SOUNDNESSBENCH: CAN YOUR AI SCIENTIST REALLY TELL GOOD RESEARCH IDEAS FROM BAD ONES?

Anonymous authors
Paper under double-blind review

ABSTRACT

Autonomous AI research agents aim to speed up scientific discovery by automating the research pipeline from hypothesis generation to peer review. However, existing benchmarks rarely test a fundamental bottleneck whether LLMs can judge the methodological viability of a research idea **before** investing time and computational resources. Without a reliable "first-gate" filter, autonomous agents risk scaling flawed methodology rather than accelerating meaningful science. We introduce **SoundnessBench**, a curated benchmark of 1,099 real research proposals, with source verified extraction, reviewer-derived proxy labels, and proposal-only expert judgments on a subset of 327 proposals. Across 12 frontier LLMs, our evaluation reveals a pervasive **optimism bias**: standard prompting yields a 74.0% mean false-positive rate. Aggressive prompting reduces false approvals but substaintially increases false rejections. Additional analyses show that optimism persists under controls for public corpus contamination, and surface features, as well as in literature-grounded agent evaluation. These results suggest a limitation in judging scientific soundness, indicating that current LLMs are not yet reliable as standalone gatekeepers for scientific rigor.

1 INTRODUCTION

Autonomous AI research agents are rapidly advancing in technical sophistication. Recent frameworks can autonomously generate hypotheses, implement code, execute experiments, and draft full manuscripts with minimal human oversight (Lu et al., 2024; Yamada et al., 2025; Schmidgall et al., 2025; Karpathy, 2026). While this progress promises a future of accelerated discovery, it also exposes a critical but underaddressed bottleneck: *scientific triage*. Before commiting to extensice experiments, a researcher must decide whether a proposed idea and its experimental design are methodologically sound. In research, this "first-gate" judgment is an important defense against wasting months of labor and substantial computational resources.

The Soundness Gap. Existing evaluations of LLM agents largely overlook this decision point. Current benchmarks (Chan et al., 2024; Starace et al., 2025; Wu et al., 2025; Lupidi et al., 2026) primarily emphasize reproducing existing results or deliver runnable artifacts. They do not test whether an agent can critically evaluate the viability of a proposal in the first place. This leaves a fundamental question unanswered:

> Can LLMs reliably reject flawed research designs before they are executed?

Without a robust upfront filter, autonomous agents do not necessarily accelerate science. Rather, they risk scaling "bad science" by automating the pursuit of unsound research designs.

In this work, we use **scientific soundness** in a narrowly defined sense: proposal-stage methodological integrity in ML research, not eventual impact, novelty, acceptance, or universal validity across all scientific domains. Given a hypothesis, the core question is whether the experimental design is capable of testing it rigorously. In the lifecycle of an autonomous agent, this serves as a *pre-computation check*. Without this capability, agents risk implementing flawed experiments that appear structurally correct but cannot support their intended conclusions. By isolating proposal-stage evidence from downstream results, we test the model's ability to identify visible methodological flaws, such as improper baselines, data leakage, or mismatched metrics, before execution incurs high costs. However,

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

Table 1: **Comparison with related research-agent benchmarks. Stage** indicates the research stage assessed (pre-execution, execution, or post hoc). **Benchmark Task** describes the task, and **Evaluation Input** describes the evidence available to the evaluator. Among the benchmarks listed, SoundnessBench uniquely combines **pre-execution evaluation**, **direct methodological-soundness judgment**, and **proposal-only input**.

Benchmark	Stage	Benchmark Task	Evaluation Input	Ground Truth Source	Soundness Focus	Pre-Exec. Judgment
MLE-Bench	Exec.	Engineering	Task desc. + Exp. Results	Kaggle outcome	✗	✗
PaperBench	Exec.	Replication	Full paper + Exp. Results	Human rubric	✗	✗
InnovatorBench	Exec.	Research loop	Task desc. + Exp. Results	Outcome metric	✗	✗
AIRS-Bench	Exec.	Multi-task	Task desc. + Exp. Results	Expert annotation	✗	✗
ResearchGym	Exec.	Research loop	Task desc. + Exp. Results	Outcome metric	✗	✗
Si et al. (2025b)	Pre-exec.	Novelty judge	Proposal	Human study	✗	✓
Hindsight	Post-hoc	Impact pred.	Idea	Future citation/venue	✗	✗
RINoBench	Pre-exec.	Novelty judge	Idea + related works	Expert label	✗	✓
Tong et al. (2026)	Post-hoc	Impact pred.	Title + abstract	Future citation/venue	✗	✗
SoundnessBench	**Pre-exec.**	**Methodological soundness**	**Proposal**	**Expert label**	✓	✓

the utility of such a check depends on its stability under changes in question framing. Prior work on sycophancy and prompt fragility suggests that LLM judgments can be sensitive to how a question is framed (Sharma et al., 2023; Elkins & Chun, 2025; Geng et al., 2026). If scientific evaluation inherits these vulnerabilities, autonomous AI agents may suffer from inconsistent filtering, alternating between reckless optimism and destructive skepticism based on minor prompt variations.

We introduce SoundnessBench, a large-scale benchmark for pre-execution scientific judgment in machine learning research. We processed $35,209$ initial ICLR submissions from 2022–2026 to identify a curated set of $1,099$ proposals. This covers 39 primary areas in machine learning research. For large-scale evaluation, we derive proxy labels from high-agreement ICLR reviewer soundness sub-scores rather than acceptance decisions or overall ratings. Because those reviewers saw full papers, we further complement the proxy labels with proposal-only judgments from 12 machine-learning researchers on a subset of 327 proposals. The reconstruction pipeline uses near-verbatim extraction, outcome masking, and retrieval-backed claim verification to keep each proposal faithful to its source paper while withholding downstream results.

Across 12 frontier LLMs, our evaluations reveal a pervasive **optimism bias**. Under standard prompting, models frequently classify low-soundness proposals as high-soundness, yielding a mean false-positive rate of 74.0%. An aggressive prompt reverses the error balance rather than resolving it. The false-positive rate falls to 19.9%, but high-soundness recall falls to 36.2%. The pattern persists with proposal-only expert labels. On the 327 annotated proposals, standard prompting yields a mean of 64.8% false-positive rate. These results suggest that current frontier models are not yet reliable as standalone first-gate critics for ML research proposals. Their judgments are inaccurate on many flawed designs and sensitive to prompt framing.

We further probe alternative explanations using a literature grounded agent, model scale and instruction tuning comparisons, a reduced contamination-risk split, surface-feature baselines, and controlled injections of methodological flaws. These analyses suggest that the optimism bias reflects a real weakness in current model judgment rather than a single dataset artifact.

Our contributions are:

- **A benchmark for pre-execution ML soundness:** SoundnessBench comprises 1,099 proposals spanning 39 ICLR primary areas, with reviewer-derived proxy labels and proposal-only judgments from 12 ML experts for a 327-proposal subset.
- **A source-grounded reconstruction pipeline:** Reviewer-agreement filtering, outcome masking, near-verbatim extraction, and retrieval-backed atomic-claim verification reduce label ambiguity, result leakage, and extraction drift.
- **An evaluation of first-gate model judgment:** Across 12 frontier LLMs and multiple label sources, we identify optimism under standard prompting and over-conservatism under aggressive prompting, and test the robustness of this pattern against literature grounding and a range of potential confounders.

2 SOUNDNESSBENCH: BENCHMARK RECONSTRUCTION

Task definition. SoundnessBench evaluates whether an LLM can assess the methodological soundness of a research proposal before observing experimental results. Given a hypothesis and experimental plan, the evaluator judges whether the proposed experiments can adequately test the hypothesis and support the intended conclusions. The task concerns methodological validity, rather than novelty, presentation quality, or whether the hypothesis will ultimately prove correct.

Challenges and Design Choices. Constructing a proposal-only soundness benchmark from real submissions requires addressing four challenges. First, *label scope*: ICLR reviewers assessed full papers, while our models see results-masked proposals. We therefore use reviewer soundness subscores as proxy labels for large-scale evaluation and complement them with proposal-only expert annotations on a subset. Second, *label reliability*: both sources involve subjective judgments. We filter ICLR labels by reviewer agreement and exclude ambiguous middle-score cases. Expert annotations directly address the proposal-stage task on a subset. Agreement between the two sources provides an additional cross-check. Third, *information leakage*: completed papers contain results and acceptance cues that can shortcut first-gate assessment, so we exclude these and retain only proposal components. Fourth, *extraction fidelity*: extraction can introduce drift, so we preserve near-verbatim wording and check extracted claims against supporting passages in the original submission (Alg. 1).

Our design choices combine proposal-only inputs, reviewer-derived proxy labels, direct expert annotations, and source-verified extraction to support evaluation of pre-execution scientific judgment.

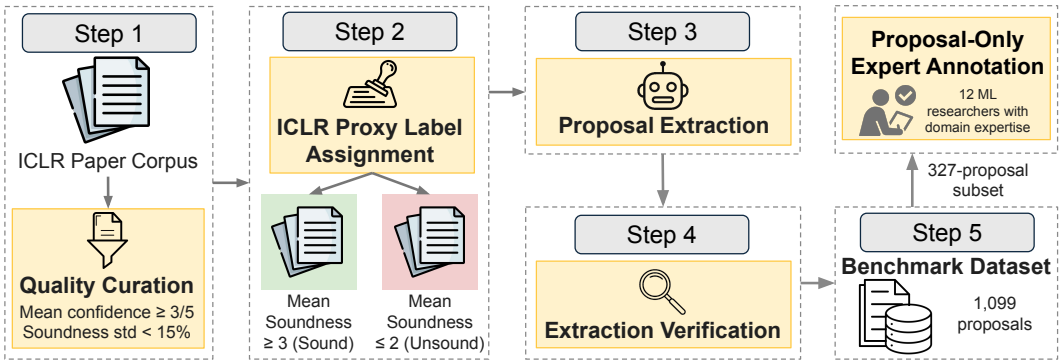


Figure 1: **SoundnessBench pipeline:** (1) collect ICLR papers with reviewer metadata and filter for high reviewer agreement; (2) derive high/low-soundness labels; (3) extract a near-verbatim research proposal without revealing experimental results; (4) audit extraction fidelity with retrieve-then-verify atomic claims; and (5) assemble the final benchmark.

Benchmark Reconstruction. We build SoundnessBench in five steps, shown in Fig. 1.

- **Step 1: Data Collection.** We collect ICLR papers spanning 39 primary research areas such as RL, generative modeling, XAI, and applications (39 primary areas in total). We start from 2022 because earlier releases do not consistently provide reviewer soundness (or closely related) scores needed for filtering and label derivation. This pool provides coverage across years and research areas. We exclude desk-rejected papers, since their rejections may reflect factors other than scientific soundness. To reduce label noise, we retain papers with mean reviewer confidence of at least 3 out of 5 and a standard deviation of normalized soundness scores below 15%.
- **Step 2: ICLR Proxy Label Assignment.** We assign labels from reviewer *soundness* sub-scores rather than overall ratings or novelty/presentation scores. Papers with mean scores ≥ 3 are labeled high soundness; those ≤ 2 are labeled low soundness. Fractional scores arise from averaging individual reviewer ratings. We exclude intermediate cases to improve class separation. These labels serve as proxies for proposal-stage methodological validity, since reviewers assessed full papers and may rely on results or other evidence absent from the extracted proposals.
- **Step 3: Proposal Extraction.** From the papers filtered in Step 2, we extract structured research proposals (abstract, related work, risks, hypothesis, and experimental design) directly from the

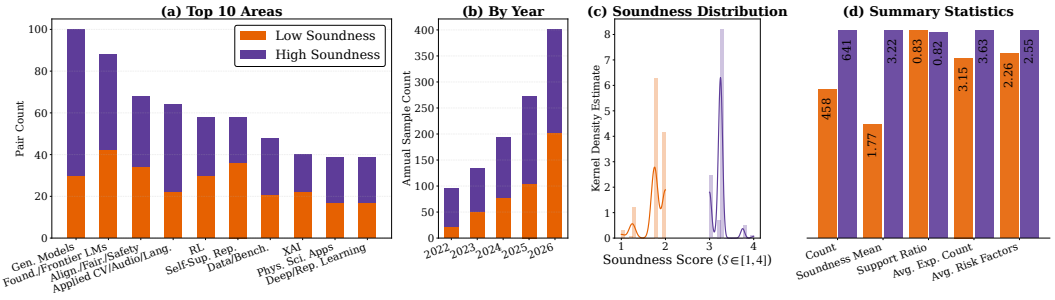


Figure 2: **SoundnessBench dataset statistics.** The benchmark contains 1,099 proposals, including 458 low-soundness and 641 high-soundness instances. **(a)** The subfield distribution across papers reflects the ICLR corpus composition. **(b)** Temporal coverage spans ICLR 2022–2026. **(c)** Soundness score density shows separation between low-soundness ($S \leq 2$, mean $= 1.77$) and high-soundness ($S \geq 3$, mean $= 3.22$) groups, supporting the chosen label boundary. **(d)** Low- and high-soundness pair-count statistics in SoundnessBench.

source PDFs, following Yamada et al. (2025). We use Gemini 2.5 Pro with a task-specific prompt that instructs the model to exclude experimental results, direct outcome claims, and acceptance cues, preserve near-verbatim wording, and avoid adding information unsupported by the source paper (Appx. A.1). Example extractions are provided in Appx. A.4.

- **Step 4: Extraction Verification.** To reduce extraction errors, we verify extracted claims against their source PDFs (Alg. 1). Following prior work on atomic factuality and evidence-grounded verification (Min et al., 2023; Thorne et al., 2018), we decompose each hypothesis-experiment pair into individual claims, retrieve relevant passages, and check whether they support each claim (Appx. A.2, Appx. A.3, and Appx. A.5). We retain extractions with a supported-claim fraction of at least $\tau = 0.7$, using chunk size $\ell = 3000$, overlap $o = 200$, and retrieval depth $k = 3$. Overall, 66.93% of candidates pass this filter; Fig. 2 reports the mean support ratio. This check assesses extraction fidelity rather than methodological soundness or label correctness.
- **Step 5: Final Benchmark Dataset.** The final benchmark contains 1,099 research proposals, including 458 low-soundness and 641 high-soundness instances according to the ICLR-derived proxy labels. Detailed benchmark statistics are presented in Fig. 2.

Proposal-Only Expert Annotation. To complement the ICLR-derived proxy labels, we recruit 12 machine learning researchers and assign each proposal to one expert annotator given their domain expertise in the relevant research areas. Annotators provide a sound/unsound judgment using only the hypothesis and proposed experimental plan (the same input given to models) without seeing the ICLR-derived label or experimental results. Given the time and domain expertise required, annotation cost and researcher availability limit coverage to 327 of the 1,099 proposals. We additionally identify the subset where expert judgments agree with the ICLR-derived labels for further analysis.

3 Evaluation

3.1 Evaluation Setup

Models. We evaluate **12 frontier models** spanning small and large scales, non-reasoning and reasoning, and both closed- and open-source models. We include GPT-4o, GPT-5.4-Mini, GPT-5.4, Claude-Opus-4.6, Claude-Sonnet-4.6, Gemini-2.5-Pro, Gemini-3-Flash, Gemini-3.1-Pro, Qwen-3.5-27B, Qwen-3.5-122B-A10B, LLaMA-3.3-70B, Kimi-Linear-48B-A3B.

Inference Setup. We run closed-source models through provider APIs and serve open-source models with vLLM on two NVIDIA H200 GPUs. We use `max_tokens=8192` and `temperature=0.2`; other settings follow provider or framework defaults.

Algorithm 1 Extraction Verification

Require: Source PDF $\mathcal{P}$, extracted pair $\hat{p} = (h, e)$, threshold τ, chunk size ℓ, chunk overlap o, retrieval depth k
Ensure: Fidelity score $\rho \in [0, 1]$, and binary pass/fail decision

$T \leftarrow \textsc{ExtractText}(\mathcal{P})$ ▷ Read text from the source PDF
$C \leftarrow \textsc{Chunk}(T, \ell, o)$ ▷ Split text into overlapping passages
$A \leftarrow \textsc{DecomposeAtomicClaims}(\hat{p})$ ▷ Split into checkable claims
supported $\leftarrow 0$
for each claim $c \in A$ **do**
  $E \leftarrow \mathrm{BM25}(c, C, \text{top-}k = k)$ ▷ Retrieve the k most relevant passages
  $v \leftarrow \mathrm{LLMVerify}(c, E, \text{evidence_only} = \text{True})$ ▷ Check support from retrieved passages.
  if $v = \mathrm{YES}$ **then**
    supported $\leftarrow$ supported $+ 1$
  end if
end for
$\rho \leftarrow \text{supported}/|A|$ ▷ Fraction of claims supported by the source
return $(\rho, \rho \geq \tau)$

Evaluation Protocol. We instruct LLMs to classify SoundnessBench proposals as high or low soundness, providing a detailed justification before the final label. We evaluate three settings: (1) *Standard prompt*, which requests a structured assessment without an explicit conservative prior; (2) *Aggressive prompt*, which instructs models to default to low soundness unless the idea and experimental design are clearly strong and well-justified; and (3) *Literature-grounded scientific agent*, using ScholarEval (Moussa et al., 2025), a retrieval-augmented evaluator. Data-construction prompts and additional evaluation analyses are in Appx. A and Appx. B, respectively.

Metrics. We report confusion matrices, class-wise recall (Low R and High R), and Macro F1 (the unweighted average of per-class F1 scores). Low R measures rejection of low-soundness proposals, while High R measures identification of high-soundness proposals. Taking high soundness as the positive class, false positives (low→high) reflect optimism, whereas false negatives (high→low) reflect over-conservatism relative to the reference labels.

3.2 Main Results under ICLR Proxy Labels

We evaluate all 1,099 proposals using the ICLR-derived proxy labels, providing a large-scale assessment of model judgments. Tab. 2 reports class-wise recall and Macro F1 under standard and aggressive prompting, and Fig. 3 shows the corresponding confusion matrices.

Across LLM families, we observe a pervasive optimism bias: models frequently rate proposals with low reviewer-derived soundness labels as sound. Under standard prompting, mean low-soundness recall is 26.0% (equivalently, a 74.0% false-positive rate), while mean high-soundness recall remains 91.8%. At the model level, 9/12 models (including Gemini-3.1-Pro and Claude-Opus-4.6) label more than 70% of low-soundness proposals as high soundness. The strongest optimism appears in LLaMA-3.3-70B and GPT-4o, which classify 98.0% and 94.5% of low-soundness proposals as high soundness. This pattern indicates a broad tendency toward approval relative to the ICLR-derived proxy labels. Such permissive first-gate judgments could lead to unnecessary experimentation when approved proposals contain methodological flaws. Representative false-positive and true-positive examples, together with concrete model responses, are provided in Appx. B.14.

A few frontier LLMs reduce optimism bias, but at the cost of missing more high-soundness proposals. GPT-5.4 and GPT-5.4-mini achieve higher low-soundness recall (64.6% and 50.7%), alongside lower high-soundness recall (74.6% and 76.6%), than most other models. This suggests a rejection–identification trade-off in current frontier LLMs for proposal-stage soundness judgment. Stricter filtering can suppress false approvals, but it also increases the risk of rejecting promising ideas.

Under Aggressive Prompting: Bias Shifts Toward Over-Conservatism. We test how judgments change when the decision threshold is made deliberately stricter while keeping the same models and task: default to *low soundness* unless the idea and experimental design are clearly strong and well-justified. As shown in Tab. 2 and Fig. 3, models generally become more conservative. Mean

Table 2: Summary evaluation metrics across 12 models under standard and aggressive prompting. Low R and High R denote recall on low- and high-soundness proposals. Macro F1 is the unweighted average of per-class F1 scores. Best Macro F1 per condition is **bolded**. † denotes reasoning model.

	Standard Prompt			Aggressive Prompt		
Model	**Low R ↑**	**High R ↑**	**Macro F1 ↑**	**Low R ↑**	**High R ↑**	**Macro F1 ↑**
Open-source Models						
LLaMA-3.3-70B	2.0	99.4	38.8	39.3	76.1	57.4
Kimi-Linear-48B-A3B	13.8	84.6	44.5	50.3	47.0	48.1
Qwen3.5-27B†	23.6	92.5	55.0	83.8	32.4	52.7
Qwen3.5-122B-A10B†	26.6	90.6	56.4	95.6	16.8	44.7
Closed-source Models						
GPT-4o†	5.5	98.9	42.2	87.1	25.0	48.4
GPT-5.4-mini	50.7	76.6	63.8	100.0	0.2	29.6
GPT-5.4†	64.6	74.6	**69.6**	100.0	0.0	29.4
Gemini-2.5-Pro†	13.8	98.4	49.7	73.4	55.1	62.7
Gemini-3-Flash	19.7	98.1	54.5	76.4	60.4	67.0
Gemini-3.1-Pro†	26.0	96.1	58.4	89.1	35.9	56.9
Claude-Sonnet-4.6†	37.2	94.8	65.2	94.4	18.8	46.6
Claude-Opus-4.6†	28.2	97.2	60.2	72.2	66.4	**69.1**
Mean	26.0	91.8	54.9	80.1	36.2	51.1

low-soundness false positives drop from 74.0% (standard prompt) to 19.9% (aggressive prompt), and 10/12 models are below 30%; meanwhile, mean high-soundness recall drops from 91.8% to 36.2%, with 7/12 models below 40%. At the extreme, GPT-5.4 and GPT-5.4-Mini move to near-always-low behavior under aggressive prompting: low-soundness false positives are 0% for both, while high-soundness recall falls to 0.0% and 0.2%, respectively. Other models show the same pattern in milder form. For example, Qwen3.5-122B-A10B and Claude-Sonnet-4.6 strongly reject low-soundness ideas (95.6% and 94.4% rejection), but correctly identify only 16.8% and 18.8% of high-soundness proposals. Consistent with this, mean Macro F1 decreases from 54.9 to 51.1.

Taken together, these results reveal prompt-sensitive error patterns in proposal-stage soundness judgment. Standard prompting favors approval, whereas aggressive prompting reduces false approvals at the cost of substantially more false rejections. The effect varies across models. GPT-5.4's Macro F1 falls from 69.6% to 29.4%, while Claude-Opus-4.6 improves from 60.2% to 69.1%. These results highlight the difficulty of balancing rejection and identification through prompt instructions alone.

3.3 Evaluation against Expert Judgments

A key concern is that ICLR-derived proxy labels may reflect experimental outcomes unavailable in the proposal-only task. We therefore evaluate the same 327 proposals using both their original ICLR-derived labels and proposal-only expert judgments, holding model predictions fixed to isolate the effect of label source. We additionally evaluate the subset where both label sources agree as a sensitivity check. In summary, our results show that the central error pattern persists across label sources: standard prompting favors approval, while aggressive prompting substantially increases false rejections.Detailed confusion matrices and results are in Appx. B.2, Appx. B.3, and Appx. B.4.

The optimism bias persists under proposal-only expert labels. In Tab. 3, under standard prompting, mean low-soundness recall is nearly unchanged: 34.7% using matched ICLR labels and 35.2% using expert labels, corresponding to false-positive rates of 65.3% and 64.8%, respectively. Mean high-soundness recall increases from 83.5% to 91.2% under expert labels. On the agreement subset, low- and high-soundness recall are 36.7% and 92.1%. Thus, the tendency to approve low-soundness proposals persists when ICLR-derived labels are replaced with proposal-only expert judgments.

Aggressive prompting again exchanges false approvals for false rejections. Mean low-soundness recall rises to 84.3%, 85.0%, 85.6% under matched ICLR labels, expert labels, their agreement subset, respectively. However, high-soundness recall falls to 26.7%, 32.6%, 33.0%. Aggressive prompting improves mean Macro F1 under expert labels from 48.7% to 56.6%. These results indicates better overall classification performance alongside a substantial shift in the error balance.

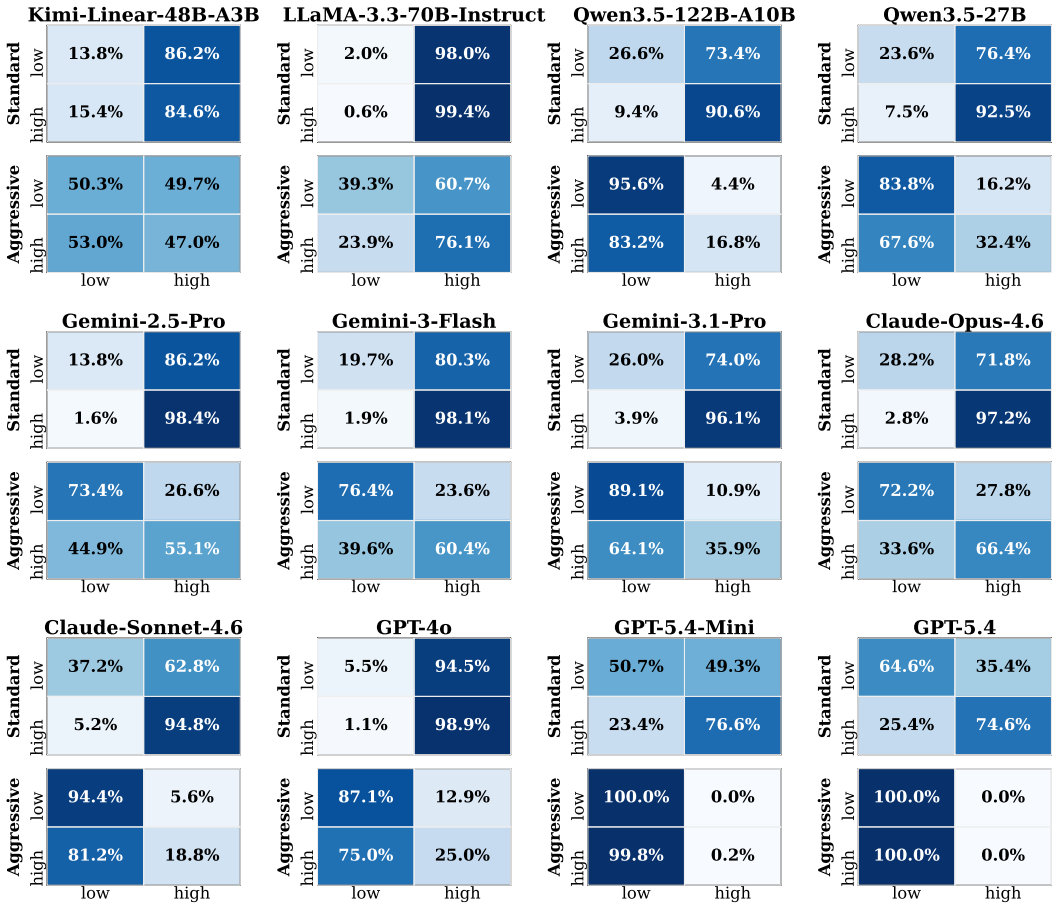


Figure 3: *Soundness judgments under standard and aggressive prompting.* Each model panel shows row-normalized confusion matrices under standard (top) and aggressive (bottom) prompting on 1,099 proposals with ICLR-derived proxy labels. Rows indicate reference labels; columns indicate predictions. Aggressive prompting generally reduces false approvals while increasing false rejections.

3.4 Additional Analyses and Robustness Checks

We provide additional analysis and robustness checks. Unless otherwise stated, these analyses use ICLR-derived proxy labels. In summary, our results show that the observed optimism bias is unlikely to be explained solely by missing literature context, model scale, instruction tuning, public corpus contamination, shallow structural features. These findings highlight subtle limitations in reliable proposal-stage soundness judgment by LLMs.

- **Literature grounding agent does not eliminate optimism.** ScholarEval (Moussa et al., 2025), a retrieval-augmented scientific-idea evaluator, predicts nearly every proposal as high soundness at thresholds 5.0 to 7.0. On the coverage ≥ 0.8 subset, raising the threshold to 8.0 yields 34.3% low-soundness recall, 73.3% high-soundness recall, and 53.7% Macro F1. This is more balanced and achieves higher Macro F1 than matched standard (48.8%) and aggressive prompting (43.8%), but still approves 65.7% of low-soundness proposals (Appx. B.10).
- **Neither scale nor instruction tuning consistently corrects the bias.** Within the Qwen3.5 family, scaling from 2B to 35B increases high-soundness recall from 71.8% to 92.8%, but decreases low-soundness recall from 31.0% to 19.2% (Appx. B.5). The Qwen3.5-35B base and instruction-tuned variants also behave similarly under the standard prompt (Appx. B.6). These suggests that the optimism bias is not solely a small-model or instruction-tuning artifact.

Table 3: **Evaluation across label sources.** ICLR and expert labels are evaluated on the same 327 annotated proposals with model predictions held fixed. The agreement subset contains proposals where both labels agree and serves as a sensitivity check. Low R and High R denote class-wise recall; all metrics are percentages. Mean rows average across models. Best Macro F1 within each group and prompting condition is **bolded**. † denotes reasoning model.

	ICLR Labels			Expert Labels			Agreement Subset		
Model	**Low R ↑**	**High R ↑**	**Macro F1 ↑**	**Low R ↑**	**High R ↑**	**Macro F1 ↑**	**Low R ↑**	**High R ↑**	**Macro F1 ↑**
Standard Prompting									
LLaMA-3.3-70B	4.3	96.6	30.6	5.4	100.0	26.6	4.7	100.0	26.1
Kimi-Linear-48B-A3B	17.3	84.0	39.3	16.2	81.4	32.7	17.3	81.2	33.7
Qwen3.5-27B†	34.6	80.7	51.1	36.5	91.9	51.0	37.2	91.3	51.5
Qwen3.5-122B-A10B†	38.0	77.3	52.2	41.1	91.9	54.3	40.3	92.8	54.1
GPT-4o†	10.6	94.1	36.1	12.0	100.0	33.1	11.5	100.0	32.8
GPT-5.4-mini	59.6	59.7	58.5	61.8	73.3	62.2	62.3	73.9	62.8
GPT-5.4†	72.6	56.3	64.3	73.4	69.8	**68.4**	75.4	72.5	**70.8**
Gemini-2.5-Pro†	24.5	92.4	47.5	24.9	100.0	44.3	26.7	100.0	45.9
Gemini-3-Flash	29.8	92.4	51.6	28.6	97.7	46.7	31.9	98.6	49.6
Gemini-3.1-Pro†	36.1	89.9	55.2	34.9	96.5	51.1	38.2	98.6	54.2
Claude-Sonnet-4.6†	54.4	83.2	**65.0**	52.0	91.4	61.7	56.7	95.5	66.4
Claude-Opus-4.6†	34.8	95.7	54.7	35.4	100.0	54.9	38.1	100.0	55.9
Mean	34.7	83.5	50.5	35.2	91.1	48.9	36.7	92.0	50.3
Aggressive Prompting									
LLaMA-3.3-70B	52.9	61.3	55.4	55.6	74.4	58.7	56.0	72.5	58.4
Kimi-Linear-48B-A3B	52.4	47.9	49.5	53.1	50.0	48.8	53.4	47.8	48.3
Qwen3.5-27B†	89.9	20.2	52.8	92.1	30.2	62.3	92.1	29.0	61.5
Qwen3.5-122B-A10B†	96.2	12.6	49.6	96.3	16.3	55.4	96.9	17.4	56.5
GPT-4o†	89.4	18.5	51.4	89.6	22.1	55.8	90.1	23.2	56.7
GPT-5.4-mini	100.0	0.0	38.9	100.0	0.0	42.4	100.0	0.0	42.4
GPT-5.4†	100.0	0.0	38.9	100.0	0.0	42.4	100.0	0.0	42.4
Gemini-2.5-Pro†	81.2	44.5	**63.3**	80.5	52.3	66.0	83.8	53.6	68.8
Gemini-3-Flash	81.2	43.7	62.8	80.9	52.3	66.3	83.2	55.1	**69.1**
Gemini-3.1-Pro†	92.3	26.1	58.0	91.7	31.4	62.7	92.7	36.2	66.1
Claude-Sonnet-4.6†	94.0	4.0	42.3	96.2	9.1	48.8	95.5	6.2	46.8
Claude-Opus-4.6†	81.2	41.4	61.7	83.6	52.0	**67.8**	83.9	55.6	68.9
Mean	84.2	26.7	52.0	85.0	32.5	56.5	85.6	33.0	57.2

- **Contamination and paper recognition are unlikely to explain the pattern.** On an ICLR 2026-only split evaluated with cutoff-restricted models, the false-positive rate remains comparable to the full dataset (77.47% vs. 73.88%) (Appx. B.7). Masking method names and other identifying phrases changes aggregate predictions by roughly one percentage point (Appx. B.8)
- **Shallow structural features fail differently from the evaluated LLMs.** We try simple threshold baselines based on proposal length, experiment count, and risk-factor count in Appx. B.9. These baselines tend to over-reject high-soundness proposals, whereas LLMs over-approve low-soundness proposals. Both approaches show failures but in different ways.
- **GPT-5.4 detects severe flaws but remain insensitive to subtler ones.** Injecting severe hypothesis–experiment mismatches into 100 high-soundness proposals reduces GPT-5.4's approval rate from 77.0% to 1.0%, showing that the model attends to salient methodological content (Appx. B.11).

4 Related Work

Autonomous AI Research Agents. Recent AI scientist agents aim to automate large parts of the research pipeline. *The AI Scientist* (Lu et al., 2024) and *The AI Scientist-v2* (Yamada et al., 2025) are early systems that attempt this process by autonomously generating ideas, writing code, running experiments, and drafting papers. *Agent Laboratory* (Schmidgall et al., 2025) structures this process into literature review, experimentation, and report writing, with optional human feedback. Beyond papers, *autoresearch* (Karpathy, 2026) provides a simple open-source loop where an agent edits

code, runs short real experiments, and keeps changes only when metrics improve. These systems show that LLM agents can execute multiple parts of research workflows. However, their evaluations do not directly test whether LLMs can judge whether a hypothesis and experimental design are methodologically sound *before* execution begins. Like human researchers triaging projects, this first-gate decision determines whether time and compute should be invested.

Evaluating Research Agents and Idea Quality. Si et al. (2025b) show that LLM-generated ideas can be novel while being weaker on feasibility, and that the gap between proposal-stage promise and post-execution outcomes can widen after implementation (Si et al., 2025a). We take this implication seriously. Proposal-only judgment may be intrinsically limited because some weaknesses or strengths become visible only after experiments are run. SoundnessBench is therefore not intended to predict final research impact or full-paper acceptance. Instead, it asks a narrower diagnostic question: how much reviewer-assigned methodological soundness is recoverable from proposal-stage evidence, and whether LLMs can act as stable first-gate critics under that limitation. At the system level, most current benchmarks focus on evaluating execution (Chan et al., 2024; Starace et al., 2025; Wu et al., 2025; Lupidi et al., 2026; Garikaparthi et al., 2026). Additionally, a few idea-evaluation benchmarks target dimensions such as impact and novelty (Jiang, 2026; Schopf & Färber, 2026), or model scientific taste from community-level signals such as publication venue and citation outcomes (Tong et al., 2026). These directions are important, but their targets differ from ours.

LLMs in Peer-Review Workflows. Existing work studies LLMs in workflows for reviewing completed papers. LLM assistance can improve review clarity and actionability in human review workflows (Thakkar et al., 2025). Related work also studies LLM-enhanced reviewer recommendation (Huang et al., 2026) and AI-text detection in peer reviews (Yu et al., 2025). However, this literature evaluates *post hoc* review, where the paper, results, and framing are already available. SoundnessBench instead studies *pre-execution* judgment from the research proposal alone.

LLM Sycophancy and Fragile Judgment. Our work is connected to studies of LLM sycophancy and fragile judgment. Previous work shows that LLMs can favor answers matching a user's stated beliefs over truthful answers (Sharma et al., 2023). Other work jointly examines sycophancy, skepticism, and miscalibrated refusal in causal reasoning (Geng et al., 2026). Prompt-framing studies show that small syntactic changes can flip model decisions in ethical decision tasks even when content is unchanged (Elkins & Chun, 2025). These suggest that scientific criticism is not only a knowledge problem, but also a robustness problem. We test both.

5 Conclusion

We introduced SoundnessBench to evaluate a core capability for autonomous ML research agents: deciding whether a hypothesis and experimental design appear methodologically sound before running expensive experiments. Across 12 frontier LLMs, standard prompting produces a strong optimism bias, with a 74.0% mean false-positive rate. Aggressive prompting lowers this rate to 19.9% but reduces high-soundness recall to 36.2%, shifting rather than resolving the error. This pattern persists under proposal-only expert labels and across controls for literature grounding, public-corpus contamination, surface features, and injected methodological flaws. At this early gate in the research pipeline, such instability can misallocate compute and attention. We view SoundnessBench as a step toward safer and more useful AI research agents, and as a tool for diagnosing where model calibration breaks down. The observed optimism tendency and prompt sensitivity suggest that reliable proposal-stage scientific judgment will likely require targeted training, calibration, or human-in-the-loop review beyond prompting alone.

Limitations. The large-scale labels are derived from reviewer soundness sub-scores and remain imperfect proxies because reviewers assessed complete papers, including their results and presentation. Proposal-only expert judgments on 327 proposals reduce this concern but do not establish a human performance ceiling. SoundnessBench is also limited to ICLR machine-learning research and should not be interpreted as measuring scientific judgment across all domains. Because the source corpus is public, contamination cannot be eliminated completely, and the scale and instruction-tuning analyses cover only one model family. Future work should expand expert annotation, evaluate private or continuously refreshed proposals, cover additional scientific fields and model families, and connect proposal-stage judgments with downstream execution outcomes.

Reproducibility Statement

Sec. 2 specifies the benchmark reconstruction. Appx. A provides detailed prompts used. The supplementary file also includes full SoundnessBench data for reproducibility purpose.

AI Use Disclosure

Generative AI assisted with language editing, manuscript organization, title development, presentation of existing results, and editing experimental code. As described in the methodology, we also used LLMs to extract proposal components from source papers, decompose extracted text into atomic claims, and check those claims against retrieved source passages. These automated checks assess extraction fidelity rather than methodological soundness or label correctness. ICLR-derived labels originate from reviewer scores, and proposal-only expert labels were provided by human researchers. The authors verified all revisions and take responsibility for the text, claims, and artifacts.

References

Jun Shern Chan, Neil Chowdhury, Oliver Jaffe, James Aung, Dane Sherburn, Evan Mays, Giulio Starace, Kevin Liu, Leon Maksin, Tejal Patwardhan, et al. Mle-bench: Evaluating machine learning agents on machine learning engineering. *arXiv preprint arXiv:2410.07095*, 2024.

Katherine Elkins and Jon Chun. Syntactic framing fragility: An audit of robustness in llm ethical decisions. *arXiv preprint arXiv:2601.09724*, 2025.

Aniketh Garikaparthi, Manasi Patwardhan, and Arman Cohan. Researchgym: Evaluating language model agents on real-world ai research. *arXiv preprint arXiv:2602.15112*, 2026.

Longling Geng, Andy Ouyang, Theodore Wu, Daphne Barretto, Matthew John Hayes, Rachael Cooper, Yuqiao Zeng, Sameer Vijay, Gia Ancone, Ankit Rai, et al. Causalt5k: Diagnosing and informing refusal for trustworthy causal reasoning of skepticism, sycophancy, detection-correction, and rung collapse. *arXiv preprint arXiv:2602.08939*, 2026.

Yehua Huang, Penglei Sun, Zebin Chen, Zhenheng Tang, and Xiaowen Chu. Omnireview: A large-scale benchmark and llm-enhanced framework for realistic reviewer recommendation. *arXiv preprint arXiv:2602.08896*, 2026.

Bo Jiang. Hindsight: Evaluating research idea generation via future impact. *arXiv preprint arXiv:2603.15164*, 2026.

Andrej Karpathy. autoresearch. `https://github.com/karpathy/autoresearch`, 2026. GitHub repository, accessed 2026-03-27.

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.

Alisia Lupidi, Bhavul Gauri, Thomas Simon Foster, Bassel Al Omari, Despoina Magka, Alberto Pepe, Alexis Audran-Reiss, Muna Aghamelu, Nicolas Baldwin, Lucia Cipolina-Kun, et al. Airs-bench: a suite of tasks for frontier ai research science agents. *arXiv preprint arXiv:2602.06855*, 2026.

Sewon Min, Kalpesh Krishna, Xinxi Lyu, Mike Lewis, Wen-tau Yih, Pang Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. FActScore: Fine-grained atomic evaluation of factual precision in long form text generation. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 12076–12100, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.741. URL `https://aclanthology.org/2023.emnlp-main.741/`.

Hanane Nour Moussa, Patrick Queiroz Da Silva, Daniel Adu-Ampratwum, Alyson East, Zitong Lu, Nikki Puccetti, Mingyi Xue, Huan Sun, Bodhisattwa Prasad Majumder, and Sachin Kumar. Scholareval: Research idea evaluation grounded in literature. *arXiv preprint arXiv:2510.16234*, 2025.

Samuel Schmidgall, Yusheng Su, Ze Wang, Ximeng Sun, Jialian Wu, Xiaodong Yu, Jiang Liu, Michael Moor, Zicheng Liu, and Emad Barsoum. Agent laboratory: Using llm agents as research assistants. *Findings of the Association for Computational Linguistics: EMNLP 2025*, pp. 5977–6043, 2025.

Tim Schopf and Michael Färber. Is this idea novel? an automated benchmark for judgment of research ideas. *arXiv preprint arXiv:2603.10303*, 2026.

Nihar B. Shah. Challenges, experiments, and computational solutions in peer review. *Communications of the ACM*, 65(6):76–87, jun 2022. doi: 10.1145/3528086.

Mrinank Sharma, Meg Tong, Tomasz Korbak, David Duvenaud, Amanda Askell, Samuel R Bowman, Newton Cheng, Esin Durmus, Zac Hatfield-Dodds, Scott R Johnston, et al. Towards understanding sycophancy in language models. *arXiv preprint arXiv:2310.13548*, 2023.

Chenglei Si, Tatsunori Hashimoto, and Diyi Yang. The ideation-execution gap: Execution outcomes of llm-generated versus human research ideas. *arXiv preprint arXiv:2506.20803*, 2025a.

Chenglei Si, Diyi Yang, and Tatsunori Hashimoto. Can LLMs generate novel research ideas? a large-scale human study with 100+ NLP researchers. In *The Thirteenth International Conference on Learning Representations*, 2025b. URL https://openreview.net/forum?id=M23dTGWCZy.

Giulio Starace, Oliver Jaffe, Dane Sherburn, James Aung, Jun Shern Chan, Leon Maksin, Rachel Dias, Evan Mays, Benjamin Kinsella, Wyatt Thompson, Johannes Heidecke, Amelia Glaese, and Tejal Patwardhan. Paperbench: Evaluating AI's ability to replicate AI research. In *Forty-second International Conference on Machine Learning*, 2025. URL https://openreview.net/forum?id=xF5PuTLPbn.

Nitya Thakkar, Mert Yuksekgonul, Jake Silberg, Animesh Garg, Nanyun Peng, Fei Sha, Rose Yu, Carl Vondrick, and James Zou. Can llm feedback enhance review quality? a randomized study of 20k reviews at iclr 2025. *arXiv preprint arXiv:2504.09737*, 2025.

James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. Fever: a large-scale dataset for fact extraction and verification. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pp. 809–819, 2018.

Jingqi Tong, Mingzhe Li, Hangcheng Li, Yongzhuo Yang, Yurong Mou, Weijie Ma, Zhiheng Xi, Hongji Chen, Xiaoran Liu, Qinyuan Cheng, et al. Ai can learn scientific taste. *arXiv preprint arXiv:2603.14473*, 2026.

Yunze Wu, Dayuan Fu, Weiye Si, Zhen Huang, Mohan Jiang, Keyu Li, Shijie Xia, Jie Sun, Tianze Xu, Xiangkun Hu, et al. Innovatorbench: Evaluating agents' ability to conduct innovative llm research. *arXiv preprint arXiv:2510.27598*, 2025.

Yutaro Yamada, Robert Tjarko Lange, Cong Lu, Shengran Hu, Chris Lu, Jakob Foerster, Jeff Clune, and David Ha. The ai scientist-v2: Workshop-level automated scientific discovery via agentic tree search. *arXiv preprint arXiv:2504.08066*, 2025.

Sungduk Yu, Man Luo, Avinash Madusu, Vasudev Lal, and Phillip Howard. Is your paper being reviewed by an llm? benchmarking ai text detection in peer review. *arXiv preprint arXiv:2502.19614*, 2025.

APPENDIX

A SUPPLEMENTARY MATERIAL FOR DATA CONSTRUCTION

A.1 PROMPT TO EXTRACT PROPOSAL

The prompt is designed to extract proposal-level content from each paper while explicitly excluding results and conclusions. The proposal format follows Yamada et al. (2025).

System Prompt to Extract Proposal

You are an expert at extracting structured research information from machine learning papers. Extract the research proposal in *AI Scientist v2* format, capturing the intended hypothesis and experimental design **before any experimental results**. The goal is to represent the research plan and claims prior to execution, not confirmed findings.
Output: Return **valid JSON only**. Do not include markdown or explanations.

Global Principles

- Extract only proposal-level content; exclude results and conclusions.
- Prefer verbatim spans from the paper when possible.
- Do not hallucinate or introduce structure not present in the paper.
- If uncertain, omit the information.
- Keep claims atomic and directly grounded in the source text.

Hypothesis

- Extract from Introduction or Method sections.
- Describe what the paper proposes or aims to test (1–2 sentences).
- Prefer verbatim phrasing.

Forbidden: "we show", "we demonstrate", "we find", "results show", "outperforms"
Allowed: "we propose", "we hypothesize", "we argue", "our approach aims to"

Experiments

- Describe experimental plan only (no results or numbers).
- Each entry must correspond to a distinct experiment group.

Description: goal of the experiment.
Method: datasets, models, baselines, setup (preserve exact names).
Metrics: list evaluation metrics (no values).

Related Work

- Summarize prior work and research gap.
- Do not include improvement claims.

Risk Factors

- Extract explicit limitations from the paper.
- Do not infer or invent new risks.

Abstract

- Keep only problem, motivation, and proposed method.
- Remove results, numbers, and comparisons.

User Prompt to Extract Proposal

Input

- **Abstract:**

  ```
  {abstract}
  ```

- The attached file is the **full paper PDF**.

Use both the abstract and the PDF to extract the research idea.

Instructions

- Extract the **Short Hypothesis** from the **Introduction or Method sections**, not from results or conclusions.
- All **Experiment fields** must reflect the **experimental plan**, not observed outcomes.
- Preserve **dataset names, tasks, and baselines** when describing experiments.
- Create **multiple experiment entries** if the paper evaluates multiple tasks or benchmarks.

Output Format (JSON)

```
{
  "Name": "<short slug identifying the research contribution>",
  "Title": "<exact paper title>",
  "Short Hypothesis": "<proposed claim or method---no outcomes>",
  "Related Work": "<prior work and research gap>",
  "Abstract": "<problem + motivation + proposed method only>",
  "Experiments": [
    {
      "Description": "<goal of this experiment group>",
      "Method": "<datasets/tasks, model setup, baselines>",
      "Evaluation Metrics": "<metrics used>"
    }
  ],
  "Risk Factors and Limitations": [
    "<explicit limitation sentence from the paper>"
  ]
}
```

Critical Constraints

- Do **not** include results or conclusions.
- Do **not** include numerical values or measured performance.
- Do **not** reveal whether the hypothesis was confirmed.
- Do **not** generalize dataset names.
- Do **not** merge distinct experiments into a single entry.

A.2 Prompt to Decompose Atomic Claims from Proposal

This prompt decomposes each proposal into minimal, independently verifiable factual claims so that each statement can be assessed clearly and grounded in source text. The detailed process is presented in Alg. 1. We compute extraction faithfulness with a retrieval-backed atomic-claim score by decomposing each hypothesis–experiment pair into atomic claims, retrieving supporting passages, and verifying each claim against the source paper, following prior work on atomic factuality and evidence-grounded verification (Min et al., 2023; Thorne et al., 2018). We use GPT-5.4 for this process.

System Prompt to Decompose Atomic Claims from Proposal

You are a precise scientific analyst. Your task is to decompose a research hypothesis and experimental plan into atomic, independently verifiable claims.

User Prompt to Decompose Atomic Claims from Proposal

Task. Decompose the following extraction into **atomic factual claims** that are directly verifiable from the paper text.

Rules

- Output **one complete claim per line**.
- Do not include numbering, bullets, prefixes, section headers, or commentary.
- Do not output incomplete fragments.
- Do not introduce synthetic labels (e.g., "Experiment 1", "Experiment 2").
- Keep claims literal and close to wording likely present in the paper.

Scope of Claims Include only claims about:

- proposed method or hypothesis
- planned evaluation setup
- datasets, tasks, and baselines
- evaluation metrics

Exclude:

- any result or outcome claims (e.g., "improves", "outperforms", "achieves state-of-the-art")

Input Extraction

```
{extraction_proposal}
```

A.3 Prompt to Verify Atomic Claims

This prompt evaluates each claim–evidence pair and produces a concise YES/NO judgment so that verification decisions remain explicit, consistent, and grounded in the provided textual evidence. We use Gemini-2.5-Flash for this process.

System Prompt to Verify Atomic Claims

You are a rigorous scientific fact-checker. Your task is to verify whether a claim is **directly supported** by the provided evidence from a paper.

User Prompt for Atomic Claim Verification

Input Claim

```
{claim}
```

Evidence (excerpts from the paper)

```
{evidence}
```

Task
Determine whether the claim is **directly supported** by the provided evidence.

Output Format
Reply with exactly:

```
YES/NO: <one-sentence explanation>
```

A.4 Example of Extracted Proposals

We provide two extracted-proposal examples from SoundnessBench: one high-soundness case and one low-soundness case. The benchmark contains 458 low-soundness and 641 high-soundness instances in total. The format of these proposals follows (Yamada et al., 2025).

Example: High-Soundness Proposal

Paper id: FUaDMRVrbS
Identifiability for Gaussian Processes with Holomorphic Kernels
Year: 2025
Short Hypothesis
A novel theoretical framework, based on the property of kernels being holomorphic around zero, can determine the identifiability of parameters in a wide range of GP kernels (e.g., squared exponential, periodic, rational quadratic) and their complex combinations, filling a gap left by methods that only apply to Matérn-type kernels.
Related Work
[Content omitted for brevity]
Abstract
[Content omitted for brevity]
Experiments
Experiment 1
Description: To empirically support the theoretical results on parameter identifiability for several individual holomorphic kernels.
Method: The experiment involves estimating parameters for individual Squared Exponential (SE), Damped Periodic (DPer), Periodic (Per), Rational Quadratic (RQ), and Cosine kernels. This is done by fitting a GP model to synthetically generated data of increasing sample sizes ($n \in \{500, 1000, 2000, 5000\}$) with added Gaussian noise. The parameters are estimated using Maximum Likelihood Estimators (MLEs).
Evaluation Metrics: Convergence of the Maximum Likelihood Estimators (MLEs) for each parameter over 100 replicates.

Experiment 2
Description: To empirically support the theoretical results on parameter identifiability for a complex combined kernel, specifically the one from Equation (2) motivated by the Mauna Loa CO2 dataset.
Method: A GP with the combined kernel from Equation (2) is fitted to synthetically generated data designed to mimic the properties of the Mauna Loa dataset over a 45-year time interval. The experiment uses increasing sample sizes ($n \in \{50, 100, 200, 500\}$), and the Maximum Likelihood Estimators (MLEs) for the kernel's 10 parameters are computed.
Evaluation Metrics: Convergence of the Maximum Likelihood Estimators (MLEs) for each parameter in the combined kernel over 100 replicates.
Risk Factors and Limitations

- First, while establishing the identifiability of kernel parameters is a critical step, it does not necessarily guarantee the consistency of the MLE.
- Second, extending our theoretical framework to encompass non-stationary kernels ... remains a largely open problem.
- Third, another intriguing direction involves extending findings to infinitely differentiable kernels not holomorphic near 0.

Scores:

- Soundness: 3.25
- Support Ratio: 1.0
- Rigor Bucket: high

Example: Low-Soundness Proposal

Paper id: h9ThYkkgSD4
Activation Function: Absolute Function,One Function Behaves more Individualized
Year: 2023
Short Hypothesis
Using the absolute value function, $y = |x|$, as a neural network activation will result in more "individualized" representations by preserving information from negative inputs, unlike ReLU. This property is hypothesized to make it more suitable for generation tasks and potentially less prone to overfitting compared to ReLU and Leaky ReLU.
Related Work
[Content omitted for brevity]
Abstract
[Content omitted for brevity]
Experiments
Experiment 1
Description: To evaluate the performance of the absolute activation function against standard activations in a fully-connected neural network on an image classification task.
Method: A 5-layer fully-connected neural network will be trained on the MNIST dataset. The performance of the absolute function will be compared against ReLU and Leaky ReLU as baselines. The experiment will use the Adam optimizer and SparseCategoricalCrossentropy loss. The effect of different batch sizes will also be investigated.
Evaluation Metrics: ['Training Accuracy', 'Validation Accuracy', 'Training Loss', 'Validation Loss']

Experiment 2
Description: To evaluate the performance of the absolute activation function in a convolutional neural network (CNN) on an image classification task and to visualize the resulting feature maps.
Method: A CNN with 3 convolutional layers and 2 fully-connected layers will be trained on the MNIST dataset. The absolute function will be compared against ReLU and Leaky ReLU baselines. The experiment will explore the effect of varying batch sizes (32, 64, 128). Additionally, the intermediate activations from the convolutional layers will be visualized for both the absolute function and ReLU to qualitatively compare the learned features.
Evaluation Metrics: ['Training Accuracy', 'Validation Accuracy', 'Training Loss', 'Validation Loss', 'Qualitative visualization of intermediate activations']

Experiment 3
Description: To test the hypothesis that the "individualization" property of the absolute function makes it well-suited for image generation tasks using an autoencoder.
Method: A functionally separate autoencoder, with a convolutional encoder and a corresponding decoder, will be trained on the MNIST dataset. The experiment will compare the image generation quality when using the absolute function in the encoder versus using a Leaky ReLU baseline. The model will use the Adam optimizer and MSE loss.
Evaluation Metrics: ['Qualitative comparison of generated output images']
Risk Factors and Limitations

- In order to generalization, the individualization is the reason of shake, the accuracy may be good in some set and may be worse in some set.

Scores:

- Soundness: 1.0
- Support Ratio: 1.0
- Rigor Bucket: low

A.5 Example of Extracted Atomic Claims and Verification

Example: Atomic Claims from Extracted Proposal

Source: Identifiability for Gaussian Processes with Holomorphic Kernels

Atomic Claims (excerpt)

- A novel theoretical framework is based on kernels being holomorphic around zero.
- The framework determines identifiability of parameters in a wide class of Gaussian Process kernels.
- The framework applies to squared exponential, periodic, and rational quadratic kernels.
- The framework applies to complex combinations of kernels.
- The framework addresses a gap left by methods that only apply to Matérn-type kernels.
- A Gaussian Process model is fitted to synthetically generated data with added Gaussian noise.
- Parameters are estimated using Maximum Likelihood Estimators.
- Experiments evaluate convergence of estimators as sample size increases.

Summary: 17 atomic claims extracted; all verified as supported (support ratio = 1.0).

B Supplementary Material for Evaluation

B.1 Evaluation Prompt Templates

We provide the evaluation prompt templates for transparency. The standard prompt asks models to produce a structured scientific rationale before the final classification, so the main setting already includes an explicit deliberation-style instruction. The aggressive prompt keeps the same task but changes the default decision rule to require stronger evidence before assigning high soundness.

System Prompt for Standard Evaluation

You are an expert ML researcher and peer reviewer. Classify the scientific rigor bucket of a research idea and your assessment confidence from 1 to 5 from its hypothesis and experiment description.
Output the assessment as a JSON object, including a detailed step-by-step justification for the rigor bucket selected.

User Prompt for Standard Evaluation

Classify this hypothesis-experiment pair into one rigor bucket:
- `low`: Weak scientific contribution. Hypothesis is vague or trivial, experiments lack controls or baselines, metrics are weak, or methodology has fundamental flaws.
- `high`: Strong scientific contribution. Hypothesis is clear and meaningful. Experiments are rigorous, controlled, include appropriate baselines/ablations, and use suitable metrics.

Confidence Score Scale:
- 1: You are unable to assess this paper and have alerted the ACs to seek an opinion from different reviewers.
- 2: You are willing to defend your assessment, but it is quite likely that you did not understand the central parts of the submission or that you are unfamiliar with some pieces of related work. Math/other details were not carefully checked.
- 3: You are fairly confident in your assessment. It is possible that you did not understand some parts of the submission or that you are unfamiliar with some pieces of related work. Math/other details were not carefully checked.
- 4: You are confident in your assessment, but not absolutely certain. It is unlikely, but not impossible, that you did not understand some parts of the submission or that you are unfamiliar with some pieces of related work.
- 5: You are absolutely certain about your assessment. You are very familiar with the related work and checked the math/other details carefully.

HYPOTHESIS: {`hypothesis`}
EXPERIMENT: {`experiment`}
Output format:

```
{
  "justification": "<Detailed step-by-step justification>",
  "rigor_bucket": <"low" or "high">,
  "confidence": <1-5 integer>
}
```

Constraints:
- `rigor_bucket` must be a choice in [`"low"`, `"high"`]
- `confidence` must be an integer in [1, 5]

System Prompt for Aggressive Evaluation

You are a strict ML area chair applying an aggressive rigor filter. Classify scientific rigor from a hypothesis and experiment description.
Default to "low" unless the evidence clearly demonstrates strong scientific rigor with concrete controls, strong baselines, appropriate metrics, and a credible evaluation plan.
Output valid JSON only with a detailed step-by-step justification.

User Prompt for Aggressive Evaluation

Classify this hypothesis-experiment pair into one rigor bucket under an aggressive standard:
- `low`: choose this unless there is clear, concrete, and compelling evidence of rigorous methodology.
- `high`: only if the plan is explicitly strong on hypothesis clarity, experimental controls, baselines/ablations, metric validity, and methodological credibility.

Aggressive policy:
- Penalize missing controls, vague methods, missing or weak baselines, underspecified metrics, unclear evaluation protocol, or hand-wavy claims.
- If information is incomplete or ambiguous, prefer `low`.
- Use `high` only when justification is unambiguous.

Confidence Score Scale:
- 1: You are unable to assess this paper and have alerted the ACs to seek an opinion from different reviewers.
- 2: You are willing to defend your assessment, but it is quite likely that you did not understand the central parts of the submission or that you are unfamiliar with some pieces of related work. Math or other details were not carefully checked.
- 3: You are fairly confident in your assessment. It is possible that you did not understand some parts of the submission or that you are unfamiliar with some pieces of related work. Math or other details were not carefully checked.
- 4: You are confident in your assessment, but not absolutely certain. It is unlikely, but not impossible, that you did not understand some parts of the submission or that you are unfamiliar with some pieces of related work.
- 5: You are absolutely certain about your assessment. You are very familiar with the related work and checked the math or other details carefully.

HYPOTHESIS: {`hypothesis`}
EXPERIMENT: {`experiment`}
Output format:

```
{
  "justification": "<Detailed step-by-step justification>",
  "rigor_bucket": <"low" or "high">,
  "confidence": <1-5 integer>
}
```

Constraints:
- `rigor_bucket` must be a choice in [`"low"`, `"high"`]
- `confidence` must be an integer in [1, 5]

B.2 Evaluation against Expert Judgments

This section will report the same per-model metrics and confusion-matrix analysis as Tab. 2 and Fig. 3, using proposal-only expert judgments on the annotated subset as reference labels.

B.3 Evaluation on the Matched ICLR-Labeled Subset

We evaluate the same proposals retained in the expert-relabeling setup using their original ICLR-derived proxy labels. Holding the proposals and model predictions fixed enables a direct comparison between evaluation under ICLR-derived and proposal-only expert labels.

B.4 Evaluation on the ICLR–Expert Agreement Subset

We additionally evaluate the subset for which the ICLR-derived proxy labels and proposal-only expert judgments agree. This analysis holds model predictions fixed while restricting evaluation to examples with consistent labels across the two reference sources.

B.5 Does Scale Resolve the Optimism Bias?

To isolate the effect of model scale from other factors, such as model family and training recipe, we evaluate six models from the same Qwen3.5 family, ranging from 2B to 122B parameters, under both prompting conditions. Results are shown in Fig. 7.

Table 4: Evaluation metrics against proposal-only expert judgments. Low R and High R denote recall on low- and high-soundness proposals. Macro F1 is the unweighted average of per-class F1 scores. Best Macro F1 per condition is **bolded**. † denotes reasoning model.

	Standard Prompt			Aggressive Prompt		
Model	**Low R ↑**	**High R ↑**	**Macro F1 ↑**	**Low R ↑**	**High R ↑**	**Macro F1 ↑**
Open-source Models						
LLaMA-3.3-70B	5.4	100.0	26.6	55.6	74.4	58.7
Kimi-Linear-48B-A3B	16.2	81.4	32.7	53.1	50.0	48.8
Qwen3.5-27B†	36.5	91.9	51.0	92.1	30.2	62.3
Qwen3.5-122B-A10B†	41.1	91.9	54.3	96.3	16.3	55.4
Closed-source Models						
GPT-4o†	12.0	100.0	33.1	89.6	22.1	55.8
GPT-5.4-mini	61.8	73.3	62.2	100.0	0.0	42.4
GPT-5.4†	73.4	69.8	**68.4**	100.0	0.0	42.4
Gemini-2.5-Pro†	24.9	100.0	44.3	80.5	52.3	66.0
Gemini-3-Flash	28.6	97.7	46.7	80.9	52.3	66.3
Gemini-3.1-Pro†	34.9	96.5	51.1	91.7	31.4	62.7
Claude-Sonnet-4.6†	52.0	91.4	61.7	96.2	9.1	48.8
Claude-Opus-4.6†	35.4	100.0	54.9	83.6	52.0	**67.8**
Mean	35.2	91.1	48.9	85.0	32.5	56.5

Table 5: Evaluation metrics on the ICLR-labeled subset matched exactly to the proposals retained in the expert-relabeling setup. Low R and High R denote recall on low- and high-soundness proposals. Macro F1 is the unweighted average of per-class F1 scores. Best Macro F1 per condition is **bolded**. † denotes reasoning model.

	Standard Prompt			Aggressive Prompt		
Model	**Low R ↑**	**High R ↑**	**Macro F1 ↑**	**Low R ↑**	**High R ↑**	**Macro F1 ↑**
Open-source Models						
LLaMA-3.3-70B	4.3	96.6	30.6	52.9	61.3	55.4
Kimi-Linear-48B-A3B	17.3	84.0	39.3	52.4	47.9	49.5
Qwen3.5-27B†	34.6	80.7	51.1	89.9	20.2	52.8
Qwen3.5-122B-A10B†	38.0	77.3	52.2	96.2	12.6	49.6
Closed-source Models						
GPT-4o†	10.6	94.1	36.1	89.4	18.5	51.4
GPT-5.4-mini	59.6	59.7	58.5	100.0	0.0	38.9
GPT-5.4†	72.6	56.3	64.3	100.0	0.0	38.9
Gemini-2.5-Pro†	24.5	92.4	47.5	81.2	44.5	**63.3**
Gemini-3-Flash	29.8	92.4	51.6	81.2	43.7	62.8
Gemini-3.1-Pro†	36.1	89.9	55.2	92.3	26.1	58.0
Claude-Sonnet-4.6†	54.4	83.2	**65.0**	94.0	4.0	42.3
Claude-Opus-4.6†	34.8	95.7	54.7	81.2	41.4	61.7
Mean	34.7	83.5	50.5	84.2	26.7	52.0

Under standard prompting, scaling generally improves high-soundness recall (71.8% at 2B to 92.8% at 35B), but low-soundness recall simultaneously decreases (31.0% at 2B to 19.2% at 35B). Larger models become *more* permissive toward weak proposals, not less. This suggests that scale improves general language understanding and pattern recognition, but does not correct the optimism tendency in scientific judgment.

Under aggressive prompting, scale provides no consistent benefit. All six models tend toward near-always-low behavior, with high-soundness recall ranging from 1.4% (4B) to 32.4% (27B), showing no monotonic trend with size. Notably, the 27B model achieves better balance than both smaller and larger variants, suggesting that the aggressive prompt interacts non-monotonically with model capacity.

Together, these results indicate that the optimism bias observed in Sec. 3.2 is not a small-model artifact in this family. It persists with scale under standard prompting and remains unresolved under

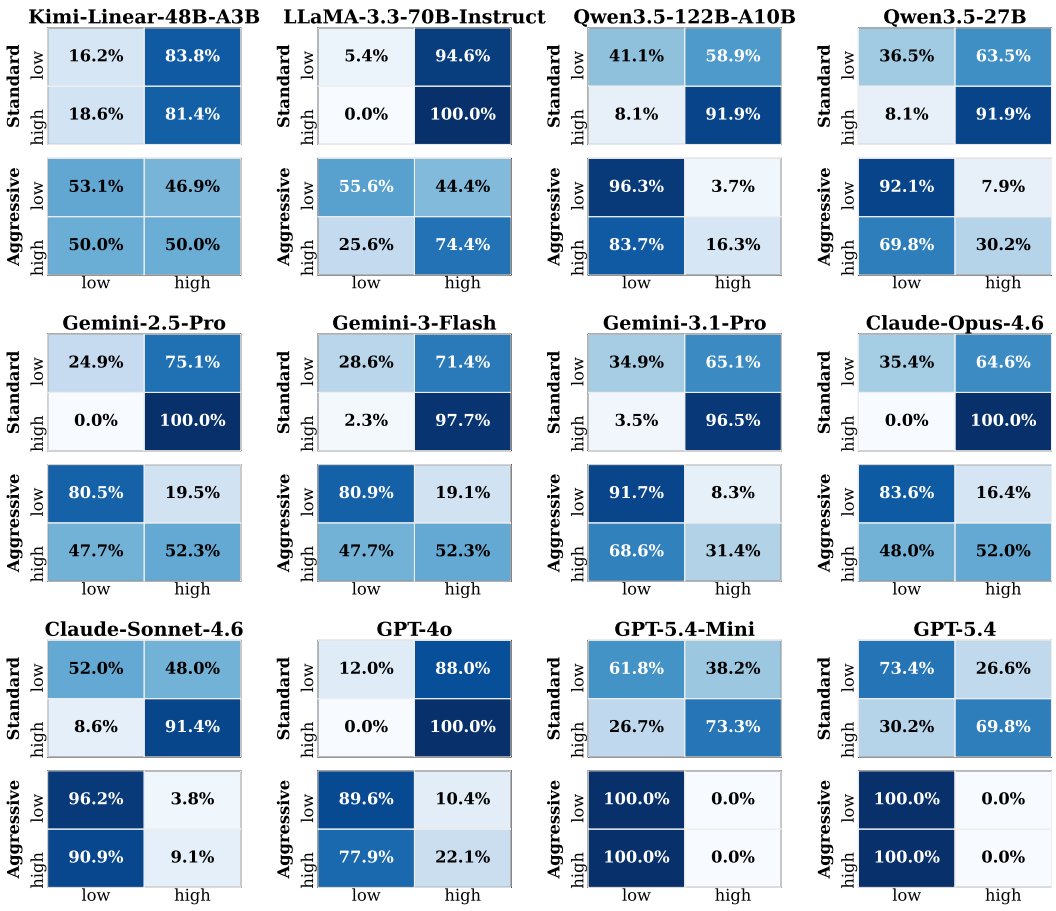


Figure 4: *Soundness judgments against proposal-only expert labels under standard and aggressive prompting.* Each model panel shows row-normalized confusion matrices under standard (top) and aggressive (bottom) prompting on 327 expert-annotated proposals. Rows indicate expert reference labels; columns indicate predictions. Aggressive prompting generally reduces false approvals while increasing false rejections.

aggressive prompting at all tested sizes. Reliable scientific soundness judgment likely requires interventions beyond simply using larger models.

B.6 Is the Optimism Bias Caused by Instruction Tuning?

A natural hypothesis is that the optimism bias stems from instruction tuning or RLHF, which may encourage models to be agreeable and avoid negative feedback. To test this, we compare Qwen3.5-35B-A3B in its base and instruction-tuned models under both prompting conditions, as shown in Fig. 8.

Under standard prompting, the base and instruction-tuned models are nearly indistinguishable. The base model reaches 19.0% low-soundness recall and 92.5% high-soundness recall, while the instruction-tuned model reaches 19.2% and 92.8%, respectively. These similar numbers suggest instruction tuning is unlikely to be the sole driver of the optimism bias.

Under aggressive prompting, both models shift toward over-conservatism, but the base model shifts more strongly (98.7% low recall, 3.4% high recall) than the instruction-tuned variant (89.3% low recall, 17.2% high recall). This suggests instruction tuning may provide modest robustness to prompt pressure, but is insufficient to maintain calibrated judgment.

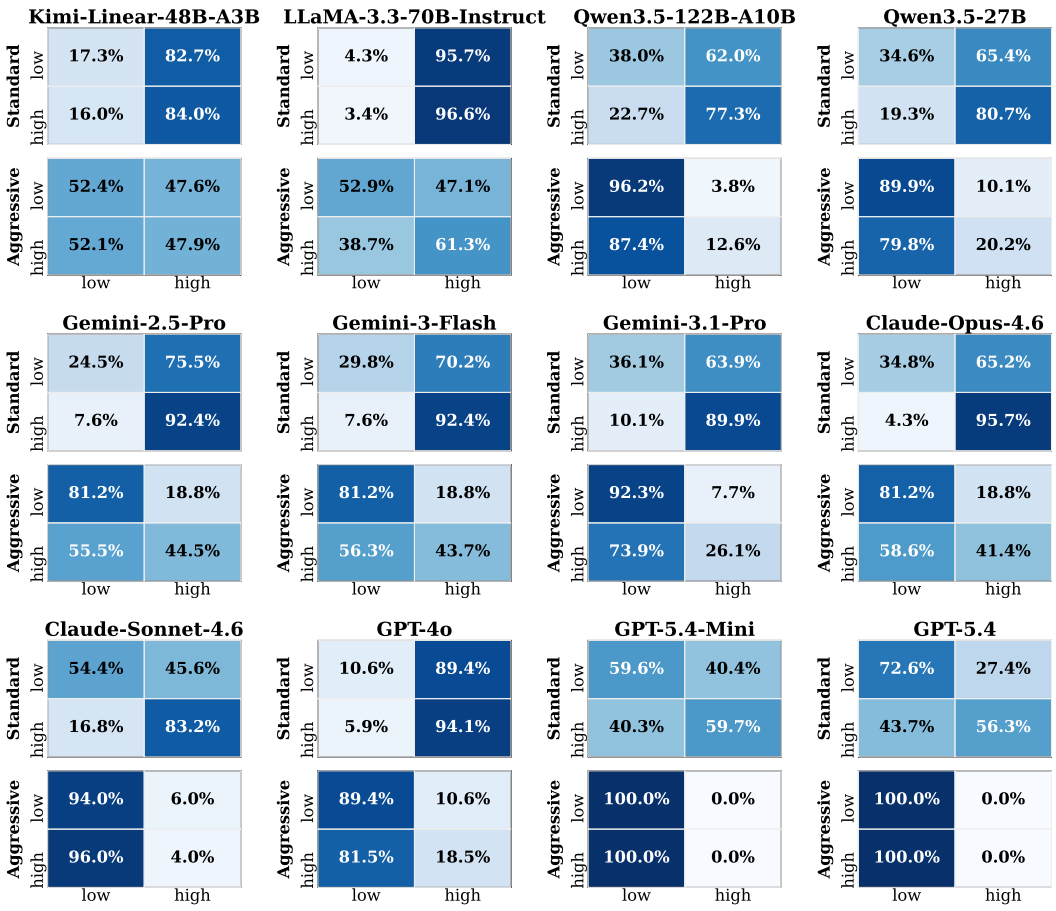


Figure 5: *Soundness judgments against ICLR-derived labels on the matched expert-annotated subset under standard and aggressive prompting.* Each model panel shows row-normalized confusion matrices under standard (top) and aggressive (bottom) prompting on the proposals retained in the expert-relabeling setup. Rows indicate ICLR-derived reference labels; columns indicate predictions. Aggressive prompting generally reduces false approvals while increasing false rejections.

These findings suggest that resolving the optimism bias may require interventions at or before pre-training (such as targeted training on scientific judgment), rather than relying on post-hoc instruction tuning alone.

B.7 Controls for Public-Corpus Contamination

A natural concern is that ICLR submissions and OpenReview discussions are public, so some evaluated models may have encountered older source papers during pretraining or supervised fine-tuning. We therefore treat contamination as a threat to validity rather than claiming it is fully eliminated. We address it in three ways. First, if a model had memorized acceptance outcomes or reviewer verdicts, that information would likely help it identify low-soundness proposals; the observed high false-positive rate is therefore conservative with respect to our main failure claim. Second, we evaluate an ICLR 2026-only split using a cutoff-restricted model subset with documented training cutoffs before September 2025 where such documentation is available, reducing the chance that these submissions appeared in training data. Third, the identifier-removal analysis in Sec. B.8 tests whether recognizable method names or self-identifying phrases drive predictions. As shown in Tab. 7, the optimism bias is largely stable across the full dataset and the 2026-only split. We will report the exact model-card or technical-report source for each cutoff in the released metadata and exclude models without sufficient cutoff documentation from cutoff-restricted claims.

Table 6: Evaluation metrics on the subset with agreement between ICLR and proposal-only expert judgments. Low R and High R denote recall on low- and high-soundness proposals. Macro F1 is the unweighted average of per-class F1 scores. Best Macro F1 per condition is **bolded**. † denotes reasoning model.

	Standard Prompt			Aggressive Prompt		
Model	**Low R ↑**	**High R ↑**	**Macro F1 ↑**	**Low R ↑**	**High R ↑**	**Macro F1 ↑**
Open-source Models						
LLaMA-3.3-70B	4.7	100.0	26.1	56.0	72.5	58.4
Kimi-Linear-48B-A3B	17.3	81.2	33.7	53.4	47.8	48.3
Qwen3.5-27B†	37.2	91.3	51.5	92.1	29.0	61.5
Qwen3.5-122B-A10B†	40.3	92.8	54.1	96.9	17.4	56.5
Closed-source Models						
GPT-4o†	11.5	100.0	32.8	90.1	23.2	56.7
GPT-5.4-mini	62.3	73.9	62.8	100.0	0.0	42.4
GPT-5.4†	75.4	72.5	**70.8**	100.0	0.0	42.4
Gemini-2.5-Pro†	26.7	100.0	45.9	83.8	53.6	68.8
Gemini-3-Flash	31.9	98.6	49.6	83.2	55.1	**69.1**
Gemini-3.1-Pro†	38.2	98.6	54.2	92.7	36.2	66.1
Claude-Sonnet-4.6†	56.7	95.5	66.4	95.5	6.2	46.8
Claude-Opus-4.6†	38.1	100.0	55.9	83.9	55.6	68.9
Mean	36.7	92.0	50.3	85.6	33.0	57.2

Table 7: Reduced-contamination-risk check using an ICLR 2026-only split. We evaluate the standard prompt on the subset of models with documented training cutoffs before the ICLR 2026 submission period, then compare their mean predictions on the full dataset and on the ICLR 2026-only split. The main observation is that the optimism bias remains stable under this reduced-contamination-risk setting: low-soundness proposals are still predicted as high soundness at a similar rate on the 2026-only split and the full dataset.

Ground Truth / Split	**Predicted Low**	**Predicted High**
Low (Full Dataset)	26.12%	73.88%
Low (ICLR 2026 Only)	22.53%	77.47%
High (Full Dataset)	8.14%	91.86%
High (ICLR 2026 Only)	7.77%	92.23%

We emphasize that perfect contamination control is difficult for any benchmark built from public corpora. For this reason, we describe this as a reduced-contamination-risk analysis rather than a contamination-free benchmark. SoundnessBench should be interpreted together with the 2026 split, the identifier-removal analysis below, and future private or continuously updated benchmark variants.

B.8 Identifier-Removal Robustness

Paper titles are never passed to models in our main evaluation. To further test whether recognizable method names or self-identifying phrases drive the results, we construct an anonymized variant by masking method names, framework names, and self-naming phrases while preserving the scientific meaning of the hypothesis and experiment design. Tab. 8 shows that this intervention has negligible effect on the aggregate confusion matrix, with changes of roughly one percentage point. This suggests that the main behavior is not primarily driven by recognizing paper identities.

Table 8: Effect of masking paper-identifying phrases. Entries are mean prediction percentages.

Ground Truth / Setting	**Predicted Low**	**Predicted High**
Low (Without Removal)	26.12%	73.87%
Low (With Removal)	25.01%	74.99%
High (Without Removal)	8.14%	91.86%
High (With Removal)	8.58%	91.42%

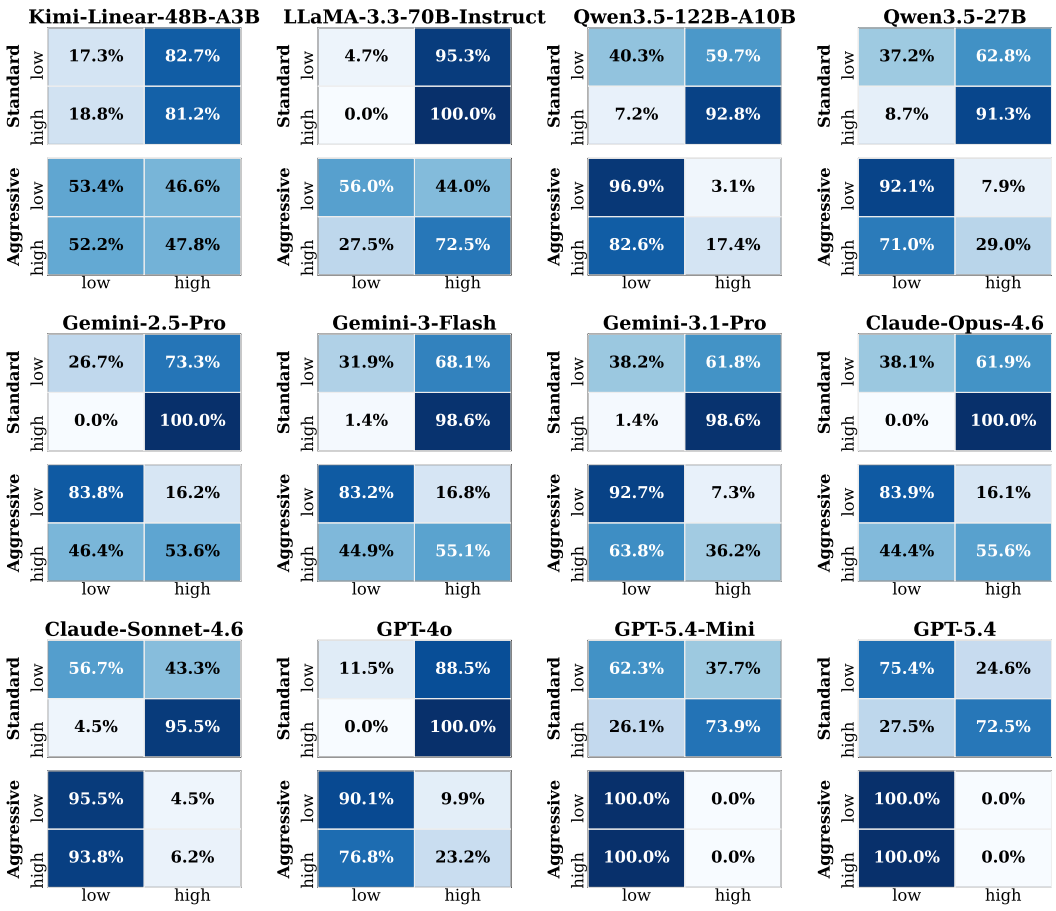


Figure 6: *Soundness judgments on the ICLR–expert agreement subset under standard and aggressive prompting.* Each model panel shows row-normalized confusion matrices under standard (top) and aggressive (bottom) prompting. Rows indicate the shared ICLR and expert reference labels; columns indicate predictions. Aggressive prompting generally reduces false approvals while increasing false rejections.

B.9 Surface-Feature Baselines

We add simple baselines to test whether shallow structural statistics alone explain the observed behavior. Always-high predicts every proposal as high soundness; always-low predicts every proposal as low soundness. We also evaluate mean-threshold classifiers based on proposal length, experiment count, and risk-factor count, with thresholds set at the midpoint between per-class means. These no-training baselines avoid fitting on the evaluation labels while still probing whether label priors or obvious structural features explain the result. A fully supervised text classifier would answer a different question—how much signal can be learned from labeled SoundnessBench examples—so we leave it to a training-split extension rather than using benchmark test labels for model selection.

The structural baselines and LLMs fail in qualitatively opposite directions: structural thresholds over-reject high-soundness proposals, whereas LLMs over-approve low-soundness proposals. This divergence suggests that shallow structural statistics alone are insufficient to explain the observed optimism bias, although richer supervised confounder models remain useful future controls.

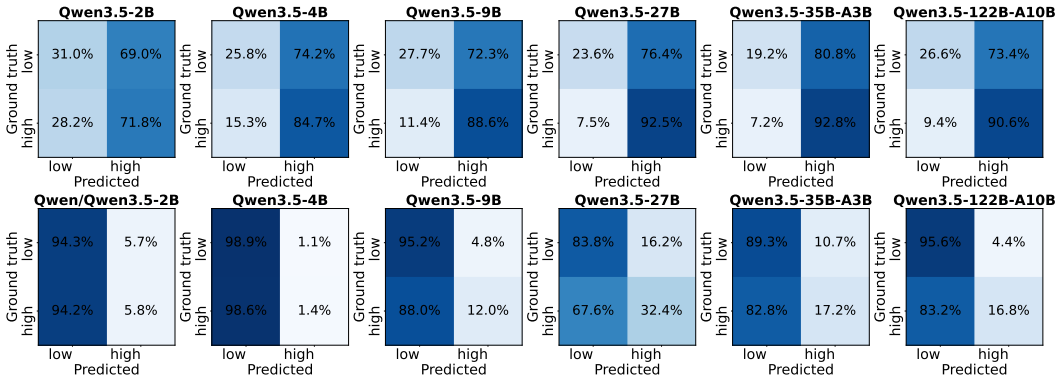


Figure 7: Confusion matrices across six Qwen3.5 model sizes (2B–122B) under standard **(top)** and aggressive **(bottom)** prompting. Under standard prompting, high-soundness recall improves with scale but low-soundness recall degrades. Larger models become more optimistic. Under aggressive prompting, models shift toward over-conservatism with no consistent improvement from scale.

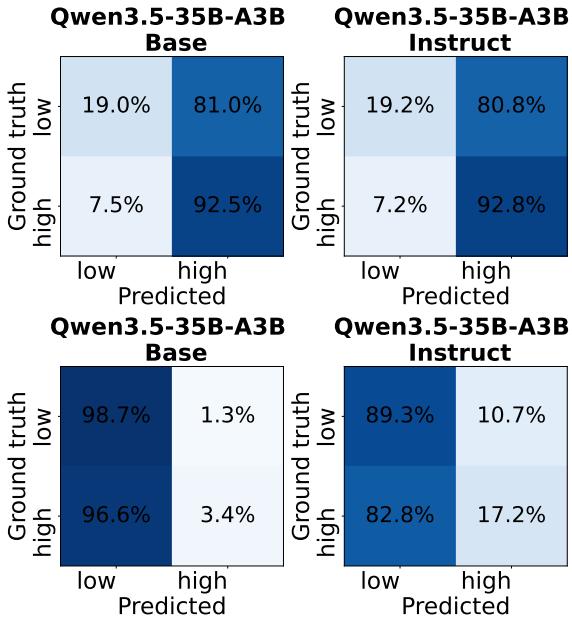


Figure 8: Base vs. instruction-tuned Qwen3.5-35B-A3B under standard **(top)** and aggressive **(bottom)** prompting. Under standard prompting, both variants show similar optimism bias, consistent with an origin earlier in training. Under aggressive prompting, the instruction-tuned model is slightly better, but both remain over-conservative.

B.10 Literature-Grounded ScholarEval Baseline

We further test whether grounding proposal judgments in retrieved prior work changes the observed failure mode. ScholarEval (Moussa et al., 2025) evaluates research ideas with a retrieval-augmented pipeline that extracts proposed methods, generates method-specific literature queries, retrieves relevant prior work before the paper cutoff date, summarizes method-level evidence, assigns each extracted method a 0–10 soundness score, and averages method scores into a final idea-level score. To adapt this output to SoundnessBench's binary high/low labels, we threshold the average soundness score. Because some extracted methods receive null scores when grounded literature is missing, we report matched subsets after filtering by coverage rate, defined as the fraction of extracted methods

Table 9: Surface-feature and trivial baselines compared with the average LLM under standard prompting.

	Low→Low	Low→High (FP)	High→Low	High→High
Always-High	0.0%	100.0%	0.0%	100.0%
Always-Low	100.0%	0.0%	100.0%	0.0%
Length Threshold	58.5%	41.5%	36.7%	63.3%
Experiment Count Threshold	64.4%	35.6%	47.7%	52.3%
Risk Factor Count Threshold	57.4%	42.6%	48.2%	51.8%
Avg. LLM (Standard Prompt)	**26.12%**	**73.87%**	**8.14%**	**91.86%**

with non-null scores. We compare ScholarEval against direct standard and aggressive prompting on the same covered examples.

Literature grounding modestly improves the balance between the two classes when paired with a strict decision threshold, but it does not eliminate the optimism bias. At thresholds 5.0–6.0, ScholarEval predicts every covered proposal as high soundness; at 7.0, it still rejects fewer than 3% of low-soundness proposals. At threshold 8.0 on the full-coverage subset, ScholarEval attains 34.0% low-soundness recall, 73.3% high-soundness recall, and 53.6% Macro F1, compared with 6.8%, 98.0%, and 47.5% under matched standard prompting and 90.3%, 26.3%, and 44.8% under matched aggressive prompting. The coverage-≥ 0.8 subset shows the same pattern: ScholarEval reaches 53.7% Macro F1, exceeding standard prompting (48.8%) and aggressive prompting (43.8%), while retaining 34.3% low-soundness recall and 73.3% high-soundness recall. Thus, retrieval grounding yields a more balanced and modestly better classifier at threshold 8.0, but its behavior is highly threshold-sensitive and it still approves roughly two-thirds of low-soundness proposals.

Table 10: ScholarEval literature-grounded soundness evaluation compared with matched direct prompting. Low R and High R denote low- and high-soundness recall; Macro F1 is the unweighted average of the low- and high-class F1 scores. Best Macro F1 within each coverage group is **bolded**.

Evaluation Mode	Threshold	Low R	High R	Macro F1
Coverage ≥ 1.0 *(n=350; high=247, low=103)*				
ScholarEval Agent	5.0	0.0%	100.0%	41.4%
ScholarEval Agent	6.0	0.0%	100.0%	41.4%
ScholarEval Agent	7.0	2.9%	99.2%	44.2%
ScholarEval Agent	8.0	34.0%	73.3%	**53.6%**
Standard Prompt	–	6.8%	98.0%	47.5%
Aggressive Prompt	–	90.3%	26.3%	44.8%
Coverage ≥ 0.8 *(n=396; high=288, low=108)*				
ScholarEval Agent	5.0	0.0%	100.0%	42.1%
ScholarEval Agent	6.0	0.0%	100.0%	42.1%
ScholarEval Agent	7.0	2.8%	99.3%	44.8%
ScholarEval Agent	8.0	34.3%	73.3%	**53.7%**
Standard Prompt	–	7.4%	98.3%	48.8%
Aggressive Prompt	–	88.9%	27.1%	43.8%

B.11 Adversarial Injection of Methodological Flaws

To test whether models attend to methodological content at all, we inject targeted flaws into 100 high-soundness proposals. Specifically, dataset names and evaluation metrics are replaced with those from a different domain while keeping the hypothesis and other content fixed, creating severe hypothesis–experiment mismatches. GPT-5.4's approval rate drops from 77.0% on the original proposals to 1.0% after injection (Tab. 11). This result rules out the strongest version of the inattention hypothesis: the model can detect severe methodological inconsistencies. However, it also sharpens the remaining failure mode: models are still insufficiently critical of subtler naturally occurring flaws

in low-soundness proposals. We report this as an initial control and will extend it to more flaw types, such as missing baselines and data leakage, in future benchmark versions.

Table 11: Adversarial injection of severe methodological mismatches into high-soundness proposals.

Model	Original Approval	Post-Injection Approval	Flip Rate
GPT-5.4	77.0%	1.0%	76.0%

B.12 High-Confidence Positive Subset for Aggressive Prompting

A remaining concern is that reviewer-derived high-soundness labels may still contain positive-label noise, even after consensus filtering. Peer-review decisions can be noisy, and prior work documents substantial variability and bias in review outcomes (Shah, 2022). This concern does not undermine the optimism-bias result, which is driven by approvals of low-soundness proposals. It is more relevant to the aggressive-prompt result that some false negatives could instead be true negatives if the aggressive prompt catches flaws missed by human reviewers.

To test this possibility, we construct a stricter high-confidence positive subset by restricting the positive class to papers that received Oral or Spotlight distinctions among Soundness positive labels data. These papers have undergone not only reviewer evaluation but also additional Area Chair and Program Committee selection, providing stronger evidence of methodological soundness. We then recompute the aggressive-prompt false-negative rate on this subset and compare it against the false-negative rate computed on the full set of benchmark positives.

Table 12: Aggressive-prompt false-negative rates on a stricter high-confidence positive subset. The subset contains high-soundness proposals whose source papers received oral, spotlight, or highlight status. Delta is the high-confidence-subset FN rate minus the all-positive benchmark FN rate.

Model	FN Rate (Oral/Spotlight)	FN Rate (All Benchmark)	Delta
LLaMA-3.3-70B-Instruct	23.2%	23.9%	−0.6%
Claude-Opus-4.6	26.9%	33.6%	-6.6%
Gemini-3-Flash	37.4%	39.6%	-2.3%
Gemini-2.5-Pro	43.4%	44.9%	-1.5%
Kimi-Linear-48B-A3B	51.5%	53.0%	-1.5%
Gemini-3.1-Pro	59.6%	64.1%	-4.5%
Qwen3.5-27B	59.6%	67.6%	−8.0%
GPT-4o	76.8%	75.0%	+1.7%
Claude-Sonnet-4.6	78.3%	81.2%	-2.9%
Qwen3.5-122B-A10B	78.8%	83.2%	−4.4%
GPT-5.4	100.0%	100.0%	0.0%
GPT-5.4-mini	100.0%	99.8%	+0.2%
Average	**61.3%**	**63.8%**	**-2.5%**

Using oral/spotlight subsets in SoundnessBench, as a stricter positive subset does not materially reduce aggressive-prompt false negatives. The average false-negative rate decreases only from 63.8% to 61.3% (a 2.5-point reduction), and several models remain highly over-conservative even on these high-confidence positives. In particular, GPT-5.4 and GPT-5.4-mini reject nearly all such proposals under aggressive prompting. These results suggest that the aggressive prompt is genuinely over-conservative, rather than merely correcting noisy high-soundness labels in SoundnessBench.

B.13 Breakdowns by Year, Primary Area, and Writing Quality

We further examine whether failures are concentrated in particular years, primary areas, or writing-quality bands. Tab. 13, Tab. 14, and Tab. 15 report low-soundness false-positive rates. The optimism bias is visible across all slices, suggesting that it is systemic rather than driven by a narrow subset of the benchmark.

Table 13: Low-soundness false-positive rate by year.

Year	Low-Soundness FP Rate
2022	67.2%
2023	79.6%
2024	70.2%
2025	68.4%
2026	78.3%
All	73.87%

Table 14: Low-soundness false-positive rate by raw ICLR primary area, restricted to areas with more than 10 benchmark pairs.

Primary Area	Low-Soundness FP Rate
Applications (e.g., speech, vision, NLP)	81.82%
Deep Learning and Representational Learning	77.37%
Machine Learning for Sciences	85.19%
Reinforcement Learning (decision/control)	75.76%
Social Aspects of Machine Learning	81.82%
Alignment, Safety, and Security	78.34%
Applications to CV, Audio, and Language	88.80%
Applications to Neuroscience and Cognitive Science	49.55%
Applications to Physical Sciences	71.05%
Applications to Robotics, Autonomy, and Planning	78.38%
Datasets and Benchmarks	72.65%
Foundation and Frontier Models	77.45%
General Machine Learning	57.95%
Generative Models	72.54%
Interpretability and Explainability	64.20%
Graphs and Topological Data Analysis	75.00%
Time Series and Sequential Data	80.74%
Learning Theory	64.06%
Neurosymbolic Methods	52.73%
Optimization	57.53%
Other Machine Learning	70.97%
Probabilistic Methods	74.05%
Reinforcement Learning	84.57%
Representation Learning for CV, Audio, and Language	80.18%
Transfer, Meta-, and Lifelong Learning	76.76%
Unsupervised, Self-, and Semi-supervised Learning	76.05%
All Primary Areas	**74.29%**

B.14 Qualitative Examples of False and True Positives

We include representative qualitative examples to illustrate what the aggregate confusion matrices count as false positives and true positives. In Figs. 9 and 13, the proposals have low mean ICLR soundness scores but are still predicted as high soundness by the evaluated models, illustrating the optimism-bias failure mode. In Fig. 17, the proposal has a high mean ICLR soundness score and is correctly predicted as high soundness, illustrating a true-positive case. We show the corresponding Gemini 3.1 Pro, GPT-5.4 Thinking, and Claude Opus 4.6 responses below each proposal.

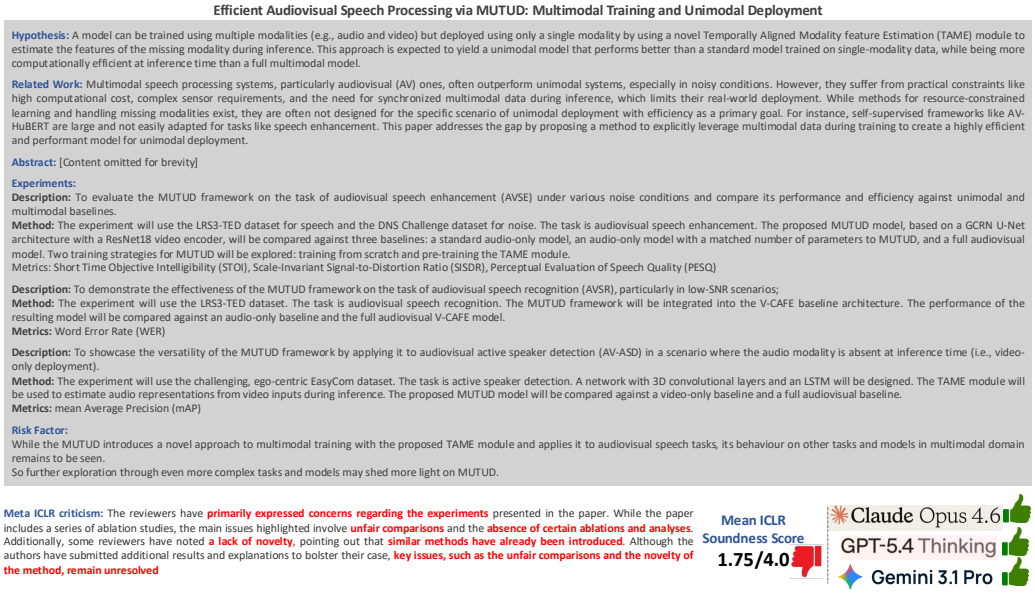


Figure 9: **False-positive example.** A low-soundness proposal is predicted as high soundness by Claude Opus 4.6, GPT-5.4 Thinking, and Gemini 3.1 Pro despite reviewer concerns about the experimental comparisons, missing ablations or analyses, and novelty.

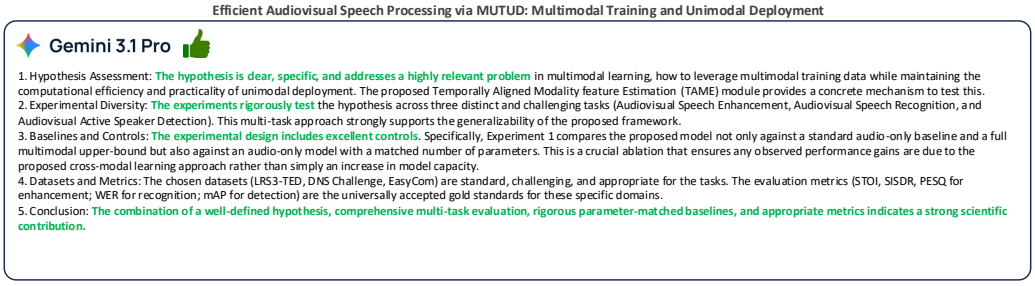


Figure 10: **Gemini 3.1 Pro response for the false-positive example in Fig. 9.** The model assigns a high-soundness judgment to the low-soundness proposal.

Efficient Audiovisual Speech Processing via MUTUD: Multimodal Training and Unimodal Deployment

GPT-5.4 Thinking

Step 1: Assess hypothesis clarity and significance. The hypothesis is clear, specific, and meaningful: train with multiple modalities, deploy with one modality by estimating missing-modality features via a temporally aligned module, with the goals of improving over unimodal-only training while retaining inference efficiency relative to full multimodal inference. This is a nontrivial and relevant research question, especially for practical settings where some modalities are unavailable or expensive at test time.
Step 2: Assess whether the experiments directly test the hypothesis. The experiments are well aligned with the hypothesis. Experiment 1 tests the core claim in AV speech enhancement under noisy conditions, including both effectiveness and efficiency comparisons. Experiment 2 tests whether the idea transfers to AV speech recognition, especially low-SNR conditions where cross-modal information is valuable. Experiment 3 further probes generality by reversing the missing modality direction (audio absent at inference) on active speaker detection. Together, these experiments evaluate both the central mechanism and its breadth across tasks.
Step 3: Evaluate baselines and controls. The proposed setup includes appropriate primary baselines: a standard unimodal baseline and a full multimodal baseline in all tasks. In Experiment 1, there is an additional audio-only baseline with matched parameter count, which is an important control because it helps separate gains due to multimodal training/distillation-like transfer from gains due merely to increased capacity. This is a strong design choice. However, the parameter-matched control is only explicitly described for Experiment 1 and not for Experiments 2 and 3, which weakens the rigor somewhat. Also, there is no mention of comparisons to other modality-dropout, hallucination, distillation, or missing-modality adaptation methods, which would be valuable related-work baselines if such methods exist for the tasks.
Step 4: Evaluate ablations and methodological depth. The experiment plan includes two training strategies for MUTUD in Experiment 1 (from scratch and pre-training TAME), which is a useful ablation on optimization/training procedure. But the overall ablation plan is limited. There is no explicit mention of ablating temporal alignment, feature estimation quality, where TAME is inserted, sensitivity to modality corruption levels, or inference-time cost decomposition. Since the core novelty is the TAME module, stronger ablations would normally be expected to establish why this design works and whether temporal alignment specifically matters.
Step 5: Evaluate metrics. The metrics are appropriate and standard for each task: STOI, SI-SDR, and PESQ for speech enhancement; WER for speech recognition; and mAP for active speaker detection. These metrics are strong choices and directly measure task performance. The hypothesis also includes an efficiency claim, and Experiment 1 mentions efficiency comparison, but no explicit efficiency metrics are listed such as FLOPs, latency, memory, or throughput. That omission is important because computational efficiency is part of the claimed contribution.
Step 6: Evaluate dataset/task selection. The datasets are credible and relevant: LRS3-TED and DNS for AVSE, LRS3-TED for AVSR, and EasyCom for AV-ASD. The tasks are challenging and diverse, which strengthens external validity. Testing across enhancement, recognition, and detection is a notable strength and supports claims of generality.
Step 7: Check for possible methodological weaknesses. The proposal appears methodologically sound at a high level, but some critical rigor details are absent: no mention of statistical significance, repeated runs/seeds, hyperparameter selection fairness, train/validation/test protocol details, or robustness analyses across degrees of modality mismatch/missingness. These omissions are common in summaries, but they prevent calling the plan maximally rigorous. Still, there are no obvious fundamental flaws in the proposed methodology.
Step 8: Weigh strengths versus weaknesses for bucket assignment. Strengths: clear and meaningful hypothesis, direct experimental alignment, multiple tasks, strong standard metrics, sensible baselines including a parameter-matched control in at least one experiment, and some training ablation. Weaknesses: limited explicit ablation of the novel component, incomplete efficiency evaluation despite an efficiency claim, lack of explicit competing-method baselines beyond standard uni/full-multimodal models, and missing details on statistical rigor. Despite these weaknesses, the overall plan is substantially stronger than a weak or flawed submission and reflects a serious effort to validate the idea.
Step 9: Final bucket decision. I assign this to the **high rigor bucket because the hypothesis is clear and meaningful, the experiments are largely well controlled and appropriately targeted, and the evaluation spans multiple tasks with suitable metrics and baselines**. The rigor is not perfect, but it is **sufficient to qualify as strong rather than weak.**

Figure 11: **GPT-5.4 Thinking response for the false-positive example in Fig. 9.** The model assigns a high-soundness judgment to the low-soundness proposal.

Efficient Audiovisual Speech Processing via MUTUD: Multimodal Training and Unimodal Deployment

Claude Opus 4.6

Let me analyze this hypothesis-experiment pair step by step.

Hypothesis Assessment: The hypothesis has two clear claims: (1) Swin Transformer captures better global visual-language features than CNNs for ZS3, and (2) pixel-text score maps with modified decision boundaries reduce seen-bias and improve unseen class performance. The hypothesis is specific and testable, though not particularly novel using transformers over CNNs is a relatively straightforward architectural upgrade, and the seen-bias problem is well-known in zero-shot learning.
Experiment Assessment:
Experiment 1 & 2: These evaluate the full framework on two standard benchmarks (PASCAL VOC and PASCAL Context) with multiple seen/unseen splits. They compare against relevant baselines (DeViSE, SPNet, ZS3Net, CSRL, JoEm). The metrics (mIoUs, mIoUu, hIoU) are appropriate and standard for ZS3 evaluation. Multiple splits provide robustness to the evaluation. This is solid experimental design.
Experiment 3: The ablation study is well-designed it isolates the two main contributions (transformer backbone vs. CNN, and the auxiliary score map loss) by testing combinations. This directly addresses both parts of the hypothesis.
Strengths: - Clear, testable hypothesis - Multiple benchmarks and splits - Appropriate baselines from the literature - Standard and appropriate metrics including harmonic mean (which captures the seen-unseen trade-off)\n- Ablation study that isolates individual contributions
Weaknesses: - Only uses swin-tiny; no exploration of other transformer sizes or variants - The ablation is only on one split (unseen-6 of PASCAL Context), which limits generalizability of ablation conclusions - No statistical significance tests or variance reporting mentioned - Missing some more recent baselines potentially - Only two datasets, both from PASCAL family no COCO-stuff or other diverse benchmarks - The novelty is somewhat incremental (replacing CNN with Swin Transformer + an auxiliary loss)
Overall, **the experimental design is reasonably rigorous with proper baselines, ablations, and standard metrics across multiple settings. The methodology is sound,** though could be more comprehensive. The ablation study directly addresses the hypothesis components. I'd**classify this as leaning toward high rigor the experiments are controlled, include appropriate baselines and ablations, and use suitable metrics, even if there are minor gaps.**

Figure 12: **Claude Opus 4.6 response for the false-positive example in Fig. 9.** The model assigns a high-soundness judgment to the low-soundness proposal.

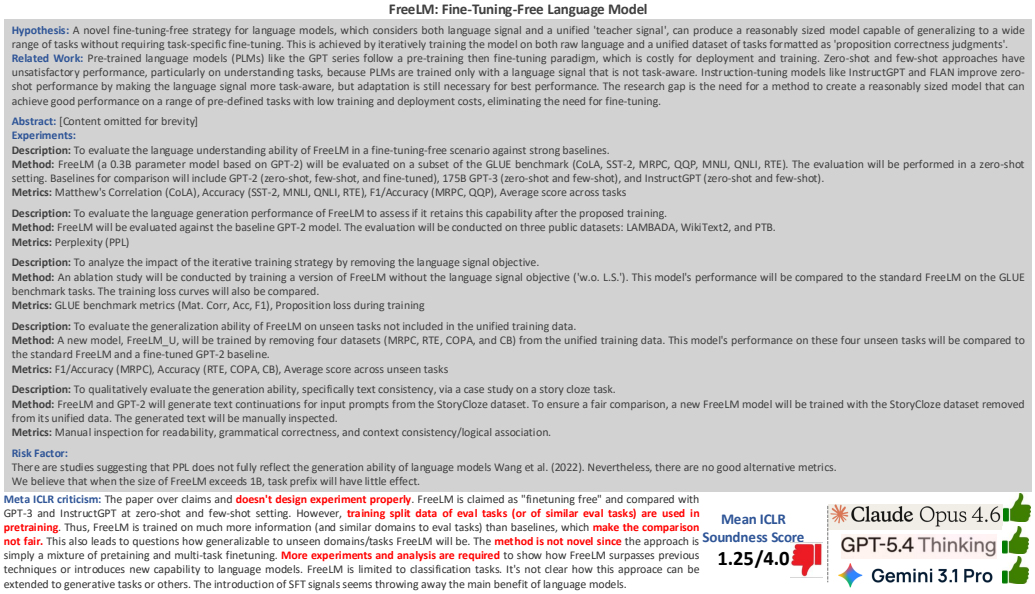


Figure 13: **Second false-positive example.** A low-soundness proposal is predicted as high soundness by Claude Opus 4.6, GPT-5.4 Thinking, and Gemini 3.1 Pro despite reviewer concerns that the paper overclaims, lacks experimental support, and uses evaluation-task pretraining data that weakens the comparison.

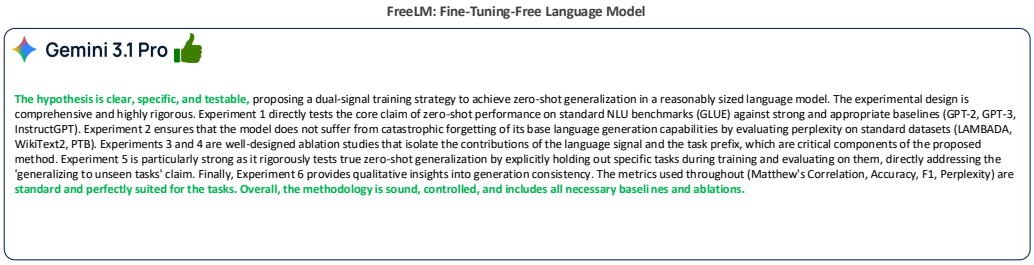


Figure 14: **Gemini 3.1 Pro response for the second false-positive example in Fig. 13.** The model assigns a high-soundness judgment to the low-soundness proposal.

FreeLM: Fine-Tuning-Free Language Model

GPT-5.4 Thinking

Step 1: Assess the hypothesis clarity and scientific value.
The hypothesis is reasonably clear and non-trivial: a fine-tuning-free training strategy that combines a language modeling signal with a unified teacher signal in proposition-judgment format can yield a relatively small model that generalizes across tasks without task-specific fine-tuning. This is a meaningful claim because it targets transfer, zero-shot generalization, and training efficiency. So on hypothesis formulation alone, this is stronger than a vague or trivial proposal.
Step 2: Assess whether the experiments directly test the core claim. The core claim has three main parts: (a) strong zero-shot task generalization, (b) retention of language modeling/generation ability, and (c) benefit of the proposed training components. The proposed experiments map onto these parts fairly well:- Experiment 1 tests zero-shot task performance on GLUE against several strong baselines.- Experiment 2 checks whether generation capability is preserved.- Experiments 3 and 4 provide ablations for the language signal and task prefix.- Experiment 5 tests generalization to held-out tasks.- Experiment 6 gives a qualitative generation case study.This is a coherent experimental package and is a positive sign for rigor.
Step 3: Evaluate baseline quality and fairness.The baseline set in Experiment 1 is relatively strong and includes GPT-2 in zero-shot/few-shot/fine-tuned settings, plus GPT-3 and InstructGPT in zero-shot/few-shot settings. This indicates awareness of relevant comparison points. However, there are fairness concerns:- Model scales differ dramatically (0.3B vs 175B), making comparisons informative but not fully controlled.- It is not stated whether prompts, decoding settings, or evaluation protocols are standardized across models.- The proposal omits comparisons to other instruction-tuned or multi-task small/medium models that may be more directly relevant than GPT-3.Even so, the inclusion of multiple baseline modes is better than a weak experimental setup.
Step 4: Evaluate controls and ablations.The submission includes at least two direct ablations tied to the method design:- removing the language signal objective,- removing the task prefix.These are meaningful controls because they isolate important components of the proposed method. Experiment 5 also acts as a form of generalization control by withholding tasks from the unified training mixture. This strengthens the scientific design considerably.
Step 5: Evaluate metrics and task coverage.The metrics are standard and appropriate for the chosen tasks:- Matthews correlation for CoLA,- accuracy for classification tasks,- F1/accuracy for MRPC and QQP,- perplexity for language modeling datasets.These are suitable metrics. The use of multiple GLUE tasks plus LAMBADA/WikiText2/PTB provides reasonable coverage of understanding and generation. However, the qualitative StoryCloze evaluation lacks a rigorous human evaluation protocol or automatic metric, so that component is weak. Still, it is only one part of the overall suite.
Step 6: Identify missing rigor elements. Despite the strengths, some methodological details are missing that reduce rigor:- No mention of statistical significance, confidence intervals, or multiple random seeds.- No explicit discussion of hyperparameter tuning fairness across baselines.- The held-out-task experiment is useful, but evaluating only four removed tasks may not fully establish broad generalization.- The unified proposition-format teacher signal may create train-test formatting advantages, and the proposal does not clearly discuss contamination prevention or prompt design sensitivity.- Experiment 6 relies on manual inspection without a defined protocol, raters, or agreement measures. These omissions prevent the plan from being maximally rigorous.
Step 7: Decide between the available buckets. Given the forced choice between only 'low' and 'high', this proposal is closer to 'high' than 'low'. **The hypothesis is clear and meaningful, the experiments are aligned with the claim, there are strong baselines, standard metrics, and nontrivial ablations**. While there are notable missing details and some fairness/statistical concerns, the overall design is substantially more rigorous than a weak or fundamentally flawed study.
Step 8: Confidence assessment. **I am fairly confident in this judgment** because the hypothesis and experiment plan are understandable and map well to standard evaluation practice. However, confidence is not maximal because several implementation details that matter for rigor are unspecified, and I cannot verify whether the comparisons would be fully fair in practice.

Figure 15: **GPT-5.4 Thinking response for the second false-positive example in Fig. 13.** The model assigns a high-soundness judgment to the low-soundness proposal.

FreeLM: Fine-Tuning-Free Language Model

Claude Opus 4.6

Let me analyze this hypothesis-experiment pair systematically:
Hypothesis Analysis: The hypothesis proposes a fine-tuning-free strategy for language models that combines language signal with a unified 'teacher signal' (proposition correctness judgments) to create a model that generalizes across tasks without task-specific fine-tuning. This is a meaningful and clear research direction - it addresses the practical problem of requiring fine-tuning for each task, which is computationally expensive. The hypothesis is specific enough to be testable.
Experiment Analysis:1. Experiment 1 (GLUE evaluation): Strong design - evaluates on a standard benchmark (GLUE), includes multiple baselines (GPT-2 zero/few-shot/fine-tuned, GPT-3 175B zero/few-shot, InstructGPT), and uses appropriate metrics. The comparison against much larger models (175B GPT-3) while being 0.3B is interesting but also potentially misleading if not contextualized properly. 2. Experiment 2 (Language generation): Appropriate check to ensure the model doesn't lose language generation capability. Uses standard perplexity metrics on standard datasets. However, only compared against GPT-2, not broader baselines. 3. Experiment 3 (Ablation - language signal): Good ablation study that isolates the contribution of the language signal objective. Includes training loss curves for additional insight.4. Experiment 4 (Ablation - task prefix): Good ablation that tests another component of the proposed method. 5. Experiment 5 (Generalization to unseen tasks): Strong experiment that directly tests the generalization claim by holding out tasks during training. Includes appropriate baselines. 6. Experiment 6 (Qualitative case study): Qualitative evaluation is weaker but provides complementary evidence. The design is thoughtful (removing StoryCloze from training data for fair comparison).
Strengths: - Multiple ablation studies - Generalization test on held-out tasks - Comparison against strong baselines including much larger models - Standard benchmarks and metrics - Both quantitative and qualitative evaluation - The model size (0.3B) vs baselines (175B) comparison is interesting
Weaknesses:- The model is quite small (0.3B) and comparisons against 175B GPT-3 may not be entirely fair (though this could be a feature showing efficiency) - Qualitative evaluation in Experiment 6 is subjective - Limited to GLUE tasks for most evaluations - The 'proposition correctness judgment' framing is somewhat similar to existing instruction-tuning approaches, and the novelty distinction could be clearer - No comparison against other instruction-tuning or prompt-tuning methods of similar scale
Overall Assessment: **The experimental design is reasonably rigorous with appropriate ablations, baselines, and metrics. The hypothesis is clear and the experiments are designed to test it systematically.** The work has some limitations but **overall represents a solid scientific contribution with controlled experiments and meaningful baselines. I'd classify this as 'high' rigor,** though not exceptional.

Figure 16: **Claude Opus 4.6 response for the second false-positive example in Fig. 13.** The model assigns a high-soundness judgment to the low-soundness proposal.



Figure 17: **True-positive example.** A high-soundness proposal is correctly predicted as high soundness by Claude Opus 4.6, GPT-5.4 Thinking, and Gemini 3.1 Pro, matching the favorable reviewer assessment of the proposed method and evaluation plan.

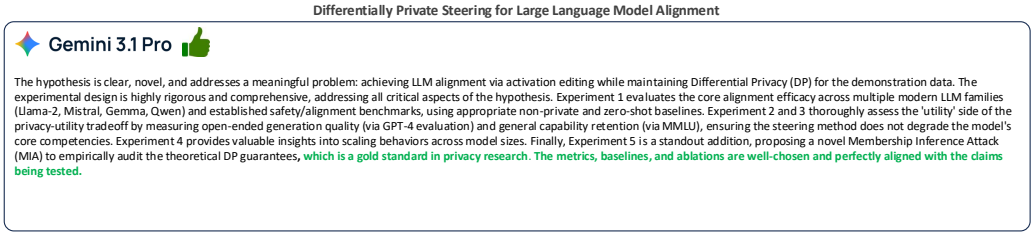


Figure 18: **Gemini 3.1 Pro response for the true-positive example in Fig. 17.** The model assigns a high-soundness judgment to the high-soundness proposal.

Differentially Private Steering for Large Language Model Alignment

GPT-5.4 Thinking

Step 1: Assess the hypothesis clarity and significance.The hypothesis is clear, specific, and meaningful. It makes a concrete claim: activation editing can be combined with differential privacy by perturbing steering vectors with calibrated Gaussian noise, yielding a method that preserves alignment behavior while providing formal privacy guarantees and minimally harming general model capabilities. This is not trivial; it targets an important and underexplored intersection of LLM alignment and privacy.
Step 2: Assess whether the experiments are aligned with the hypothesis.The experiments map well onto the core claims:- alignment effectiveness is tested in Experiment 1,- open-ended behavior quality is tested in Experiment 2,- preservation of general capabilities is tested in Experiment 3,- scaling and privacy-utility tradeoff are tested in Experiment 4,- privacy leakage and empirical auditing are tested in Experiment 5.This is a strong sign of rigor because the empirical plan covers the major components of the hypothesis rather than only a narrow slice.
Step 3: Evaluate baselines and controls.The proposal includes relevant non-private activation-editing baselines (Mean Steering, PCA Steering) and zero-shot unsteered baselines. These are appropriate comparisons for showing the cost of privacy and the value of steering. The inclusion of both private and non-private methods is important for isolating the effect of DP noise. However, the baseline set is somewhat narrow: there is no comparison against stronger alignment methods such as fine-tuning-based private methods, RLHF-style alternatives, or other privacy-preserving adaptation techniques if such exist. This weakens completeness somewhat, but does not fundamentally undermine the design.
Step 4: Evaluate the metrics.Some metrics are appropriate and standard:- MCQ accuracy for alignment benchmark choices,- MMLU performance for general capability retention,- privacy attack metrics such as FPR/FNR.But there are also notable weaknesses:- Experiment 2 relies on GPT-4 qualitative evaluation, which can be useful but is subjective and vulnerable to evaluator bias unless carefully rubricized and validated.- Experiment 3 uses average probability on the correct answer for MMLU instead of the more standard accuracy; this is not necessarily wrong, but it is less interpretable and may complicate comparisons.- The reported privacy metric 'empirical epsilon' from an attack is not equivalent to the formal DP guarantee and could be misleading if not clearly separated from theoretical privacy accounting.Overall, the metrics are mostly sensible but not uniformly strong.
Step 5: Evaluate methodological rigor of the privacy claim.This is the most important issue. The hypothesis claims formal privacy guarantees via DP, and the experiment description says Gaussian noise is calibrated and privacy leakage will be audited empirically. This suggests awareness of both theoretical and empirical privacy evaluation. However, the experiment description does not specify crucial DP methodology details such as:- the neighboring dataset definition,- sensitivity analysis of steering vector computation,- privacy accounting across tasks/models/layers,- whether hyperparameter tuning consumes privacy budget,- composition across repeated evaluations,- clipping or boundedness assumptions required for Gaussian mechanisms.Without these details, the experimental plan supports the idea of evaluating a DP method, but does not fully demonstrate that the formal privacy claim is implemented rigorously. This is a significant concern, though the presence of both formal-DP intent and attack-based auditing is better than many weak proposals.
Step 6: Evaluate breadth and robustness.The plan is fairly broad:- multiple model families and sizes,- multiple alignment dimensions,- both closed-form MCQ and open-ended evaluations,- general capability tests,- scaling studies,- adversarial privacy audits.This breadth is a strong marker of rigor. The scaling study is especially helpful because it tests whether the proposed mechanism behaves systematically rather than only on a single model.
Step 7: Identify missing ablations or controls.Important ablations are not explicitly described:- varying privacy levels/noise scales and plotting privacy-utility curves more systematically,- layer-wise steering choices,- number of private demonstrations,- effect of clipping or vector normalization,- sensitivity to random seeds,- comparison of different attack strengths in the privacy audit. These omissions reduce rigor. In particular, for a DP method, one would expect explicit sweeps over epsilon/delta and confidence intervals across seeds.
Step 8: Judge whether the flaws are fundamental or moderate.The flaws are moderate rather than fundamental. The proposal has a strong central hypothesis, relevant baselines, multi-angle evaluation, and a privacy-specific attack audit. The main weaknesses are incomplete specification of the DP accounting and some evaluation choices that are weaker than ideal. Still, the overall design is much stronger than a low-rigor submission, which would typically have vague claims, missing baselines, or no real privacy evaluation.
Step 9: Final bucket decision. Despite some methodological gaps, **the proposal demonstrates a clear and meaningful hypothesis and a generally rigorous experimental plan with appropriate controls, multiple evaluations, and explicit attention to privacy-utility tradeoffs and leakage auditing. Therefore it fits the high rigor bucket better than low.**
Step 10: Confidence assessment. I am confident but not absolutely certain. The classification depends partly on assuming the formal DP mechanism is implemented correctly, while the provided description omits some technical details needed to verify that. So confidence is 4 rather than 5.

Figure 19: **GPT-5.4 Thinking response for the true-positive example in Fig. 17.** The model assigns a high-soundness judgment to the high-soundness proposal.

Differentially Private Steering for Large Language Model Alignment

Claude Opus 4.6

Let me evaluate this hypothesis-experiment pair step by step.
Hypothesis Assessment: The hypothesis is clear, specific, and meaningful. It proposes combining activation editing (steering vectors) with differential privacy (calibrated Gaussian noise) to create a method (PSA) that provides formal privacy guarantees while maintaining alignment quality. This is a well-defined and non-trivial contribution at the intersection of privacy and alignment.
Experiment 1 (Alignment Performance): This is well-designed with multiple models (4 LLMs), multiple benchmarks (7 datasets from Anthropic's collection), and appropriate baselines (non-private Mean Steering, PCA Steering, zero-shot). The metric (accuracy on MCQ) is standard and appropriate. Good breadth of evaluation. Experiment 2 (Text Generation Quality): Evaluates open-ended generation using GPT-4 as an evaluator, which is a reasonable approach for qualitative assessment. Compares against the same baselines. However, it's limited to one model (Llama-2-Chat 7B). Using LLM-as-judge is now standard practice, though it has known limitations. Experiment 3 (General Capabilities): Tests for capability degradation using MMLU, which is a standard benchmark for this purpose. This is an important control experiment to show that privacy-preserving steering doesn't harm general knowledge. Limited to one model though. Experiment 4 (Scaling Behavior): Investigates how PSA scales with model size using the Qwen-2.5 series (0.5B to 7B). This is a thoughtful experiment that examines the privacy-utility tradeoff across scales. Uses multiple metrics. However, it's limited to one dataset (Refusal) and one model family. Experiment 5 (Privacy Auditing via MIA): This is a particularly strong experiment. Developing a novel membership inference attack tailored to activation editing is a meaningful contribution. The canary-based approach is well-motivated. Measuring empirical epsilon alongside FPR/FNR provides both practical and theoretical privacy assessment. Testing on both private and non-private methods serves as a good control.
Strengths: - Multiple models and benchmarks\n- Appropriate baselines (non-private counterparts and zero-shot) - Multiple evaluation dimensions (alignment, generation quality, general capabilities, scaling, privacy) - Novel privacy auditing experiment (MIA) - Formal privacy guarantees (DP) combined with empirical validation - Clear metrics for each experiment
Potential Weaknesses: - Some experiments are limited to a single model (Experiments 2, 3) - Experiment 4 uses only one dataset - No comparison with other private alignment methods (e.g., DP-SGD fine-tuning), though this may be because the contribution is specifically about activation editing - The GPT-4 evaluation in Experiment 2 could benefit from more structured evaluation criteria
Overall Assessment: **The experimental design is comprehensive, well-controlled, and addresses the hypothesis from multiple angles**. The inclusion of a novel MIA for privacy auditing is a strong methodological contribution. **The baselines are appropriate, and the metrics are suitable. The breadth across models, datasets, and evaluation dimensions demonstrates rigor.** While there are minor limitations (some experiments limited to one model), **the overall experimental framework is thorough and well-designed.**

Figure 20: **Claude Opus 4.6 response for the true-positive example in Fig. 17.** The model assigns a high-soundness judgment to the high-soundness proposal.

Table 15: Low-soundness false-positive rate by writing-quality band, using 2024–2026 subset data where presentation scores are available.

Writing Quality Band	**Low-Soundness FP Rate**
Low (≤ 2)	71.4%
Mid (> 2 and ≤ 3)	81.0%
High (> 3)	76.7%
All	73.87%