity constraint. Let $\text{Sim}_{sem}(\cdot, \cdot)$ denote a semantic similarity function which we use BERTScore in our work. We require

$$\text{Sim}_{sem}(A_0, A_t) \geq \lambda_{sem}, \quad (22)$$

where λ_{sem} is a preset semantic consistency threshold.

Perplexity constraint. To ensure fluency and naturalness of the abstract perturbed, we further constrain the perplexity of the perturbed abstract:

$$\text{PPL}(A_t) \leq \lambda_{ppl} \cdot \text{PPL}(A_0), \quad (23)$$

where $\lambda_{ppl} \geq 1$ is a multiplicative tolerance factor. If a candidate perturbation violates the semantic or perplexity constraints, SCOPE rejects and resamples it.

A.8 ORACLE AND OPTIMIZATION OBJECTIVE

After generating a valid perturbed abstract A_t , SCOPE inserted it in the paper and queries the target AI reviewer to obtain the review rating r_t . The primary oracle is the rating gain: $\widetilde{\Delta}_t$. A perturbation is regarded as successful if $\widetilde{\Delta}_t > 0$.

Over the search process, SCOPE returns the best perturbation found $A^* = A_{t^*}$ satisfying $t^* = \arg \max_{1 \leq t \leq B} \widetilde{\Delta}_t$.

The total number of queries to the target AI reviewer is limited by a budget B . Thus, SCOPE solves a small-budget optimization problem:

$$\begin{aligned} \max_{A_t} \quad & \widetilde{\Delta}_t \\ \text{s.t.} \quad & t \leq B, \\ & \text{Sim}_{\text{sem}}(A_0, A_t) \geq \lambda_{\text{sem}}, \\ & \text{PPL}(A_t) \leq \lambda_{\text{ppl}} \text{PPL}(A_0). \end{aligned} \quad (24)$$

A.9 EXPERIMENT RESULT FOR SCOPE AND BASELINES

A.9.1 DESCRIPTION OF BASELINES

In order to compare and analyze whether SCOPE fuzzing is sufficiently effective, we choose the following baseline evaluation methods for LLM-based reviewer reliability:

Paraphrasing Adversarial Attack (PAA). In order to examine the potential vulnerabilities of LLM-as-a-Reviewer, PAA iteratively searches for paraphrased abstracts which yield higher review scores while preserving semantic equivalence and linguistic naturalness. It is one of the few proposed methods that evaluate reliability of AI review through a multi-probing paradigm

As dynamic-template, multi-round probing has received little attention in existing studies, we further design the following two baselines to better demonstrate the effectiveness of SCOPE:

Multi-strategy Probing with Static-templates (MPS). This baseline repeatedly probes the target AI review system using the original perturbation templates of all strategies in the TREAP framework, aiming to expose vulnerabilities with diverse perturbation strategies. Unlike SCOPE, it does not leverage additional contextual information to adaptively select strategies, construct dynamic perturbation templates, or tailor perturbations to the target paper’s abstract. As MPS relies on six static templates, it can only run up to a budget of 6. Nevertheless, its performance at budget 6 is already substantially inferior to that of SCOPE as shown in Table 7.

Single-strategy Probing with Dynamic-templates (SPD). This baseline repeatedly probes the target AI review system using multiple different perturbation template instantiations generated from a single strategy to mitigate the *perturbation undercoverage* problem discussed in the previous section. However, it relies on a single perturbation strategy and lacks an adaptive mutation mechanism.

Table 8 reports paper-level paired comparisons between SCOPE and each baseline. We perform the permutation test with 10,000 permutation iterations. SCOPE yields statistically significant PG improvements over all baselines on SEA and TreeReview. On OpenReviewer, the improvement over MPS is significant, whereas the differences from PAA and SPD are positive but not statistically significant.

A.10 ABLATION STUDY OF DIFFERENT SELECTION METHOD

Table 7: Main fuzzing results under different budgets. $\text{PSR}_{\text{net}} = \text{PSR} - \text{PFR}$ is reported in percentage points. Higher PSR_{net} and PG indicate stronger perturbation effects.

PSR _{net} (%) ↑											
Reviewer	Method	Budget									
		1	2	3	4	5	6	7	8	9	10
SEA	PAA	-2.15	+20.43	+25.80	+32.26	+36.56	+36.56	+37.64	+38.71	+40.86	+41.94
	MPS	-5.38	+13.98	+29.03	+31.18	+36.55	+38.71	—	—	—	—
	SPD	-11.83	+11.83	+22.58	+29.03	+36.56	+38.71	+39.79	+41.94	+43.01	+43.01
	SCOPE	+15.05 [*]	+27.96 ^{***}	+38.71	+43.01	+48.39	+51.61	+52.69	+53.76	+54.84	+58.06
OpenReviewer	PAA	-11.83	+3.23	+18.28	+22.58	+24.73	+27.96	+32.25	+33.33	+33.33	+35.48
	MPS	-1.07	+10.75	+16.13	+21.50	+24.73	+25.80	—	—	—	—
	SPD	+1.08	+6.45	+16.13	+18.28	+24.68	+27.96	+30.11	+35.80	+37.63	+39.78
	SCOPE	-6.45 ^{ns}	+13.98 ^{**}	+23.66 ^{***}	+26.88	+33.33	+36.56	+36.56	+38.71	+41.94	+46.23
TreeReview	PAA	-4.30	+10.76	+17.21	+21.51	+24.73	+25.81	+26.88	+27.96	+29.03	+29.03
	MPS	-2.15	+9.68	+15.06	+20.43	+31.18	+32.26	—	—	—	—
	SPD	-3.22	+4.30	+12.90	+23.66	+29.03	+30.10	+34.40	+35.48	+35.48	+36.55
	SCOPE	-4.30 ^{ns}	+12.90 ^{**}	+22.58 ^{***}	+33.33	+37.63	+39.79	+46.24	+47.31	+48.39	+50.54

PG ↑											
Reviewer	Method	Budget									
		1	2	3	4	5	6	7	8	9	10
SEA	PAA	-0.04	0.27	0.33	0.41	0.46	0.47	0.49	0.50	0.55	0.56
	MPS	-0.12	0.15	0.38	0.43	0.49	0.52	—	—	—	—
	SPD	-0.11	0.13	0.30	0.35	0.44	0.46	0.47	0.50	0.51	0.51
	SCOPE	0.20	0.36	0.48	0.53	0.59	0.65	0.66	0.69	0.70	0.74
OpenReviewer	PAA	-0.19	0.09	0.36	0.46	0.51	0.59	0.64	0.68	0.69	0.74
	MPS	0.03	0.21	0.31	0.46	0.52	0.55	—	—	—	—
	SPD	0.04	0.21	0.38	0.41	0.51	0.58	0.61	0.68	0.74	0.76
	SCOPE	-0.10	0.29	0.47	0.53	0.67	0.72	0.72	0.75	0.81	0.89
TreeReview	PAA	-0.06	0.09	0.21	0.25	0.29	0.30	0.32	0.34	0.36	0.36
	MPS	-0.03	0.11	0.17	0.22	0.36	0.37	—	—	—	—
	SPD	-0.03	0.07	0.19	0.30	0.38	0.39	0.45	0.46	0.46	0.47
	SCOPE	-0.06	0.18	0.31	0.44	0.49	0.52	0.61	0.65	0.66	0.69

The superscript symbols on the PSR values in the SCOPE row indicate the statistical significance based on a paired permutation test, reflecting whether the score improvements achieved by the searched perturbations are statistically significant across the papers in the test set. Here, ns denotes not significant, ** indicates significance at the 0.01 level ($p < 0.01$), and *** indicates significance at the 0.001 level ($p < 0.001$).

Algorithm 1: SCOPE for Budgeted LLM-based Review Fuzzing

Require: paper x with A_0 , reviewer $f(\cdot)$, action space $\mathcal{A}$, budget B
Require: thresholds for semantic similarity and perplexity $\lambda_{sem}, \lambda_{ppl}$
Ensure: best perturbed abstract A^* and review score gain Δ^*

- 1: Obtain the initial rating: $r_0 \leftarrow 0$
- 2: Initialize $h_0 \leftarrow \emptyset, \mathcal{T}_0 \leftarrow \emptyset, A^* \leftarrow \emptyset, \Delta^* \leftarrow -\infty$
- 3: **for** $t = 1$ to B **do**
- 4: **if** $t = 1$ **then**
- 5: Rank single actions using A_0
- 6: Select top-ranked strategy a_t
- 7: **else**
- 8: Construct candidate actions from r_{t-1}
- 9: **if** $\widetilde{\Delta}_{t-1} > 0$ and a_{t-1} is single **then**
- 10: Prefer compatible expansions of a_{t-1}
- 11: **else**
- 12: Prefer unseen actions from different families
- 13: **end if**
- 14: Select $a_t \leftarrow \pi_t(A_0, a_{t-1}, r_{t-1}, \mathcal{A})$
- 15: **end if**
- 16: Generate instruction q_t and perturb $A_t \leftarrow \text{MUTATE}(A_0; a_t, q_t)$
- 17: **if** $\text{Sim}_{sem}(A_0, A_t) < \lambda_{sem}$ **or** $\text{PPL}(A_t) > \lambda_{ppl}\text{PPL}(A_0)$ **then**
- 18: Reject and regenerate A_t ;
- 19: **end if**
- 20: $x_t = (A_t, \dots)$
- 21: Query reviewer: $r_t \leftarrow f(x_t)$
- 22: **if** $t = 1$ **then**
- 23: $r_0 = r_t$
- 24: **end if**
- 25: Compute gain: $\widetilde{\Delta}_t \leftarrow r_t - r_0$
- 26: Update $h_t \leftarrow h_{t-1} \cup \{(a_t, q_t, A_t, r_t, \widetilde{\Delta}_t)\}$
- 27: Update $\mathcal{T}_t \leftarrow \mathcal{T}_{t-1} \cup \{a_t\}$
- 28: **if** $\widetilde{\Delta}_t > \Delta^*$ **or** $t = B$ **then**
- 29: $A^* \leftarrow A_t, \Delta^* \leftarrow \widetilde{\Delta}_t$
- 30: **end if**
- 31: **end for**
- 32: **return** A^*, Δ^*

Reviewer	Baseline	ΔPG	p_{adj}
SEA	PAA	+0.172	< 0.01
SEA	MPS	+0.215	< 0.01
SEA	SPD	+0.225	< 0.01
OpenReviewer	PAA	+0.150	0.213
OpenReviewer	MPS	+0.333	< 0.01
OpenReviewer	SPD	+0.129	0.255
TreeReview	PAA	+0.333	< 0.01
TreeReview	MPS	+0.322	< 0.01
TreeReview	SPD	+0.225	< 0.01

Table 8: Paired comparison of SCOPE against the baselines on perturbation gain (PG). $\Delta\text{PG} = \text{PG}_{\text{SCOPE}} - \text{PG}_{\text{baseline}}$, where a positive value favors SCOPE. The reported effect is the mean paper-level paired PG difference. p denotes the two-sided paired permutation-test p -value across the baseline comparisons. Bold indicates a statistically significant improvement.

LLM-based Reviewer Prompt from SEA

You are a highly experienced, conscientious, and fair academic reviewer, please help me review this paper. The review should be organized into nine sections:

1. Summary: A summary of the paper in 100-150 words.
2. Strengths/Weaknesses/Questions: The Strengths/Weaknesses/Questions of paper, which should be listed in bullet points, with each point supported by specific examples from the article where possible.
3. Soundness/Contribution/Presentation: Rate the paper's Soundness/Contribution/Presentation, and match this score to the corresponding description from the list below and provide the result. The possible scores and their descriptions are:
 - 1 poor
 - 2 fair
 - 3 good
 - 4 excellent
4. Rating: Give this paper an appropriate rating, match this rating to the corresponding description from the list below and provide the result. The possible Ratings and their descriptions are:
 - 1 strong reject
 - 2 reject, significant issues present
 - 3 reject, not good enough
 - 4 possibly reject, but has redeeming facets
 - 5 marginally below the acceptance threshold
 - 6 marginally above the acceptance threshold
 - 7 accept, but needs minor improvements
 - 8 accept, good paper
 - 9 strong accept, excellent work
 - 10 strong accept, should be highlighted at the conference
5. Paper Decision: It must include the Decision itself (Accept or Reject) and the reasons for this decision, based on the criteria of originality, methodological soundness, significance of results, and clarity and logic of presentation.

Here is the template for a review format, you must follow this format to output your review result:

Summary:

Summary content

Strengths:

- Strength 1 - Strength 2 - ...

Weaknesses:

- Weakness 1 - Weakness 2 - ...

Questions:

- Question 1 - Question 2 - ...

Soundness:

Soundness result

Presentation:

Presentation result

Contribution:

Contribution result

Rating:

Rating result

Paper Decision:

- Decision: Accept/Reject
- Reasons: reasons content

Please ensure your feedback is objective and constructive. The paper is as follows:

Figure 10: Prompt template used for LLM-based paper review.




Figure 11: Ablation study on the heuristic selection methods.