

# Do LLMs Favor LLMs? Quantifying Interaction Effects in Peer Review

VIBHHU SHARMA, Cornell University, USA

THORSTEN JOACHIMS, Cornell University, USA

SARAH DEAN, Cornell University, USA

There are increasing indications that LLMs are not only used for producing scientific papers, but also as part of the peer review process. In this work, we provide the first comprehensive analysis of LLM use across the peer review pipeline, with particular attention to interaction effects: not just whether LLM-assisted papers or LLM-assisted reviews are different in isolation, but whether LLM-assisted reviews evaluate LLM-assisted papers differently. In particular, we analyze over 125,000 paper-review pairs from ICLR, NeurIPS, and ICML. We initially observe what appears to be a systematic interaction effect: LLM-assisted reviews seem especially kind to LLM-assisted papers compared to papers with minimal LLM use. However, controlling for paper quality reveals a different story: LLM-assisted reviews are simply more lenient toward lower quality papers in general, and the over-representation of LLM-assisted papers among weaker submissions creates a spurious interaction effect rather than genuine preferential treatment of LLM-generated content. By augmenting our observational findings with reviews that are fully LLM-generated, we find that fully LLM-generated reviews exhibit severe rating compression that fails to discriminate paper quality, while human reviewers using LLMs substantially reduce this leniency. Finally, examining metareviews, we find that LLM-assisted metareviews are more likely to render accept decisions than human metareviews given equivalent reviewer scores, though fully LLM-generated metareviews tend to be harsher. This suggests that meta-reviewers do not merely outsource the decision-making to the LLM. These findings provide important input for developing policies that govern the use of LLMs during peer review, and they more generally indicate how LLMs interact with existing decision-making processes.

CCS Concepts: • **Applied computing** → **Publishing**; • **Human-centered computing** → *Empirical studies in collaborative and social computing*; • **Computing methodologies** → *Causal reasoning and diagnostics*; Natural language generation.

Additional Key Words and Phrases: peer-review, bias, large language model, causal, causal inference, academic publishing, machine learning, interaction

## ACM Reference Format:

Vibhhu Sharma, Thorsten Joachims, and Sarah Dean. 2026. Do LLMs Favor LLMs? Quantifying Interaction Effects in Peer Review. In *The 2026 ACM Conference on Fairness, Accountability, and Transparency (FAccT '26)*, June 25–28, 2026, Montreal, QC, Canada. ACM, New York, NY, USA, 39 pages. <https://doi.org/10.1145/3805689.3806505>

## 1 Introduction

Since ChatGPT's release in late 2022, LLMs have become a fixture in research workflows, with 81% of researchers now incorporating them into some aspect of their work (Liao et al. [11]). Recent work has shown evidence of an increasing trend in LLM use by both authors and reviewers at conferences: Liang et al. [9] found that up to 17% of reviews at top machine learning conferences show evidence of LLM modification, while Liang et al. [10] observed that the abstracts of up to 17.5% of recent Computer Science papers showed evidence of LLM use—a significantly higher number than that for papers in other domains. Anecdotally, some researchers have attributed

---

Authors' Contact Information: Vibhhu Sharma, Cornell University, Ithaca, USA, [vs552@cornell.edu](mailto:vs552@cornell.edu); Thorsten Joachims, Cornell University, Ithaca, USA; Sarah Dean, Cornell University, Ithaca, USA.



This work is licensed under a Creative Commons Attribution 4.0 International License.

*FAccT '26, Montreal, QC, Canada*

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2596-8/2026/06

<https://doi.org/10.1145/3805689.3806505>




Fig. 1. (a) LLM use trends over time at ICLR. (b-c) Rating distributions for human vs. LLM-aided papers and reviews for ICLR 2024 and 2025 submissions. LLM-aided papers receive a higher distribution of lower scores while LLM-aided reviews more frequently give middling scores relative to their fully human counterparts. Here LLM-aided means the text showed sufficient similarity to LLM vocabulary ( $\alpha > 0.15$  in Section A.3).

subpar reviews they have received in recent years to LLM use, indicating a general sentiment of distrust in the use of LLMs in peer review.

Figure 1a shows the year by year increase in LLM use, in both papers and reviews submitted to the International Conference on Learning Representations (ICLR), by using the technique from Liang et al. [9] to estimate the fraction  $\alpha$  of LLM generated text in a document (described in Section 3). Notably, LLM use is substantially more common in writing reviews than in writing papers: in the latest edition of the conference, 3.31% of papers and 26.65% of reviews show substantial LLM modification (detected fraction over  $\alpha_{\text{threshold}} = 15\%$ ). The increase in LLM use in review writing is particularly steep, almost doubling in magnitude in a year.

Both LLM-assisted reviews and LLM-assisted papers show interesting differences from those that do not have detected LLM assistance (henceforth referred to as "human"). LLM-assisted papers are frequently given lower scores (Fig. 1b), with '3' being the most commonly assigned quality score, relative to a mode of '6' for human-written papers. Meanwhile LLM assisted reviews concentrate toward assigning intermediate ratings and give fewer harsh ratings (Fig. 1c), suggesting a tendency to hedge rather than provide extreme assessments on either side.

An interesting question that naturally arises from LLM use on both sides of the peer review process regards the existence of a systemic interaction effect between LLM/Human papers and LLM/Human Reviews. More directly, *does a human reviewer behave differently when presented a fully human paper vs an LLM-assisted paper?* Similarly, *does an LLM-assisted reviewer behave differently when presented an LLM-assisted paper vs. a fully human paper?* Interestingly, ICML has instituted a new policy for their 2026 edition (ICML Policy [2]) requiring both reviewers and authors to declare either conservative or permissive LLM use, with reviewer-author pairing done based on these declarations. If any sort of systematic bias does exist in LLM-LLM interactions within peer review, such a policy could inadvertently amplify it. Moreover, authors could potentially game the system if they understand these interaction effects and their implications for review outcomes.

As an initial datapoint, Table 1 divides (paper, review) pairs into 4 quadrants based on LLM usage—both human-written, human paper with LLM-assisted review, LLM-assisted paper with human review, and both LLM-assisted—and presents the mean and standard deviations of the scores for these pairs. We find that LLM-assisted papers receive lower ratings on average while LLM-assisted reviews give higher ratings. Both human and LLM-assisted papers receive higher scores when reviewed by an LLM-assisted reviewer, but notably the magnitude differs: *a 0.25-point jump for human papers versus 0.63 points for LLM-assisted papers*. Meanwhile, reviewers show higher confidence for LLM-assisted papers, with human reviewers slightly more confident than LLM-assisted reviewers.

Quadrant ( <i>Paper, Review</i> )	Rating	Confidence	Presentation	Contribution	Soundness
(Human, Human)	5.12 ± 1.74	3.66 ± 0.80	2.61 ± 0.76	2.35 ± 0.70	2.59 ± 0.70
(Human, LLM)	5.37 ± 1.55	3.61 ± 0.82	2.72 ± 0.67	2.46 ± 0.67	2.68 ± 0.63
(LLM, Human)	4.15 ± 1.88	3.85 ± 0.80	2.30 ± 0.84	2.04 ± 0.74	2.23 ± 0.76
(LLM, LLM)	4.78 ± 1.87	3.80 ± 0.79	2.48 ± 0.76	2.29 ± 0.75	2.48 ± 0.72

Table 1. Summary Statistics for (Paper, Review) pairs from ICLR 2024 and 2025 combined. (Paper, Review) pairs are divided into 4 quadrants based on the LLM status of the paper and the review.  $\alpha$  value of 0.15 is used as a threshold: papers and reviews are classified as LLM-assisted if they have  $\alpha > 0.15$ . Cell values are presented as Mean ± Std Dev.

These observations suggest potential systematic effects worth exploring, though summary statistics alone cannot verify this due to confounding. It is wholly possible that the papers in one quadrant are fundamentally different papers from the papers in other quadrants, making rating comparisons for these papers moot. This necessitates a causally grounded analysis to eliminate paper-level confounding factors.

Motivated by the above observations, our paper uses observational data (paper text, metadata, reviews, scores and metareviews) from top machine learning conferences over the last 3 years, in addition to synthetic data generated by prompting LLMs, to answer the following key research questions:

- (1) *Are LLM-aided reviews kinder to LLM-aided papers on average?*
- (2) *Does the kindness of LLM-aided reviews to LLM-aided papers vary with paper quality?*
- (3) *How does an LLM-aided review differ from a fully LLM-generated review?*
- (4) *How does LLM-assistance in metareviewing relate to a paper’s final acceptance decision?*

Our work differs from prior studies in several key ways. Unlike research that examines LLM use in papers or reviews in isolation, we provide the first comprehensive analysis of interaction effects between LLM use on both sides of the peer-review process. We combine observational analysis of over 125,000 paper-review pairs from multiple top-tier conferences with controlled synthetic experiments to disentangle genuine LLM effects from the effect of the human in the loop. Our findings on metareview leniency, the quality-dependent nature of LLM review bias, and the moderating role of human-in-the-loop mediation are novel contributions to understanding how AI tools are reshaping research dissemination and decision making.

## 2 Related Work

The rapid adoption of LLMs in research has motivated the development of detection methods to quantify their use. Liang et al. [9] introduced a statistical approach finding that up to 17% of reviews at top ML conferences showed significant LLM modification, while Liang et al. [10] documented 17.5% of recent Computer Science paper abstracts showed evidence of LLM use. However, they do not study how these distinct forms of LLM use interact with each other.

Several studies have examined how LLM-generated reviews differ from human reviews. Zhu et al. [20] identified systematic biases in LLM reviewers, including inflated ratings for weaker papers and reduced sensitivity to paper quality. The susceptibility of LLM reviews to manipulation has been documented by Ye et al. [18] and Keuper [6], who demonstrated successful prompt injection attacks. However, these studies do not explicitly account for LLM-aided reviews. Kocak et al. [7] advocate for LLMs as complementary tools rather than replacements for human judgment.

Beyond peer review, several studies have investigated whether LLMs exhibit bias toward their own outputs. Zheng et al. [19] adopt the term ‘self-enhancement bias’ from social cognition literature. Panickssery et al. [14] found evidence that LLMs preferentially rate their own generated text higher, while Laurito et al. [8] extended this to binary choice scenarios, showing LLM-based assistants consistently prefer LLM-presented options. This is

an important motivation for our work, and we argue that the setting of peer review serves as a valuable testbed for understanding the implications of two-sided LLM use in practice. OpenReview’s transparency policies provide rich, structured data that allows us to study authentic high-stakes author-reviewer interactions where both parties may have used LLMs.

### 3 Datasets

We conduct experiments on paper and review data from OpenReview for ICLR (2024, 2025), the Conference on Neural Information Processing Systems (NeurIPS) (2024, 2025) and the International Conference on Machine Learning (ICML) (2025). These conferences conduct double-blind peer review through OpenReview. Each paper typically receives 3-4 reviews from reviewers assigned based on expertise and preferences. After initial reviews, authors submit rebuttals, followed by a discussion phase between authors and reviewers, where reviewers can observe co-reviewers’ assessments. Area chairs then write metareviews recommending decisions to the program committee. ICLR is unique because of its policy to keep public OpenReview pages for every paper after the conference, irrespective of its acceptance or rejection at the conference. The papers from NeurIPS and ICML are largely those that were accepted, as only a small fraction of researchers chose to keep the OpenReview pages for their rejected papers publicly accessible. For each available paper, we collect full paper and review text, the ratings assigned to the paper, and paper metadata (paper area and decision). This is an average of 6.5k papers and 25k reviews per conference (full data statistics are included in Appendix A.1).

*LLM Detection.* To detect LLM-generated text in papers and reviews, we use the statistical method from Liang et al. [9, 10], which models LLM-generated and human text as distinct distributions. The method assigns each vocabulary token a probability under each distribution based on its frequency in representative training corpora. It then estimates the fraction  $\alpha$  of LLM-generated text via maximum likelihood estimation: the method posits any text to have been drawn from a weighted distribution of LLM text and Human text, and returns  $\alpha$ , which is the mixture weight of the LLM text distribution in the solution to the MLE problem (see Appendix A.3 for details). In this way,  $\alpha$  depends on the bag-of-words representation of each paper or review. We set  $\alpha_{\text{threshold}} = 0.15$ , meaning that any text with  $\alpha > \alpha_{\text{threshold}}$  is classified as LLM-aided and human otherwise. This method offers key advantages: ease of use, fast inference and robustness to text that is merely "LLM aided" as compared to fully LLM generated. The reference LLM distribution we use is based on GPT-3.5 and GPT-4o outputs, which were the predominant models during our study period. This approach may underestimate usage of other contemporaneous models (e.g., Claude, Gemini), though GPT-4o’s dominant market share at the time suggests this is a minor limitation. In order to add robustness to our results, we repeat our analysis using Pangram 3.0, another state of the art LLM detector by Thai et al. [16], and provide these results in Appendix A.11.

## 4 Results

We now turn to our core empirical questions, analyzing interaction effects between LLM use in papers and reviews. Our analysis proceeds in four stages. First, we examine whether LLM reviews exhibit differential kindness to LLM papers in the full ICLR dataset. Then, we investigate how this varies by paper quality across venues. Next, we disentangle the behavior of fully LLM-generated reviews from LLM-assisted reviews found in the wild. Finally, we examine whether LLM assistance in metareviews affects their propensity to render accept decisions.

### 4.1 Are LLM-aided reviews kinder to LLM-aided papers?

While the summary statistics in Table 1 suggest the existence of a differential kindness exhibited by LLM-aided reviews to LLM-aided papers, these statistics cannot paint the whole picture by themselves. Comparing average ratings across quadrants implicitly assumes papers in each quadrant are comparable. However, several sources of bias may undermine this assumption: certain research areas may have both higher LLM adoption rates and

different baseline rating distributions, while less experienced authors may be more likely to use LLMs while also producing lower quality work for reasons unrelated to LLM use. This is a classic case of confounding, where aggregate comparisons can be misleading.

*4.1.1 Regression with area correction.* To address these concerns, we formally model this problem with the causal model in Figure 2. Solid arrows represent hypothesized causal effects; dashed arrows indicate unobserved confounding paths. For each paper-review pair  $(i, j)$ , we define the treatment  $T_{ij} \in \{0, 1\}$  as whether review slot  $j$  for paper  $i$  was assigned to a reviewer that used LLM assistance,  $X_i$  as paper covariates including  $X_i^{\text{LLM}} \in \{0, 1\}$  (whether paper  $i$  contains LLM-assisted text) and  $X_i^{\text{area}} \in \{0, 1\}^K$  (indicator variables for which of  $K$  subject areas paper  $i$  belongs to), and the outcome  $Y_{ij}$  as the overall rating (or auxiliary scores like presentation, contribution, or soundness).

Let  $Y_{ij}(1)$  and  $Y_{ij}(0)$  denote the potential outcomes under treatment and control, respectively. Our key identifying assumption is **conditional ignorability**:

$$T_{ij} \perp \{Y_{ij}(0), Y_{ij}(1)\} \mid X_i \quad (1)$$

That is, conditioned on available covariates (in our case, the paper’s area and the paper’s own LLM-assistance status), treatment assignment is independent of potential outcomes. We condition on area because certain subject areas may have both higher rates of LLM-assisted reviewing and systematically different rating distributions, acting as a potential confounder. We are specifically interested in whether the treatment effect of LLM-assisted reviewing differs depending on whether the paper itself was written with LLM assistance. Each paper has multiple review slots, which we treat as interchangeable, and receives some number of unique reviews. As a consequence of OpenReview overwriting original ratings and making them inaccessible after the rebuttal stage, we can only observe the final ratings provided by each review. This raises concerns about violations of the Stable Unit Treatment Value Assumption (SUTVA), which requires that (1) a unit’s potential outcomes depend only on its own treatment assignment, not on the treatment of other units (no interference), and (2) there are no hidden variations in treatment (Rubin [15]). However, several factors suggest these violations are unlikely to create stronger (hence spurious) effects than would exist in pre-rebuttal scores. First, most reviews (75-81%) remain unchanged (Kargaran et al. [5]), meaning our data predominantly reflects initial, independent assessments. Second, co-reviewer influence tends to push reviews toward consensus (Kargaran et al. [5]), which would dampen rather than amplify any differential treatment patterns. While access to pre-rebuttal scores would enable stronger causal claims, these considerations suggest post-rebuttal interference is unlikely to generate the interaction effects we document. We still provide analysis based on available pre-rebuttal scores for ICLR 2024-2025 collected by Kargaran et al. [5] in Appendix A.9 and observe that our takeaways remain unchanged.

We now frame the problem as a regression, which allows us to isolate the effects of interest while controlling for potential confounders. We use the conventional way of including interactions in a linear model by modeling the rating  $Y_{ij}$  for paper  $i$  and review  $j$  as

$$Y_{ij} = \beta_0 + \beta_1 X_i^{\text{LLM}} + \beta_2 T_{ij} + \beta_3 (X_i^{\text{LLM}} \times T_{ij}) + \gamma' Z_i + \epsilon_{ij}, \quad (2)$$

where  $Z_i$  is a vector of confounders we control for by including them in our regression analysis (for our regression, this is the group of indicator variables  $X_i^{\text{area}}$  that take a value of 1 if the paper belongs to a particular area and 0 otherwise) and  $\epsilon_{ij}$  is the noise term. This framing is useful because it naturally lends significant meanings to the regression coefficients:

- $\beta_0$ : The expected rating for a human-written paper receiving a human-written review (the baseline).
- $\beta_1$ : The difference in expected rating for LLM-assisted vs. human-written papers, holding review type constant.

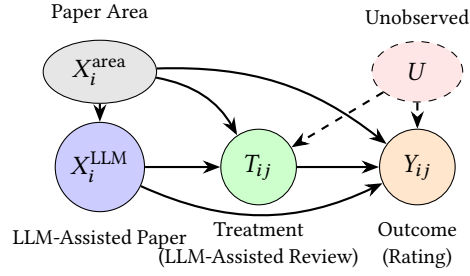


Fig. 2. Causal diagram for LLM use in peer review.  $T_{ij}$ : treatment (LLM-assisted review).  $Y_{ij}$ : outcome (rating).  $X_i^{\text{LLM}}$ : paper LLM use.  $X_i^{\text{area}}$ : paper area (confounder). Solid arrows are hypothesized causal effects; dashed arrows are potential unobserved confounding.

- $\beta_2$ : This is the Conditional Average Treatment Effect (CATE) of receiving an LLM review instead of a human review, for papers written without LLM assistance:  $\text{CATE}(X_i^{\text{LLM}} = 0) = \beta_2$ .
- Similarly, the CATE for LLM-assisted papers is given by:  $\text{CATE}(X_i^{\text{LLM}} = 1) = \beta_2 + \beta_3$ . This is the additional score an LLM-assisted paper receives if it is reviewed with LLM-assistance instead of entirely by a human.

The interaction coefficient  $\beta_3$  is our primary quantity of interest. It captures the *relative* leniency of LLM reviews toward LLM-assisted papers, compared to their leniency toward human-written papers. More precisely, it answers the question: how much more (or less) does switching from a human reviewer to a reviewer using LLM assistance benefit an LLM-assisted paper than it benefits a human-written paper? This connects directly to our earlier summary analysis, where we observed that LLM reviews tended to assign higher ratings than human reviews to *both* LLM-assisted and human-written papers on average, but the magnitude of this gap was notably larger for LLM-assisted papers (0.63 points) than for human-written papers (0.25 points). The interaction coefficient  $\beta_3$  formalizes this comparison:

$$\beta_3 = \text{CATE}(X_i^{\text{LLM}} = 1) - \text{CATE}(X_i^{\text{LLM}} = 0) \quad (3)$$

A positive  $\beta_3$  would imply a positive differential kindness shown by LLM-aided reviews to LLM-aided papers as compared to human reviewers.

The results of our regression analysis on ICLR Papers and Reviews from 2024 and 2025 are detailed in Table 2.

Quantity	Interpretation	Estimate	95% CI	p-value
$\beta_0$	Baseline: Human paper, Human review	5.145	[5.032, 5.259]	<0.001
$\beta_2$	CATE: LLM-aided review effect on Human papers	+0.253	[0.223, 0.283]	<0.001
$\beta_2 + \beta_3$	CATE: LLM-aided review effect on LLM papers	+0.629	[0.476, 0.782]	<0.001
$\beta_3$	Differential: Extra kindness to LLM-aided papers	+0.376	[0.221, 0.532]	<0.001

Table 2. Regression Analysis for determining LLM Review Kindness using the model from Eq. (2) on ICLR 2024+ICLR 2025 combined.

These results imply that LLM-aided reviewers are kinder than human reviewers to both human papers and LLM-aided papers, but their relative kindness, determined by  $\beta_3$ , is higher for LLM-aided papers than for human papers. This is very close to the observation we made in Table 1, where the differential kindness of LLM-aided reviews to LLM-assisted papers was  $0.63 - 0.25 = 0.38$ . Although regression controls for observable confounders, unobserved paper-level characteristics (e.g., topic trendiness, writing clarity) may still bias our estimates if they jointly influence both the likelihood of receiving an LLM review and the rating itself. Adding additional terms to the regression formulation would eventually run into the curse of dimensionality, particularly if interactions

between those terms matter. We also run into the risk of multicollinearity issues or accidentally conditioning on colliders and creating spurious associations. To address this limitation, we turn to a more stringent identification strategy that eliminates between-paper confounding without the typical drawbacks associated with potentially overspecified regressions.

**4.1.2 Within Paper Paired Analysis.** The regression framework, despite controlling for area, still compares ratings across different papers. We now leverage a unique feature of peer review: papers receive multiple reviews, allowing us to observe both potential outcomes for the same paper. By restricting to papers that received at least one human review and at least one LLM-assisted review, we can compute the actual treatment effect for each paper without relying on conditional independence assumptions. While restricting our corpus to only papers that received both types of reviews is a limitation that can introduce a selection bias, it provides an alternative method for estimating the CATE that can corroborate the previous analysis. The restricted corpus consists of 10945 papers.

For each paper  $i$ , we compute the within-paper difference:

$$\Delta_i = \bar{Y}_i^{\text{LLM}} - \bar{Y}_i^{\text{Human}} \quad (4)$$

where  $\bar{Y}_i^{\text{LLM}}$  and  $\bar{Y}_i^{\text{Human}}$  are the average ratings from LLM-assisted and human reviews, respectively.  $\Delta_i$  is the paper-specific treatment effect, holding all paper characteristics constant. We partition papers by LLM status and estimate:

$$\text{CATE}(X_i^{\text{LLM}} = 0) = \mathbb{E}[\Delta_i \mid \text{paper is human}], \quad \text{CATE}(X_i^{\text{LLM}} = 1) = \mathbb{E}[\Delta_i \mid \text{paper is LLM-assisted}] \quad (5)$$

If LLM reviews exhibit differential leniency toward LLM-assisted papers, then  $\text{CATE}(X_i^{\text{LLM}} = 1) > \text{CATE}(X_i^{\text{LLM}} = 0)$ . Beyond comparing means, we test whether the entire distribution of treatment effects differs between paper types using the Mann-Whitney U test:

$$H_0 : \text{The distributions of } \Delta_i \text{ for human and LLM-assisted papers are identical} \quad (6)$$

$$H_1 : \Delta_i \text{ for LLM-assisted papers is different from } \Delta_i \text{ for human papers} \quad (7)$$

Table 3 shows the results. The mean treatment effect is 0.340 for LLM-assisted papers versus 0.244 for human papers (difference: +0.096), though the p-value of 0.058 provides only weak evidence against the null hypothesis. Comparing this number to  $\beta_3$  (0.376) from Table 2 and the estimate from Table 1 (0.38), we find the effect size to be much smaller upon controlling for paper-level confounding factors.

Metric	CATE (LLM-aided papers)	CATE (Human papers)	Difference	p-value
Rating	0.340	0.244	+0.096	0.058

Table 3. Within-paper treatment effects (ICLR 2024-2025). Mann-Whitney U test (one-tailed) for distributional difference.

**4.1.3 How influential are LLM-aided reviews in final decisions?** Beyond examining rating patterns, we investigate whether LLM-aided reviews carry comparable weight to human reviews in determining final acceptance decisions. We measure review influence by alignment with the final decision: a review "aligns" if its rating matches the paper's outcome (rating  $> 5$  for accepted papers, rating  $\leq 5$  for rejected papers). We additionally stratify by paper quality to examine whether LLM-aided reviews show differential alignment for borderline versus clear-cut cases, as influence at decision boundaries is most consequential for outcomes.

Using the ICLR 2024 and 2025 combined corpus, we stratify papers by quality using a hold-out human review as the quality indicator, then compute alignment rates for LLM-aided and human reviews (excluding the quality-indicator review) across six quality buckets that correspond to the six ratings reviewers are able to choose from. Table 4 shows the results. Human reviews consistently show higher alignment rates than LLM-aided reviews

across most quality levels, with alignment scores between 75% – 87% compared to 69% – 81% for LLM-aided reviews.

Paper Quality	Human Align %	LLM Align %
1	87.4	80.7
3	80.8	73.5
5	75.0	69.5
6	76.0	74.9
8	78.8	79.2
10	79.9	74.6
Overall	77.7	73.8

Table 4. Review alignment rates with final decisions by paper quality for ICLR 2024-2025. Alignment defined as: rating > 5 for accepted papers, rating  $\leq$  5 for rejected papers. Quality determined by hold-out human review.

These patterns suggest human reviews remain more influential in decision-making, even for borderline papers where the accept/reject decision is most consequential. The lower alignment of LLM-aided reviews in the medium-quality range (ratings 3-6) can be an artifact of their tendency to inflate ratings for lower quality papers relative to human reviews, which leads to misalignment in rejection decisions. In fact, 91.6% of LLM-aided review misalignments in quality bucket 5 are due to recommended acceptances for ultimately rejected papers.

At this stage, the results indicate evidence for LLM-aided reviews favoring LLM-aided papers, though their lower alignment rates seem to suggest limited influence on final decisions. We now turn to accepted papers to see if this pattern persists, since these papers form the most important paper subset in the corpus.

#### 4.2 Does the kindness of LLM-aided reviews to LLM-aided papers vary with paper quality?

Interestingly, when we restrict the ICLR corpus to only accepted papers, we observe markedly different patterns (Table 5). The quadrants now show similar scores with no apparent differential LLM kindness. We observe similar patterns for NeurIPS and ICML (Appendix A.2), where by default rejected papers are private on OpenReview (though some authors choose to make their papers public).

*4.2.1 Regression with area correction.* In order to investigate this, we run the same regression analysis from Section 4.1.1 on the accepted only corpus of ICLR 2024-2025, in Table 6. The previously observed differential kindness ( $\beta_3$ ) is no longer significant ( $p > 0.30$ ), and individual CATEs are near zero. Similar null results appear for NeurIPS and ICML (Appendix A.6.2, A.6.3).

To understand this discrepancy between results on the accepted-only corpus and the all papers corpus, we examine whether LLM-aided reviews differ from human reviews across the quality spectrum. We stratify papers into four quality buckets (rating  $\leq 3$ ,  $= 5$ ,  $= 6$ ,  $\geq 8$ ) using a randomly selected held-out human review as the quality

Quadrant (Paper, Review)	Rating	Confidence	Presentation	Contribution	Soundness
(Human, Human)	6.467 $\pm$ 1.319	3.573 $\pm$ 0.792	2.902 $\pm$ 0.681	2.720 $\pm$ 0.632	2.918 $\pm$ 0.605
(Human, LLM)	6.409 $\pm$ 1.226	3.535 $\pm$ 0.811	2.932 $\pm$ 0.604	2.756 $\pm$ 0.621	2.924 $\pm$ 0.555
(LLM, Human)	6.482 $\pm$ 1.309	3.671 $\pm$ 0.752	2.879 $\pm$ 0.752	2.725 $\pm$ 0.610	2.839 $\pm$ 0.598
(LLM, LLM)	6.552 $\pm$ 1.242	3.606 $\pm$ 0.853	2.945 $\pm$ 0.566	2.836 $\pm$ 0.656	2.903 $\pm$ 0.597

Table 5. Summary Statistics for (Paper, Review) pairs from ICLR 2024 and 2025 combined (accepted papers only). (Paper, Review) pairs are divided into 4 quadrants based on the LLM status of the paper and the review.  $\alpha$  value of 0.15 is used as a threshold: Papers and Reviews are classified as LLM-aided if they have  $\alpha > 0.15$  and Human otherwise. Cell values are presented as Mean  $\pm$  Std Dev

Quantity	Interpretation	Estimate	95% CI	p-value
$\beta_0$	Baseline: Human paper, Human review	6.544	[6.389, 6.699]	<0.001
$\beta_2$	CATE: LLM-aided review effect on Human papers	-0.051	[-0.091, -0.012]	0.011
$\beta_2 + \beta_3$	CATE: LLM-aided review effect on LLM-aided papers	+0.082	[-0.168, +0.332]	0.520
$\beta_3$	Differential: Extra kindness to LLM-aided papers	+0.134	[-0.120, +0.387]	0.301

Table 6. Regression Analysis for ICLR accepted papers only (2024+2025 Combined). Model: Rating =  $\beta_0 + \beta_1 X_i^{LLM} + \beta_2 T_{ij} + \beta_3 (X_i^{LLM} \times T_{ij}) + \gamma' Z_i$ . Area fixed effects included. The interaction effect  $\beta_3$  is not statistically significant, in contrast to the full dataset (Table 2).

indicator, then compute average ratings from LLM-assisted and human reviews (excluding the quality-indicator review) within each bucket. Figure 3 reveals two key patterns. First, LLM-aided reviews provide substantially higher ratings than human reviews for low-quality papers, but this gap narrows and disappears as paper quality increases. For papers with quality indicator rating  $\leq 3$ , LLM-aided reviews average approximately 0.4-0.5 points higher than human reviews, shrinking to near-zero for medium and high-quality papers. Second, LLM-assisted papers are heavily overrepresented in the low-quality bucket: 173 of 2,768 papers (6.2%) in the rating  $\leq 3$  bucket are LLM-assisted, compared to only 29 of 1,316 papers (2.2%) in the rating  $\geq 8$  bucket—mirroring the rating distribution patterns in Figure 1b. The combination of these two factors could potentially be creating the appearance of differential kindness in aggregate analyses.




Fig. 3. Average ratings from LLM-aided reviews vs. human reviews across quality buckets. Shaded area is 1.96\*std-error.

Fig. 4. Regression coefficients by quality bucket. Each point represents a separate regression run on papers within that quality bucket. Shaded area is 95% CI. The dashed red line shows the estimate from the full regression for comparison. Zero is indicated with a black dotted line.

We probe this quality-dependent relationship further by running separate regressions within each quality bucket using all remaining reviews (excluding the quality-indicator review). Figure 4 plots the regression coefficients for each bucket. The leftmost panel shows that  $\beta_1$  (the effect of a paper being LLM-assisted on ratings from human reviewers) is strongly negative for low-quality papers ( $\approx -1.0$  for rating  $\leq 3$ ) but approaches zero for high-quality papers (rating  $\geq 8$ )—human reviewers penalize low-quality LLM-aided papers but show no bias for high-quality LLM-aided papers. The center panel shows the interaction effect  $\beta_3$  follows an inverted-U pattern: it is small for the lowest-quality papers ( $\approx 0.1$ ), peaks around medium quality (rating 5-6, where  $\beta_3 \approx 0.5$ ), and decreases again for high-quality papers ( $\approx 0.2$ ). Notably,  $\beta_3$  remains positive across all buckets but only achieves statistical significance in the middle-quality range, where the regression estimate on the full population (dashed red line) aligns with the bucket-specific estimates. This suggest LLM-aided reviews show differential kindness to LLM-aided papers in the intermediate quality range, but not otherwise. The rightmost panel shows  $\beta_1 + \beta_3$  (the effect of a paper being LLM-aided on ratings from LLM-aided reviewers) is strongly negative for

low-quality papers but becomes slightly positive for medium-quality papers before returning to near-zero for high-quality papers - LLM-aided reviewers penalize low-quality LLM-aided papers (but to a lesser extent than human reviewers) and only show a very small bias for high-quality LLM-aided papers.

These patterns reconcile our earlier findings: the positive  $\beta_3$  estimated by regression over the complete paper corpus reflects the middle-quality range where the effect is strongest. Restricting to the accepted-papers corpus is similar to restricting to the high-quality range where  $\beta_3$  approaches zero.

**4.2.2 Within Paper Paired Analysis.** To further validate these findings, we repeat the within-paper paired analysis stratified by quality bucket. For each bucket, we compute treatment effects  $\Delta_i$  using only the non-quality-indicator reviews and test whether  $\Delta_i$  differs between LLM and human papers. Table 7 shows the results. Interestingly, none of the quality buckets show statistically significant differences in treatment effects between LLM and human papers, though point estimates vary. We observe that the effect magnitudes are also slightly smaller than for the entire paper corpus (0.096 points), and even change sign in the ‘Med-High’ bucket, except for papers in the ‘High Quality’ bucket, which have a low frequency of LLM-aided papers.

Quality Bucket	N <sub>LLM</sub>	N <sub>Human</sub>	CATE Diff.	p-value	Sig.
Low (rating $\leq 3$ )	173	2,595	+0.038	0.332	no
Medium (rating = 5)	66	2,723	+0.082	0.292	no
Med-High (rating = 6)	65	2,831	-0.090	0.524	no
High (rating $\geq 8$ )	29	1,287	+0.369	0.137	no

Table 7. Within-paper paired analysis by quality bucket. CATE Diff. =  $\text{mean}(\Delta_i \text{ for LLM papers}) - \text{mean}(\Delta_i \text{ for human papers})$ . N<sub>Human</sub> and N<sub>LLM</sub> represent the number of Human and LLM-aided papers in the bucket. Mann-Whitney U test (one-tailed).

The lack of significance within buckets, despite regression showing varying  $\beta_3$ , suggests aggregate patterns reflect between-paper heterogeneity rather than systematic differential treatment. Combined with LLM-aided reviews’ uniform leniency toward low-quality papers (Figure 3), this supports the interpretation that apparent LLM-LLM interaction is driven by LLM-aided papers’ overrepresentation in low-quality ranges where LLM-aided reviews are generally lenient, not by specific favoritism toward LLM-aided content.

### 4.3 How do fully LLM-generated reviews differ from LLM-aided reviews?

Up to this point, our analysis has relied on observational data from real papers and reviews, capturing authentic peer review dynamics. However, this conflates two distinct agents: the LLM and the human using it. When we classify a review as *LLM-assisted*, we identify reviews with LLM-aided text, but the final rating reflects a human reviewer who may accept, modify, or ignore any LLM suggestions. Our observational analysis captures *humans who use LLMs*, not LLMs themselves. This is an important distinction since identical prompts can produce markedly different responses across sessions and user contexts Wang et al. [17]. This raises the following question: if we see some form of bias in LLM-assisted reviews, is this bias due to the LLM or due to the human using the LLM?

To disentangle these effects, we augment our analysis with experiments on fully LLM-generated reviews. In particular, we prompt GPT-4o to generate reviews for papers in our corpus, obtaining ratings reflecting the LLM’s judgment directly, without human mediation. We avoid GPT-5 because of the strong possibility that it has been trained on the original reviews already. This isolates the *pure LLM* effect, allowing comparison to both human reviewers and *humans using LLMs*. Note that the fully LLM-generated reviews are not subject to confounding or selection biases, since they are synthetically generated.

We generate reviews for randomly sampled papers from the 2025 editions of ICLR, ICML, and NeurIPS. Prompts follow the template used by Lu et al. [12] with conference-specific guidelines and no additional instructions to be lenient or harsh (details in Appendix A.12). We sample 60 LLM-aided papers and 480 human papers from the

ICLR 2025 corpus. Furthermore, we sample 75 LLM-aided papers and 300 human papers from the ICLR 2025 accepted corpus, the NeurIPS 2025 corpus, and the ICML 2025 corpus. Note that this oversamples LLM-aided papers to ensure reliable statistical analysis. These numbers are restricted by the number of LLM-aided papers available in the corpus, since ICML 2025 only has 75 papers classified as LLM-aided. We then perform the same within-paper paired analysis on these generated reviews against the human reviews for the same papers. Table 8 has these results for the "Overall Rating" field while additional results are included in Appendix A.8. These

Conference	Range	CATE (LLM Papers)	CATE (Human Papers)	Difference	p-value
ICLR 2025	1-10	2.248	1.623	+0.625	0.0018
ICLR 2025 (Accepted)	1-10	0.674	0.712	-0.038	0.5670
NeurIPS 2025	1-6	0.408	0.383	+0.025	0.3616
ICML 2025	1-5	0.774	0.759	+0.015	0.3315

Table 8. Result of Mann-Whitney U test (one-tailed) on 2025 conference data with fully generated reviews. Kindness difference = Generated Review rating - Human rating.

results are interesting in two respects. First, we notice that the magnitude of difference in mean kindness for LLMs over Humans on LLM Papers as compared to Human Papers is actually a lot larger than what we saw with observational data in 3 where it was only about 0.1. This could be explained by a few hypotheses: for example, people using LLMs might either lower the rating the LLM provides, prompt the LLM to directly provide a more critical review, or merely use the LLM to expand upon their own critical review. Second, the generated reviews provide substantially higher ratings than human reviews for both human and LLM-aided papers but again show no differential treatment when we restrict the corpus to only accepted papers.

We next examine the distribution of scores assigned by these fully LLM-generated reviews (Figure 5). We observe that fully LLM-generated reviews assign ratings almost exclusively between 6 and 7 regardless of paper quality, exhibiting minimal slope across quality buckets. This compression has serious implications: fully LLM-generated reviews fail to discriminate paper quality, systematically inflating ratings for weak papers while providing little signal about paper merit. The gap between LLM and human ratings is largest for low-quality papers, where automated reviews would recommend acceptance for papers humans would reject. These observations suggest that fully automated LLM reviewing would undermine the accuracy of peer review. Human-in-the-loop moderation (as seen in observational LLM-assisted reviews) partially mitigates this compression: the line corresponding to LLM-aided reviews in Figure 3 has a lower slope and higher intercept than the line corresponding to human reviews, but still shows some discriminative power.

The results of the significance test indicate that fully LLM-generated reviews are not especially kinder to LLM-aided papers than they are to Human papers. Instead, they are simply less harsh to papers in general, and the over-representation of LLM aided papers in low quality buckets only misleads us into thinking this differential kindness exists in the complete ICLR corpus.

#### 4.4 How does LLM-assistance in metareviewing relate to a paper’s final acceptance decision?

Metareviews are written by Area Chairs based on the reviews to render a final acceptance decision. While metareviewers can add their own analysis and arguments, part of their role is summarizing the arguments provided by the reviewers. Summarization is a domain where LLMs typically excel, making metareviews a tempting target for LLM assistance.

Applying the detection method from Liang et al. [9] to ICLR metareviews from 2023-2025, we observe increasing LLM use (Figure 6), with ICLR 2025 showing evidence of significant LLM involvement in up to 49% of metareviews. This indicates a growing rate of LLM adoption even among senior researchers in the field.

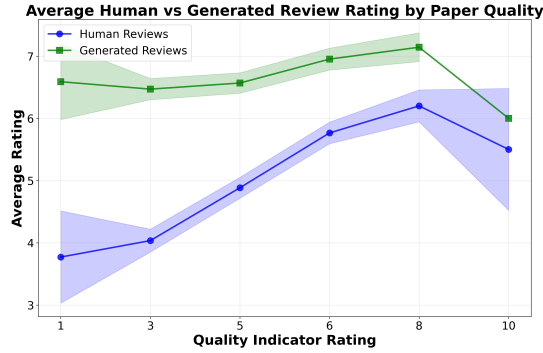


Fig. 5. Average Rating provided by Human reviews and fully LLM reviews to randomly sampled ICLR 2025 papers, bucketed by quality. The datapoints correspond to quality bucket 10 are noisy because only 1 such paper was sampled.




Fig. 6. Occurrence of LLM-aided metareviews in ICLR. An LLM-aided metareview has  $\alpha > 0.15$ . We see an increasing trend with over 49% of metareviews in 2025 satisfying this criterion.

Based on the role of a metareview in the peer review process, a key question is whether an LLM-assisted metareview is more or less likely to render an accept decision than a human metareview observing the same set of scores.

4.4.1 *Regression Analysis.* We frame this as a logistic regression problem for every individual paper:

$$\text{logit}(P_{\text{accepted}}) = \beta_1(\text{mean rating}) + \beta_2(\text{std rating}) + \beta_3 I_{\text{metareview}} + \gamma'Z, \quad (8)$$

where  $P_{\text{accepted}}$  is the probability of paper acceptance, mean rating captures the central tendency of reviewer scores, std rating measures the disagreement among reviewers,  $I_{\text{metareview}}$  is a binary indicator for LLM-aided metareviews ( $\alpha > 0.15$ ), and  $Z$  represents area fixed effects. Our coefficient of interest,  $\beta_3$ , measures whether LLM-aided metareviews associate with different acceptance probabilities after controlling for both rating level, agreement, and area.

Table 9 presents the results for ICLR 2024 and 2025 combined. As expected, higher mean ratings strongly predict acceptance ( $\beta_1 = +3.327$ ,  $p < 0.001$ , 95% CI: [3.323, 3.331]), while higher rating variance reduces acceptance likelihood ( $\beta_2 = -0.288$ ,  $p < 0.001$ , 95% CI: [-0.297, -0.280]). Importantly,  $\beta_3 = +0.214$  ( $p < 0.001$ , 95% CI: [+0.204, +0.224]), indicating that papers with LLM-aided metareviews are significantly more likely to be accepted, after controlling for both the mean and variance of ratings. This suggests that LLM assistance in metareview writing is associated with more favorable acceptance decisions beyond what would be expected from the underlying review scores alone.

Parameter	Interpretation	Coefficient	95% CI	p-value
$\beta_1$	Contribution of mean rating	+3.327	[3.323, 3.331]	< 0.001
$\beta_2$	Contribution of std deviation of ratings	-0.288	[-0.297, -0.280]	< 0.001
$\beta_3$	Contribution of metareview being LLM-assisted	+0.214	[+0.204, +0.224]	< 0.001

Table 9. Logistic regression results for ICLR 2024-2025. Model: Eq. 8. Target: Accept (1) vs Reject (0). Area fixed effects included. Higher mean ratings increase acceptance probability, higher rating variance decreases it, and LLM-aided metareviews are associated with increased acceptance likelihood controlling for both rating level and consensus.

Actual Metareview Type	N	Alignment %	Synthetic Accept Rate	Actual Accept Rate
Human ( $\alpha \leq 0.15$ )	108	76.9	17.6%	40.7%
LLM-aided ( $\alpha > 0.15$ )	122	68.9	20.5%	51.6%
<b>Overall</b>	230	72.6	19.1%	46.5%

Table 10. Alignment between fully LLM-generated metareviews and actual decisions (ICLR 2025 sample). Synthetic metareviews are substantially harsher than actual decisions, recommending acceptance at less than half the rate.

**4.4.2 Fully LLM-generated Metareviews.** To again disentangle whether the decisions that humans make with LLM-assistance align with the decisions that LLMs make on their own, we generated fully LLM-generated metareviews for 300 randomly sampled ICLR 2025 papers. 70 of these papers do not have metareviews because they were withdrawn from the conference after initial reviews, which raises the acceptance rate in the remaining corpus of 230 papers to 46.5%. We prompt GPT-4o with all reviews a paper received and the conference rating and confidence scales (see Appendix A.12 for full prompts). We then extract accept/reject decisions from the generated metareviews and compare these fully LLM-generated decisions to actual decisions, stratified by whether the actual metareview was LLM-aided ( $\alpha > 0.15$ ).

Table 10 shows the results. Overall, fully LLM-generated metareviews align with actual decisions 72.6% of the time. We actually see that fully LLM-generated metareviews are substantially harsher than actual decisions: they recommend acceptance for only 44 of 230 papers (19.1%), while actual decisions accepted 107 papers (46.5%). Alignment rates differ by actual metareview type: synthetic decisions align with human metareviews 76.9% of the time versus 68.9% for LLM-aided metareviews (which tended to be more positive).

For both human and LLM-aided metareviews, fully LLM-generated decisions achieve 100% precision (when they recommend accept, the paper was actually accepted) but low recall (43.2% for human metareviews, 39.7% for LLM-aided). Fully LLM-generated metareviews systematically under-predict acceptances: they correctly identify strong accepts but fail to recommend acceptance for many borderline papers that human metareviewers ultimately accepted. While we can only speculate about how metareviewers use the LLM as part of their decision process, they do not appear to be following the LLM acceptance recommendations without further reflection.

The harshness of fully LLM metareviews also reveals how LLMs interpret reviewing instructions. Despite receiving the same reviews and guidelines that led human metareviewers to accept 40.7% of papers, the LLM recommends acceptance for only 17.6%. This suggests LLMs may have stricter interpretations of decision thresholds, or a lack of context typically associated with scores, for instance, a rating of '8' is only termed an "Accept" but is generally considered a very good paper. This would also explain their tendency to give out higher ratings as reviewers when recommending acceptances, even when following identical instructions.

## 5 Key Takeaways

We return to the research questions introduced in Section 1 with empirically-grounded answers:

*Are LLM-aided reviews kinder to LLM-aided papers on average and does this vary by paper quality?* Our threshold-varying analysis (Figure 4) demonstrates that the interaction effect is strongly modulated by paper quality. For papers in the top quality buckets, we observe minimal differential treatment ( $\beta_3 \approx 0$ ). The positive interaction coefficient observed in the full corpus is driven primarily by papers in the lower quality buckets, where LLM reviews exhibit greater leniency than human reviews, a leniency that benefits LLM-assisted papers more simply because they are overrepresented in this range. As shown in Figure 3, LLM-assisted reviews provide systematically higher ratings to lower quality papers regardless of whether those papers used LLM assistance, but this compression of the rating scale disproportionately affects the (LLM Paper, LLM Review) quadrant due to its compositional skew toward weaker submissions. The within-paper paired analyses with quality buckets (Table 7) confirms this observation.

*How does an LLM-aided review differ from a fully LLM-generated review? What is the effect of the human in the loop?* Our synthetic experiments (Table 8) reveal substantial differences between unmediated LLM reviews and human-mediated LLM-assisted reviews. Fully LLM-generated reviews exhibit markedly compressed rating distributions, assigning scores almost exclusively in the 6–7 range (on a 1–10 scale) regardless of paper quality (Figure 5). LLM-aided reviews have a broader and more discriminating rating distribution that better tracks paper quality. The kindness differential between LLM and human reviews for LLM-assisted papers is approximately 0.625 points in the fully synthetic condition (ICLR 2025 full corpus) but only 0.096 points in the observational data (Table 3), a reduction of roughly 85%.

*How does LLM use in a metareview relate to a paper’s final decision?* The regression analyses in Section 4.4 reveals that given the same set of scores, an LLM-aided metareviewer is more likely to render an accept decision than a human reviewer. However, synthetic experiments show that fully LLM metareviews are actually significantly harsher than their human and LLM-aided counterparts.

## 6 Discussion and Conclusions

This paper provides empirical insights into how LLM use on both sides of peer review affects ratings and paper outcomes. We find asymmetric patterns: while LLM-aided papers receive lower ratings overall, LLM-aided reviews show stronger positive effects when evaluating LLM-aided papers. Critically, these patterns largely disappear when restricting analysis to accepted papers, indicating effects are driven by lower quality submissions where LLM-aided reviews exhibit general leniency. The degree of rating compression, both in LLM-aided and fully LLM-generated reviews (the latter corroborating Zhu et al. [20]’s findings) is particularly concerning since these reviews ultimately provide less information and introduce noise into the decision process. Beyond review ratings, we find that LLM-aided metareviews are associated with higher acceptance rates despite fully LLM-generated metareviews being overly conservative, and that LLM-aided reviews align less well with final decisions than human reviews, even for borderline papers where decision-making is most consequential. We also make observations about how LLM-aided reviews differ from their fully LLM-generated counterparts, noting the latter’s lack of discriminative power.

Periodic reviews of this kind are essential as this landscape continues to evolve. Interesting future directions include comparing our results to those from submissions to similarly reputed journals, which are often less susceptible to some of the issues conferences face: reviewers are not authors competing for the same limited spots, and review timelines are less stressed. We also aim to examine how this relationship varies across subjects, not limited to computer science. However, conferences in other domains (and even journals within CS) do not provide numerical scores with their reviews, which is a major component of our present analysis.

Our results on LLM-aided and fully LLM-generated reviews/metareviews together suggest potential patterns in the way people use LLMs in peer review. For instance, metareviewers might selectively deploy LLMs when writing accept decisions, where the summarization task is more straightforward, but write rejection metareviews themselves, where detailed critical reasoning may be required. Alternatively, LLM use could be more popular among metareviewers who are more lenient. Similarly, the gap between LLM-aided reviews and fully LLM-generated reviews could be a result of reviewers asking LLMs to be more critical in their prompts, or they could simply be using LLMs only to expand upon their own salient reviews. We encourage full-scale user studies to test any such hypotheses.

Multiple major ML conferences have now instituted LLM policies beyond simple disclosure of use. On the most permissive end, the Association for the Advancement of Artificial Intelligence started a pilot program (AAAI Policy [1]) to provide supplementary LLM-generated reviews to authors alongside human reviews, and provide LLM generated summaries of author-reviewer discussions to aid metareviewers in their decision. As mentioned in Section 1, ICML has instituted a new policy for their 2026 edition requiring both reviewers and authors to declare

either conservative or permissive LLM use, with reviewer-author pairing done based on these declarations. We see this as problematic, since our results indicate that this could lead to a larger number of low-quality LLM-aided papers being accepted, because of the general leniency of LLM-assisted and fully LLM generated reviews. The academic community needs to direct their focus to identifying ways to use these tools in a way that is beneficial to the peer review process, and to counteract use that is detrimental. We hope the findings of this paper provide valuable insights to conference organizers as they design new policies and tools.

## 7 Generative AI usage statement

Generative AI tools have been used to assist in formatting  $\LaTeX$  (creating tables and improving alignment).

## References

- [1] AAAI Policy. 2025. AAAI Launches AI-Powered Peer Review Assessment System. <https://aaai.org/aaai-launches-ai-powered-peer-review-assessment-system/>.
- [2] ICML Policy. 2025. ICML 2026 LLM Policy. <https://icml.cc/Conferences/2026/LLM-Policy>.
- [3] International Conference on Learning Representations. 2025. ICLR 2025 Reviewer Guide. <https://iclr.cc/Conferences/2025/ReviewerGuide>.
- [4] International Conference on Machine Learning. 2025. ICML 2025 Reviewer Instructions. <https://icml.cc/Conferences/2025/ReviewerInstructions>.
- [5] Amir Hossein Kargaran, Nafiseh Nikeghbal, Jing Yang, and Nedjma Ousidhoum. 2025. Insights from the ICLR Peer Review and Rebuttal Process. arXiv:2511.15462 [cs.CY] <https://arxiv.org/abs/2511.15462>
- [6] Janis Keuper. 2025. Prompt Injection Attacks on LLM Generated Reviews of Scientific Publications. arXiv:2509.10248 [cs.LG] <https://arxiv.org/abs/2509.10248>
- [7] Burak Kocak, Mehmet Ruhi Onur, Seong Ho Park, Pascal Baltzer, and Matthias Dietzel. 2025. Ensuring peer review integrity in the era of large language models: A critical stocktaking of challenges, red flags, and recommendations. *European Journal of Radiology Artificial Intelligence* 2 (2025), 100018. doi:10.1016/j.ejrai.2025.100018
- [8] Walter Laurito, Benjamin Davis, Peli Grietzer, Tomáš Gavenčíak, Ada Böhm, and Jan Kulveit. 2025. AI–AI bias: Large language models favor communications generated by large language models. *Proceedings of the National Academy of Sciences* 122, 31 (2025), e2415697122. arXiv:<https://www.pnas.org/doi/pdf/10.1073/pnas.2415697122> doi:10.1073/pnas.2415697122
- [9] Weixin Liang, Zachary Izzo, Yaohui Zhang, Haley Lepp, Hancheng Cao, Xuandong Zhao, Lingjiao Chen, Haotian Ye, Sheng Liu, Zhi Huang, Daniel A. McFarland, and James Y. Zou. 2024. Monitoring AI-modified content at scale: a case study on the impact of ChatGPT on AI conference peer reviews. In *Proceedings of the 41st International Conference on Machine Learning (Vienna, Austria) (ICML '24)*. JMLR.org, Article 1192, 46 pages.
- [10] Weixin Liang, Yaohui Zhang, Zhengxuan Wu, Haley Lepp, Wenlong Ji, Xuandong Zhao, Hancheng Cao, Sheng Liu, Siyu He, Zhi Huang, Diyi Yang, Christopher Potts, Christopher D Manning, and James Y. Zou. 2024. Mapping the Increasing Use of LLMs in Scientific Papers. arXiv:2404.01268 [cs.CL] <https://arxiv.org/abs/2404.01268>
- [11] Zhehui Liao, Maria Antoniak, Inyoung Cheong, Evie Yu-Yen Cheng, Ai-Heng Lee, Kyle Lo, Joseph Chee Chang, and Amy X. Zhang. 2024. LLMs as Research Tools: A Large Scale Survey of Researchers' Usage and Perceptions. arXiv:2411.05025 [cs.CL] <https://arxiv.org/abs/2411.05025>
- [12] Chris Lu, Cong Lu, Robert Tjarko Lange, Jakob Foerster, Jeff Clune, and David Ha. 2024. The AI Scientist: Towards Fully Automated Open-Ended Scientific Discovery. arXiv:2408.06292 [cs.AI] <https://arxiv.org/abs/2408.06292>
- [13] Neural Information Processing Systems. 2025. NeurIPS 2025 Reviewer Guidelines. <https://neurips.cc/Conferences/2025/ReviewerGuidelines>.
- [14] Arjun Panickssery, Samuel R. Bowman, and Shi Feng. 2024. LLM evaluators recognize and favor their own generations. In *Proceedings of the 38th International Conference on Neural Information Processing Systems (Vancouver, BC, Canada) (NIPS '24)*. Curran Associates Inc., Red Hook, NY, USA, Article 2197, 31 pages.
- [15] Donald B. Rubin. 1980. Randomization Analysis of Experimental Data: The Fisher Randomization Test Comment. *J. Amer. Statist. Assoc.* 75, 371 (1980), 591–593. <http://www.jstor.org/stable/2287653>
- [16] Katherine Thai, Bradley Emi, Elyas Masrouf, and Mohit Iyyer. 2025. EditLens: Quantifying the Extent of AI Editing in Text. arXiv:2510.03154 [cs.CL] <https://arxiv.org/abs/2510.03154>
- [17] Angelina Wang, Daniel E. Ho, and Sanmi Koyejo. 2025. The Inadequacy of Offline LLM Evaluations: A Need to Account for Personalization in Model Behavior. arXiv:2509.19364 [cs.CL] <https://arxiv.org/abs/2509.19364>

- [18] Rui Ye, Xianghe Pang, Jingyi Chai, Jiaao Chen, Zhenfei Yin, Zhen Xiang, Xiaowen Dong, Jing Shao, and Siheng Chen. 2024. Are We There Yet? Revealing the Risks of Utilizing Large Language Models in Scholarly Peer Review. arXiv:2412.01708 [cs.CL] <https://arxiv.org/abs/2412.01708>
- [19] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. Judging LLM-as-a-Judge with MT-Bench and Chatbot Arena. arXiv:2306.05685 [cs.CL] <https://arxiv.org/abs/2306.05685>
- [20] Changjia Zhu, Junjie Xiong, Renkai Ma, Zhicong Lu, Yao Liu, and Lingyao Li. 2025. When Your Reviewer is an LLM: Biases, Divergence, and Prompt Injection Risks in Peer Review. arXiv:2509.09912 [cs.CY] <https://arxiv.org/abs/2509.09912>

## A Appendix

### A.1 Dataset

Table 11 summarizes paper and review counts across venues that we use for analyses. Note that the NeurIPS and ICML data is heavily biased towards accepted papers.

Conference	Total		Accepted		Rejected	
	Papers	Reviews	Papers	Reviews	Papers	Reviews
ICLR 2024	7,404	28,028	2,260	8,742	5,144	19,286
ICLR 2025	11,672	46,748	3,704	14,940	7,968	31,808
NeurIPS 2024	4,236	16,644	4,035	15,855	201	789
NeurIPS 2025	5,526	22,313	5,275	21,301	251	1,012
ICML 2025	3,422	13,102	3,257	12,478	165	624

Table 11. Dataset statistics: Submission and review counts across venues.

### A.2 Summary Statistics

Tables 12, 13, and 14 present summary statistics for NeurIPS 2024, NeurIPS 2025, and ICML 2025. This includes every field reviewers were required to provide a rating on.

Quadrant ( <i>Paper, Review</i> )	n	Rating	Confidence	Presentation	Contribution	Soundness
(Human, Human)	13,020	5.87 ± 1.20	3.60 ± 0.84	2.86 ± 0.71	2.72 ± 0.65	2.92 ± 0.63
(Human, LLM)	3,432	5.86 ± 1.14	3.62 ± 0.88	2.91 ± 0.61	2.78 ± 0.61	2.94 ± 0.56
(LLM, Human)	65	5.66 ± 1.33	3.85 ± 0.73	2.82 ± 0.61	2.52 ± 0.66	2.78 ± 0.62
(LLM, LLM)	58	5.97 ± 1.31	3.78 ± 0.92	2.84 ± 0.67	2.79 ± 0.61	2.95 ± 0.63

Table 12. Summary Statistics for (Paper, Review) pairs from NeurIPS 2024. (Paper, Review) pairs are divided into 4 quadrants based on the LLM status of the paper and the review.  $\alpha$  value of 0.15 is used as a threshold: Papers and Reviews are classified as LLM-aided if they have  $\alpha > 0.15$  and Human otherwise. Cell values are presented as Mean ± Std Dev

Quadrant ( <i>Paper, Review</i> )	n	Rating	Confidence	Quality	Clarity	Significance	Originality
(Human, Human)	18,135	4.31 ± 0.74	3.52 ± 0.82	2.87 ± 0.61	2.88 ± 0.70	2.72 ± 0.65	2.78 ± 0.64
(Human, LLM)	3,686	4.29 ± 0.68	3.54 ± 0.84	2.91 ± 0.56	2.92 ± 0.61	2.78 ± 0.63	2.84 ± 0.62
(LLM, Human)	264	4.20 ± 0.72	3.69 ± 0.80	2.74 ± 0.59	2.77 ± 0.70	2.61 ± 0.61	2.63 ± 0.63
(LLM, LLM)	111	4.34 ± 0.69	3.65 ± 0.87	2.86 ± 0.60	2.87 ± 0.63	2.83 ± 0.57	2.81 ± 0.60

Table 13. Summary Statistics for (Paper, Review) pairs from NeurIPS 2025. (Paper, Review) pairs are divided into 4 quadrants based on the LLM status of the paper and the review.  $\alpha$  value of 0.15 is used as a threshold: Papers and Reviews are classified as LLM-aided if they have  $\alpha > 0.15$  and Human otherwise. Cell values are presented as Mean ± Std Dev

Quadrant ( <i>Paper, Review</i> )	Rating
(Human, Human)	3.21 ± 0.80
(Human, LLM)	3.28 ± 0.76
(LLM, Human)	3.11 ± 0.84
(LLM, LLM)	3.25 ± 0.75

Table 14. Summary Statistics for (Paper, Review) pairs from ICML 2025. (Paper, Review) pairs are divided into 4 quadrants based on the LLM status of the paper and the review.  $\alpha$  value of 0.15 is used as a threshold: Papers and Reviews are classified as LLM-aided if they have  $\alpha > 0.15$  and Human otherwise. Cell values are presented as Mean ± Std Dev. The ‘Rating’ field is on a scale from 1 to 5.

### A.3 Detecting LLM Use

We use the method proposed by Liang et al. [9, 10] to detect the presence of LLM generated text in both papers and reviews. The method frames the detection problem as a Maximum Likelihood Estimation problem as follow:

- (1) Let  $P$  and  $Q$  denote the probability distribution of documents written by humans and LLMs respectively. Given a new corpus of text  $x$ , we posit that  $x$  was drawn from a weighted sum of these two distributions, i.e.,  $x$  was generated from the mixture distribution

$$(1 - \alpha)P + \alpha Q$$

- (2) The method finds the  $\alpha$  that maximizes

$$\mathcal{L}(\alpha) = \sum_{i=1}^n \log((1 - \alpha)P(x_i) + \alpha Q(x_i)).$$

where  $x_1, x_2, \dots, x_n$  are sentences in text corpus  $x$ .

- (3) The base distributions  $P$  and  $Q$  are determined using unigram frequency models estimated from training data. Specifically,  $P$  is estimated from human-written text, while  $Q$  is estimated from LLM-generated paraphrases of the same text. The authors of Liang et al. [10] provide pre-computed word frequency distributions trained on computer science paper abstracts from arXiv (categories cs.CL and cs.LG) published before November 2023, paired with their LLM generated paraphrases. We use these provided distributions for our analysis.

### A.4 LLM aided and Human Ratings to papers in different quality bins

Figure 7 shows the difference in average scores given by LLM-aided reviews and Human reviews to papers in different quality bins. We only consider papers that received at least one human review. Here, quality is determined by using a left out human review received by the paper, and the remaining human reviews are then used to plot the human curve on the graph. *Contribution, Soundness and Presentation* follow the same trend we saw Rating follow in Figure 3: LLM-aided reviews give slightly higher scores to low quality papers than human reviews but they give similar scores to medium and high quality papers. The *Confidence* trend lines are interesting: both LLM-aided reviews and human reviews show lower confidence when providing ratings to higher quality papers. The confidence scores given by human reviews are slightly higher than those given by LLM-aided reviews, one speculative reason for which could be these reviewers manually assigning a lower confidence rating since they used an LLM to help with the review.




Fig. 7. Average Rating provided by Human reviews and LLM-aided reviews to Papers bucketed by paper quality. Clockwise from top-left, we have Confidence, Contribution, Soundness, and Presentation.

### A.5 Fully LLM and Human Ratings to papers in different quality bins

Figure 8 shows the difference in average scores given by fully LLM generated reviews and Human reviews to papers in different quality bins. Here, quality is determined by using a left out human review received by the paper, and the remaining human reviews are then used to plot the human curve on the graph. Again, *Contribution*, *Soundness* and *Presentation* follow the same trend we saw Rating follow in Figure 5. The fully LLM generated reviews give substantially higher scores to papers in every quality bin and the generated review curve in these three graphs barely has a positive slope. Meanwhile, these reviews seem to uniformly assign a confidence score of 4 to every rating, again implying the reviewers using an LLM post-hoc lowering the confidence score they provide in their review. Note that the quality bucket with rating ‘10’ only has a one paper, which is why the curve trend abruptly deviates at that point.

### A.6 Full Regression Analyses

#### A.6.1 ICLR. Regression with area correction: Using the framing from Equation 2

$$Y_{ij} = \beta_0 + \beta_1 X_i^{\text{LLM}} + \beta_2 T_{ij} + \beta_3 (X_i^{\text{LLM}} \times T_{ij}) + \gamma' Z_i + \epsilon_{ij} \quad (9)$$

where:

- $X_i^{\text{LLM}} \in \{0, 1\}$  is a binary indicator that indicates whether paper  $i$  was written with LLM assistance

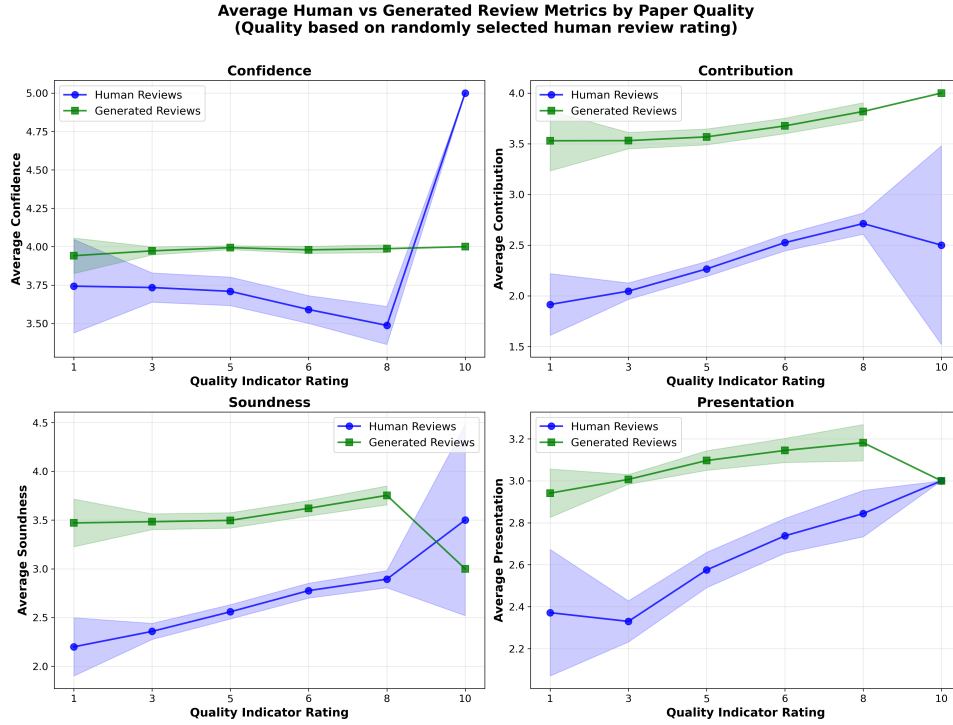


Fig. 8. Average Rating provided by Human reviews and LLM-aided reviews to Papers ordered by their quality percentile. Clockwise from top-left, we have Confidence, Contribution, Soundness, and Presentation.

- $T_{ij} \in \{0, 1\}$  is a binary indicator that indicates whether review  $j$  was written with LLM assistance (the treatment)
- $Z_i$  is a vector of control variables (e.g., area fixed effects). These are the confounders we control for by including them in our regression analysis. For our regression, this is a group of indicator variables that take a value of 1 if the paper belongs to a particular area and 0 otherwise.
- $\epsilon_{ij}$  is the noise term

For ICLR 2024, there are 20 unique research areas and we add an indicator term for all but one reference area in our regression formulation. For ICLR 2025, there are 21 unique research areas. The overlap between paper areas across the two years is not clean, requiring us to develop a mapping to combine the datasets. After mapping equivalent areas, we are left with 23 unique research areas across both years.

Table 15 shows the research areas for each year and the combined dataset. The four area mappings we apply are: (1) “general machine learning” (2024) maps to “other topics in machine learning” (2025), (2) “societal considerations including fairness/safety/privacy” (2024) maps to “alignment/fairness/safety/privacy/societal considerations” (2025), (3) “infrastructure, software libraries, hardware, etc.” (2024) maps to “infrastructure, software libraries, hardware, systems, etc.” (2025), and (4) “representation learning for computer vision/audio/language/other modalities” (2024) maps to “applications to computer vision/audio/language/other modalities” (2025).

ICLR 2024 (20 areas)	ICLR 2025 (21 areas)	Combined (23 areas)
Causal reasoning	Causal reasoning	Causal reasoning
Datasets & benchmarks	Datasets & benchmarks	Datasets & benchmarks
General ML	Other topics in ML	General/other ML
Generative models	Generative models	Generative models
Graph & geometry learning	Graph & geometry learning	Graph & geometry learning
Infrastructure	Infrastructure	Infrastructure
–	Foundation models	Foundation models
–	Interpretability & explainable AI	Interpretability & explainable AI
Learning theory	Learning theory	Learning theory
Metric/kernel/sparse coding	–	Metric/kernel/sparse coding
Representation learning (CV/audio/language)	Multimodal applications	Multimodal applications
Neuroscience & cognitive	Neuroscience & cognitive	Neuroscience & cognitive
Neurosymbolic & hybrid	Neurosymbolic & hybrid	Neurosymbolic & hybrid
Optimization	Optimization	Optimization
Physical sciences	Physical sciences	Physical sciences
Probabilistic methods	Probabilistic methods	Probabilistic methods
Reinforcement learning	Reinforcement learning	Reinforcement learning
Representation learning	Representation learning	Representation learning
Robotics & autonomy	Robotics & autonomy	Robotics & autonomy
Societal considerations	Alignment/fairness/safety/privacy	Societal considerations
–	Time series & dynamical systems	Time series & dynamical systems
Transfer/meta/lifelong	Transfer/meta/lifelong	Transfer/meta/lifelong
Visualization/interpretation	–	Visualization/interpretation

Table 15. Research areas for ICLR 2024, ICLR 2025, and the combined dataset. Areas unique to one year are marked with “–” in the other column. Four area pairs were mapped together for the combined analysis (see text).

*ICLR 2024 Results.* For ICLR 2024, we analyze 27,859 (paper, review) pairs. Among these, 743 papers (2.7%) are classified as LLM-aided, 4,327 reviews (15.5%) are classified as LLM-aided, and 168 pairs (0.6%) have both paper and review classified as LLM-aided.

Table 16 presents the regression coefficients for the main variables. The intercept  $\beta_0 = 5.114$  ( $p < 0.0001$ ) represents the baseline rating for human-written papers with human-written reviews in the baseline area (causal reasoning). The coefficient  $\beta_1 = -1.058$  ( $p < 0.0001$ ) indicates that LLM-aided papers receive substantially lower ratings on average, holding review type and area constant.

The CATE estimates reveal an asymmetric relationship. For human-written papers ( $\text{CATE}(\text{Human}) = \beta_2$ ), LLM-aided reviews increase ratings by 0.271 points ( $p < 0.0001$ , 95% CI: [0.215, 0.328]). For LLM-aided papers ( $\text{CATE}(\text{LLM}) = \beta_2 + \beta_3$ ), LLM-aided reviews increase ratings by 0.615 points ( $p < 0.0001$ , 95% CI: [0.322, 0.909]). The difference in treatment effects ( $\beta_3 = 0.344$ ,  $p = 0.024$ , 95% CI: [0.046, 0.643]) indicates that LLM-aided reviews have a significantly stronger positive effect when evaluating LLM-aided papers.

"Significant" area coefficients are presented in Table 17. Among the area effects, learning theory shows the strongest positive effect (0.583,  $p < 0.0001$ ), while transfer/meta/lifelong learning shows a significant negative effect ( $-0.232$ ,  $p = 0.017$ ). Generative models (0.203,  $p = 0.027$ ) and probabilistic methods (0.223,  $p = 0.037$ ) also show significant positive effects relative to the baseline.

Parameter	Coefficient	Std Error	t-statistic	p-value
$\beta_0$ (Intercept)	5.114	0.086	59.32	< 0.0001
$\beta_1$ (Paper LLM)	-1.058	0.072	-14.66	< 0.0001
$\beta_2$ (Review LLM)	0.271	0.029	9.40	< 0.0001
$\beta_3$ (Both LLM)	0.344	0.152	2.26	0.024
<i>CATE Estimates</i>				
CATE(Human) = $\beta_2$	0.271	0.029	9.40	< 0.0001
CATE(LLM) = $\beta_2 + \beta_3$	0.615	0.150	4.11	< 0.0001
Difference = $\beta_3$	0.344	0.152	2.26	0.024

Table 16. Main regression results for ICLR 2024 (N = 27,859). CATE(Human) represents the effect of LLM-aided reviews on human-written papers. CATE(LLM) represents the effect of LLM-aided reviews on LLM-aided papers.

Area	Coefficient	Std Error	t-statistic	p-value
Causal Reasoning (baseline)	0.000	—	—	—
Generative Models	0.203	0.092	2.22	0.027
Learning Theory	0.583	0.105	5.54	< 0.0001
Probabilistic Methods	0.223	0.107	2.09	0.037
Transfer/Meta/Lifelong	-0.232	0.097	-2.39	0.017

Table 17. Significant area coefficients for ICLR 2024 regression (N = 27,859). Only areas with  $p < 0.05$  are shown. All coefficients are relative to the baseline (causal reasoning).

*ICLR 2025 Results.* For ICLR 2025, we analyze 46,555 (paper, review) pairs. Among these, 1,482 papers (3.2%) are classified as LLM-aided, 12,417 reviews (26.7%) are classified as LLM-aided, and 517 pairs (1.1%) have both paper and review classified as LLM-aided.

Table 18 presents the regression results. The intercept  $\beta_0 = 5.168$  ( $p < 0.0001$ ) is slightly higher than 2024. The coefficient for LLM-aided papers is  $\beta_1 = -0.902$  ( $p < 0.0001$ ), somewhat smaller in magnitude than 2024 but still highly significant and negative.

The CATE estimates show similar patterns to 2024. For human-written papers, LLM-aided reviews increase ratings by 0.257 points ( $p < 0.0001$ , 95% CI: [0.221, 0.292]). For LLM-aided papers, LLM-aided reviews increase ratings by 0.617 points ( $p < 0.0001$ , 95% CI: [0.436, 0.798]). The difference in treatment effects is  $\beta_3 = 0.360$  ( $p = 0.0001$ , 95% CI: [0.175, 0.545]), confirming the asymmetric effect found in 2024.

"Significant" area coefficients are presented in Table 19. Among area coefficients, learning theory again shows the strongest positive effect (0.266,  $p = 0.003$ ). Transfer/meta/lifelong learning ( $-0.370$ ,  $p < 0.0001$ ), time series & dynamical systems ( $-0.279$ ,  $p = 0.003$ ), representation learning ( $-0.246$ ,  $p = 0.003$ ), and general ML ( $-0.243$ ,  $p = 0.005$ ) all show significant negative effects relative to the baseline.

*Combined ICLR 2024-2025 Results.* Combining both years yields 74,414 (paper, review) pairs. Among these, 2,225 papers (3.0%) are classified as LLM-aided, 16,744 reviews (22.5%) are classified as LLM-aided, and 685 pairs (0.9%) have both paper and review classified as LLM-aided.

Table 20 presents the regression results for the combined dataset. The coefficient for LLM-aided papers is  $\beta_1 = -0.962$  ( $p < 0.0001$ ), falling between the 2024 and 2025 values. The CATE estimates show LLM-aided reviews increase ratings by 0.253 points for human-written papers ( $p < 0.0001$ , 95% CI: [0.223, 0.283]) and 0.629 points for LLM-aided papers ( $p < 0.0001$ , 95% CI: [0.476, 0.782]). The difference in treatment effects is  $\beta_3 = 0.376$  ( $p < 0.0001$ , 95% CI: [0.221, 0.532]).

Parameter	Coefficient	Std Error	t-statistic	p-value
$\beta_0$ (Intercept)	5.168	0.078	66.11	< 0.0001
$\beta_1$ (Paper LLM)	-0.902	0.055	-16.27	< 0.0001
$\beta_2$ (Review LLM)	0.257	0.018	14.09	< 0.0001
$\beta_3$ (Both LLM)	0.360	0.094	3.82	0.0001
<i>CATE Estimates</i>				
CATE(Human) = $\beta_2$	0.257	0.018	14.09	< 0.0001
CATE(LLM) = $\beta_2 + \beta_3$	0.617	0.092	6.68	< 0.0001
Difference = $\beta_3$	0.360	0.094	3.82	0.0001

Table 18. Main regression results for ICLR 2025 (N = 46,555). CATE(Human) represents the effect of LLM-aided reviews on human-written papers. CATE(LLM) represents the effect of LLM-aided reviews on LLM-aided papers.

Area	Coefficient	Std Error	t-statistic	p-value
Causal Reasoning (baseline)	0.000	—	—	—
General ML	-0.243	0.086	-2.83	0.005
Learning Theory	0.266	0.090	2.96	0.003
Representation Learning	-0.246	0.083	-2.95	0.003
Time Series & Dynamical	-0.279	0.094	-2.98	0.003
Transfer/Meta/Lifelong	-0.370	0.090	-4.13	< 0.0001

Table 19. Significant area coefficients for ICLR 2025 regression (N = 46,555). Only areas with  $p < 0.05$  are shown. All coefficients are relative to the baseline (causal reasoning).

Parameter	Coefficient	Std Error	t-statistic	p-value
$\beta_0$ (Intercept)	5.145	0.058	88.87	< 0.0001
$\beta_1$ (Paper LLM)	-0.962	0.044	-21.87	< 0.0001
$\beta_2$ (Review LLM)	0.253	0.015	16.52	< 0.0001
$\beta_3$ (Both LLM)	0.376	0.080	4.74	< 0.0001
<i>CATE Estimates</i>				
CATE(Human) = $\beta_2$	0.253	0.015	16.52	< 0.0001
CATE(LLM) = $\beta_2 + \beta_3$	0.629	0.078	8.06	< 0.0001
Difference = $\beta_3$	0.376	0.080	4.74	< 0.0001

Table 20. Main regression results for Combined ICLR 2024-2025 (N = 74,414). CATE(Human) represents the effect of LLM-aided reviews on human-written papers. CATE(LLM) represents the effect of LLM-aided reviews on LLM-aided papers.

Area coefficients are presented in Table 21. The combined analysis shows learning theory with the strongest positive effect (0.383,  $p < 0.0001$ ), while transfer/meta/lifelong learning (-0.304,  $p < 0.0001$ ), time series & dynamical systems (-0.251,  $p = 0.001$ ), representation learning (-0.200,  $p = 0.001$ ), and general ML (-0.163,  $p = 0.009$ ) show significant negative effects. Generative models shows a modest positive effect (0.128,  $p = 0.037$ ).

*Accepted Papers Only.* Meanwhile, when we only consider accepted papers, the patterns change substantially. For ICLR 2024 accepted papers, we analyze 8,706 (paper, review) pairs. Among these, only 96 papers (1.1%) are classified as LLM-aided, 1,447 reviews (16.6%) are classified as LLM-aided, and 31 pairs (0.4%) have both paper and review classified as LLM-aided. For ICLR 2025 accepted papers, we analyze 14,884 pairs, with 349 papers (2.3%)

Area	Coefficient	Std Error	t-statistic	p-value
Causal Reasoning (baseline)	0.000	—	—	—
General ML	-0.163	0.062	-2.61	0.009
Generative Models	0.128	0.061	2.09	0.037
Learning Theory	0.383	0.068	5.62	< 0.0001
Representation Learning	-0.200	0.061	-3.26	0.001
Time Series & Dynamical	-0.251	0.078	-3.24	0.001
Transfer/Meta/Lifelong	-0.304	0.066	-4.62	< 0.0001

Table 21. Significant area coefficients for Combined ICLR 2024-2025 regression (N = 74,414). Only areas with  $p < 0.05$  are shown. All coefficients are relative to the baseline (causal reasoning).

classified as LLM-aided, 4,092 reviews (27.5%) classified as LLM-aided, and 133 pairs (0.9%) with both classified as LLM-aided. The combined accepted papers dataset contains 23,590 pairs, with 445 papers (1.9%) classified as LLM-aided, 5,539 reviews (23.5%) classified as LLM-aided, and 164 pairs (0.7%) with both classified as LLM-aided.

Table 22 presents the regression results for accepted papers across all three datasets. Notably, the coefficient for LLM-aided papers ( $\beta_1$ ) is no longer significantly negative. For ICLR 2024 accepted papers,  $\beta_1 = 0.217$  ( $p = 0.188$ ); for ICLR 2025,  $\beta_1 = -0.045$  ( $p = 0.609$ ); and for the combined dataset,  $\beta_1 = 0.012$  ( $p = 0.880$ ). None of these coefficients are statistically significant, indicating that among accepted papers, LLM-aided papers do not receive systematically different ratings than human-written papers after controlling for research area.

The CATE estimates for accepted papers differ markedly from the full dataset. For ICLR 2024 accepted papers, neither CATE(Human) ( $-0.016$ ,  $p = 0.676$ , 95% CI:  $[-0.091, 0.059]$ ) nor CATE(LLM) ( $-0.102$ ,  $p = 0.724$ , 95% CI:  $[-0.665, 0.462]$ ) is statistically significant. For ICLR 2025 accepted papers, CATE(Human) is small but significant ( $-0.063$ ,  $p = 0.009$ , 95% CI:  $[-0.110, -0.016]$ ), while CATE(LLM) is not significant ( $0.140$ ,  $p = 0.324$ , 95% CI:  $[-0.138, 0.418]$ ). The difference in treatment effects ( $\beta_3$ ) is not significant in either year:  $-0.086$  for 2024 ( $p = 0.768$ ) and  $0.203$  for 2025 ( $p = 0.158$ ).

For the combined accepted papers dataset, CATE(Human) is small but significant ( $-0.051$ ,  $p = 0.011$ , 95% CI:  $[-0.091, -0.012]$ ), indicating that LLM-aided reviews are associated with slightly lower ratings for human-written papers among those that were accepted. However, CATE(LLM) is not significant ( $0.082$ ,  $p = 0.520$ , 95% CI:  $[-0.168, 0.332]$ ), and the difference in treatment effects ( $\beta_3 = 0.134$ ,  $p = 0.301$ , 95% CI:  $[-0.120, 0.387]$ ) is also not significant.

These results indicate that the strong negative effect of LLM-aided papers and the asymmetric positive effect of LLM-aided reviews observed in the full dataset are largely driven by rejected papers. Among accepted papers, LLM usage in papers does not predict significantly different ratings, and the effect of LLM-aided reviews is either negligible or slightly negative. The asymmetric treatment effect that was prominent in the full dataset disappears when restricting to accepted papers.

We also find that most area effects are not statistically significant in the accepted papers analysis. For ICLR 2024 accepted papers, none of the area coefficients reach statistical significance. For ICLR 2025 accepted papers, only transfer/meta/lifelong learning shows a significant negative effect ( $-0.303$ ,  $p = 0.015$ ). For the combined accepted papers dataset, transfer/meta/lifelong learning shows a marginally significant negative effect ( $-0.173$ ,  $p = 0.063$ ). We omit these tables for brevity.

**A.6.2 NeurIPS.** For NeurIPS 2024, there are 31 unique research areas and we add an indicator term for all but one reference area in our regression formulation. For NeurIPS 2025, there are 13 unique research areas. The area taxonomies differ substantially between years, preventing direct combination of the datasets. We therefore analyze each year separately and do not create a combined NeurIPS model.

Dataset	N	$\beta_0$	$\beta_1$	$\beta_2$	$\beta_3$
<i>Main Coefficients</i>					
ICLR 2024 (Accepted)	8,706	6.528 *** (0.120)	0.217 (0.165)	-0.016 (0.038)	-0.086 (0.290)
ICLR 2025 (Accepted)	14,884	6.556 *** (0.106)	-0.045 (0.089)	-0.063 ** (0.024)	0.203 (0.144)
Combined (Accepted)	23,590	6.544 *** (0.079)	0.012 (0.078)	-0.051* (0.020)	0.134 (0.129)
<i>CATE Estimates</i>					
ICLR 2024 (Accepted):					
CATE(Human) = $\beta_2$		-0.016 (SE: 0.038, $p = 0.676$ , 95% CI: [-0.091, 0.059])			
CATE(LLM) = $\beta_2 + \beta_3$		-0.102 (SE: 0.287, $p = 0.724$ , 95% CI: [-0.665, 0.462])			
ICLR 2025 (Accepted):					
CATE(Human) = $\beta_2$		-0.063 (SE: 0.024, $p = 0.009$ , 95% CI: [-0.110, -0.016])			
CATE(LLM) = $\beta_2 + \beta_3$		0.140 (SE: 0.142, $p = 0.324$ , 95% CI: [-0.138, 0.418])			
Combined (Accepted):					
CATE(Human) = $\beta_2$		-0.051 (SE: 0.020, $p = 0.011$ , 95% CI: [-0.091, -0.012])			
CATE(LLM) = $\beta_2 + \beta_3$		0.082 (SE: 0.128, $p = 0.520$ , 95% CI: [-0.168, 0.332])			

Table 22. Regression results for ICLR accepted papers only. Standard errors in parentheses. \*\*\* $p < 0.001$ , \*\* $p < 0.01$ , \* $p < 0.05$ . All regressions control for research area fixed effects. The reduction in statistical significance compared to the full dataset (Table 20) suggests that the patterns observed in the full data are driven primarily by rejected papers.

*NeurIPS 2024 Results.* For NeurIPS 2024, we analyze 16,575 (paper, review) pairs. Among these, 123 papers (0.7%) are classified as LLM-aided, 3,467 reviews (20.9%) are classified as LLM-aided, and 58 pairs (0.3%) have both paper and review classified as LLM-aided.

Table 23 presents the regression results. The intercept  $\beta_0 = 5.782$  ( $p < 0.0001$ ) represents the baseline rating for human-written papers with human-written reviews in the baseline area (active learning). The coefficient for LLM-aided papers is  $\beta_1 = -0.148$  ( $p = 0.316$ ), which is negative but not statistically significant.

The CATE estimates show no significant effects. For human-written papers, LLM-aided reviews increase ratings by 0.022 points ( $p = 0.340$ , 95% CI: [-0.023, 0.067]). For LLM-aided papers, LLM-aided reviews increase ratings by 0.303 points ( $p = 0.156$ , 95% CI: [-0.115, 0.720]). Neither effect is statistically significant. The difference in treatment effects is  $\beta_3 = 0.281$  ( $p = 0.190$ , 95% CI: [-0.139, 0.701]), also not statistically significant.

Parameter	Coefficient	Std Error	t-statistic	p-value
$\beta_0$ (Intercept)	5.782	0.121	47.95	< 0.0001
$\beta_1$ (Paper LLM)	-0.148	0.147	-1.00	0.316
$\beta_2$ (Review LLM)	0.022	0.023	0.95	0.340
$\beta_3$ (Interaction)	0.281	0.214	1.31	0.190
<i>CATE Estimates</i>				
CATE(Human) = $\beta_2$	0.022	0.023	0.95	0.340
CATE(LLM) = $\beta_2 + \beta_3$	0.303	0.213	1.42	0.156
Difference = $\beta_3$	0.281	0.214	1.31	0.190

Table 23. Main regression results for NeurIPS 2024 (N = 16,575). CATE(Human) represents the effect of LLM-aided reviews on human-written papers. CATE(LLM) represents the effect of LLM-aided reviews on LLM-aided papers.

Significant area coefficients are presented in Table 24. Among the area effects, neuroscience and cognitive science shows the strongest positive effect (0.313,  $p = 0.019$ ), followed by learning theory (0.306,  $p = 0.015$ ). Algorithmic game theory (0.301,  $p = 0.037$ ) and probabilistic methods (0.282,  $p = 0.031$ ) also show significant positive effects relative to the baseline.

Area	Coefficient	Std Error	t-statistic	p-value
Active Learning (baseline)	0.000	—	—	—
Algorithmic Game Theory	0.301	0.144	2.09	0.037
Learning Theory	0.306	0.126	2.43	0.015
Neuroscience & Cognitive Science	0.313	0.133	2.35	0.019
Probabilistic Methods	0.282	0.131	2.15	0.031

Table 24. Significant area coefficients for NeurIPS 2024 regression (N = 16,575). Only areas with  $p < 0.05$  are shown. All coefficients are relative to the baseline (active learning).

*NeurIPS 2025 Results.* For NeurIPS 2025, we analyze 22,196 (paper, review) pairs. Among these, 367 papers (1.7%) are classified as LLM-aided, 3,750 reviews (16.9%) are classified as LLM-aided, and 107 pairs (0.5%) have both paper and review classified as LLM-aided.

Table 25 presents the regression results. The intercept  $\beta_0 = 4.291$  ( $p < 0.0001$ ) represents the baseline rating for human-written papers with human-written reviews in the baseline area (applications). The coefficient for LLM-aided papers is  $\beta_1 = -0.081$  ( $p = 0.074$ ), which is negative and marginally significant.

The CATE estimates reveal an interesting pattern. For human-written papers, LLM-aided reviews have essentially no effect ( $-0.006$  points,  $p = 0.663$ , 95% CI:  $[-0.032, 0.020]$ ). For LLM-aided papers, LLM-aided reviews increase ratings by 0.177 points ( $p = 0.034$ , 95% CI:  $[0.013, 0.341]$ ), a statistically significant effect. The difference in treatment effects is  $\beta_3 = 0.183$  ( $p = 0.031$ , 95% CI:  $[0.017, 0.348]$ ), indicating that LLM-aided reviews have a significantly stronger positive effect when evaluating LLM-aided papers compared to human-written papers.

Parameter	Coefficient	Std Error	t-statistic	p-value
$\beta_0$ (Intercept)	4.291	0.012	362.95	< 0.0001
$\beta_1$ (Paper LLM)	-0.081	0.046	-1.79	0.074
$\beta_2$ (Review LLM)	-0.006	0.013	-0.44	0.663
$\beta_3$ (Interaction)	0.183	0.085	2.16	0.031
<i>CATE Estimates</i>				
CATE(Human) = $\beta_2$	-0.006	0.013	-0.44	0.663
CATE(LLM) = $\beta_2 + \beta_3$	0.177	0.083	2.12	0.034
Difference = $\beta_3$	0.183	0.085	2.16	0.031

Table 25. Main regression results for NeurIPS 2025 (N = 22,196). CATE(Human) represents the effect of LLM-aided reviews on human-written papers. CATE(LLM) represents the effect of LLM-aided reviews on LLM-aided papers.

Significant area coefficients are presented in Table 26. Among area coefficients, theory shows the strongest positive effect (0.128,  $p < 0.0001$ ), followed by probabilistic methods (0.070,  $p = 0.010$ ) and neuroscience and cognitive science (0.069,  $p = 0.023$ ). The area labeled "other" shows a significant negative effect ( $-0.143$ ,  $p = 0.003$ ).

Area	Coefficient	Std Error	t-statistic	p-value
Applications (baseline)	0.000	—	—	—
Neuroscience & Cognitive Science	0.069	0.030	2.28	0.023
Other	-0.143	0.048	-3.00	0.003
Probabilistic Methods	0.070	0.027	2.58	0.010
Theory	0.128	0.020	6.44	< 0.0001

Table 26. Significant area coefficients for NeurIPS 2025 regression (N = 22,196). Only areas with  $p < 0.05$  are shown. All coefficients are relative to the baseline (applications).

**A.6.3 ICML.** For ICML 2025, there are 8 unique research areas and we again add an indicator term for all but one reference area in our regression formulation. The research areas are: applications, deep learning, general machine learning, optimization, probabilistic methods, reinforcement learning, social aspects, and theory.

We analyze 13,063 (paper, review) pairs from ICML 2025. Among these, 292 papers (2.2%) are classified as LLM-aided, 2,451 reviews (18.8%) are classified as LLM-aided, and 110 pairs (0.8%) have both paper and review classified as LLM-aided.

Table 27 presents the regression results. The intercept  $\beta_0 = 3.181$  ( $p < 0.0001$ ) represents the baseline rating for human-written papers with human-written reviews in the baseline area (applications). The coefficient for LLM-aided papers is  $\beta_1 = -0.095$  ( $p = 0.111$ ), which is negative but not statistically significant, indicating that LLM-aided papers do not receive systematically different ratings after controlling for research area and review type.

For human-written papers, LLM-aided reviews increase ratings by 0.079 points ( $p < 0.0001$ , 95% CI: [0.043, 0.115]). For LLM-aided papers, LLM-aided reviews increase ratings by 0.164 points ( $p = 0.089$ , 95% CI: [-0.025, 0.354]), though this effect is only marginally significant. The difference in treatment effects is  $\beta_3 = 0.085$  ( $p = 0.386$ , 95% CI: [-0.107, 0.278]), which is not statistically significant, suggesting that unlike ICLR, the effect of LLM-aided reviews does not differ significantly between human-written and LLM-aided papers in ICML.

Parameter	Coefficient	Std Error	t-statistic	p-value
$\beta_0$ (Intercept)	3.181	0.017	188.21	< 0.0001
$\beta_1$ (Paper LLM)	-0.095	0.060	-1.59	0.111
$\beta_2$ (Review LLM)	0.079	0.018	4.33	< 0.0001
$\beta_3$ (Interaction)	0.085	0.098	0.87	0.386
<i>CATE Estimates</i>				
CATE(Human) = $\beta_2$	0.079	0.018	4.33	< 0.0001
CATE(LLM) = $\beta_2 + \beta_3$	0.164	0.097	1.70	0.089
Difference = $\beta_3$	0.085	0.098	0.87	0.386

Table 27. Main regression results for ICML 2025 (N = 13,063). CATE(Human) represents the effect of LLM-aided reviews on human-written papers. CATE(LLM) represents the effect of LLM-aided reviews on LLM-aided papers.

Area coefficients are presented in Table 28. Several areas show significant positive effects relative to the baseline (applications): general machine learning (0.085,  $p = 0.001$ ), probabilistic methods (0.132,  $p = 0.002$ ), and theory (0.108,  $p = 0.0001$ ). No areas show significant negative effects.

Area	Coefficient	Std Error	t-statistic	p-value
Applications (baseline)	0.000	—	—	—
General Machine Learning	0.085	0.024	3.47	0.001
Probabilistic Methods	0.132	0.043	3.08	0.002
Theory	0.108	0.027	3.97	0.0001

Table 28. Significant area coefficients for ICML 2025 regression. Only areas with  $p < 0.05$  are shown. All coefficients are relative to the baseline (applications).

### A.7 Full Within Paper Paired Analyses

We present the full results of our within-paper paired analyses below, including every rating field available to us, like presentation, soundness, contribution, quality, clarity, significance and confidence.

Metric	CATE (LLM-aided Paper)	CATE (Human Paper)	Difference	p-value	Effect
Confidence	-0.012	-0.116	+0.104	0.0073	Small
Soundness	0.189	0.110	+0.080	0.0288	V. Small
Rating	0.340	0.244	+0.096	0.0576	V. Small
Presentation	0.142	0.106	+0.036	0.0897	V. Small
Contribution	0.155	0.117	+0.038	0.1888	V. Small

Table 29. Result of Mann-Whitney U test (one-tailed) on ICLR data (all fields) from 2024 and 2025 combined. Kindness difference = LLM rating – Human rating. Total papers: 10,945 (LLM aided papers: 405, Human papers: 10,540).

Metric	CATE (LLM-aided Paper)	CATE (Human Paper)	Difference	p-value	Effect
Presentation	0.215	0.038	+0.178	0.0122	Small
Soundness	0.148	0.034	+0.114	0.0785	Small
Rating	0.091	-0.029	+0.121	0.1617	Small
Contribution	0.111	0.044	+0.067	0.3070	V. Small
Confidence	-0.105	-0.101	-0.004	0.3963	V. Small

Table 30. Result of Mann-Whitney U test (one-tailed) on ICLR data (accepted papers only) from 2024 and 2025 combined. Kindness difference = LLM rating – Human rating. Total papers: 3,543 (LLM aided papers: 94, Human papers: 3,449).

Metric	CATE (LLM-aided Paper)	CATE (Human Paper)	Difference	p-value	Effect
Contribution	0.336	0.062	+0.274	0.0258	Small
Soundness	0.216	0.059	+0.157	0.0798	Small
Rating	0.310	0.041	+0.269	0.1428	Small
Presentation	0.129	0.057	+0.072	0.2668	V. Small
Confidence	-0.261	-0.081	-0.180	0.8612	Small

Table 31. Result of Mann-Whitney U test (one-tailed) on NeurIPS 2024 data. Kindness difference = LLM rating – Human rating. Total papers: 2,285 (LLM aided papers: 29, Human papers: 2,256).

Metric	CATE (LLM-aided Paper)	CATE (Human Paper)	Difference	p-value	Effect
Significance	0.254	0.072	+0.182	0.0315	Small
Rating	0.154	0.003	+0.151	0.0955	Small
Quality	0.139	0.071	+0.068	0.0997	V. Small
Clarity	0.110	0.075	+0.035	0.2380	V. Small
Originality	0.164	0.097	+0.067	0.2725	V. Small
Confidence	-0.005	-0.058	+0.053	0.3476	V. Small

Table 32. Result of Mann-Whitney U test (one-tailed) on NeurIPS 2025 data. Kindness difference = LLM rating – Human rating. Total papers: 2,697 (LLM aided papers: 66, Human papers: 2,631).

### A.8 Full Synthetic Within Paper Paired Analyses

Metric	CATE (LLM-aided Paper)	CATE (Human Paper)	Difference	p-value	Effect
Rating	2.248	1.623	+0.625	0.0018	Large
Presentation	0.771	0.479	+0.292	0.0001	Small
Contribution	1.521	1.274	+0.246	0.0024	Small
Soundness	1.057	0.960	+0.097	0.1383	V. Small
Confidence	0.142	0.359	-0.217	0.9958	Small

Table 33. Result of Mann-Whitney U test (one-tailed) on ICLR 2025 data with synthetic (fully LLM-generated) reviews. Kindness difference = LLM-generated score – average human score. Total papers: 540 (LLM aided papers: 60, Human papers: 480).

Metric	CATE (LLM-aided Paper)	CATE (Human Paper)	Difference	p-value	Effect
Rating	0.674	0.712	-0.038	0.5670	V. Small
Confidence	0.252	0.450	-0.198	0.9978	Small
Soundness	0.689	0.772	-0.083	0.9205	V. Small
Presentation	0.321	0.322	-0.001	0.7216	V. Small
Contribution	1.037	1.052	-0.015	0.5890	V. Small

Table 34. Result of Mann-Whitney U test (one-tailed) on ICLR 2025 data with synthetic (fully LLM-generated) reviews, accepted papers only. Kindness difference = LLM-generated score – average human score. Total papers: 375 (LLM aided papers: 75, Human papers: 300).

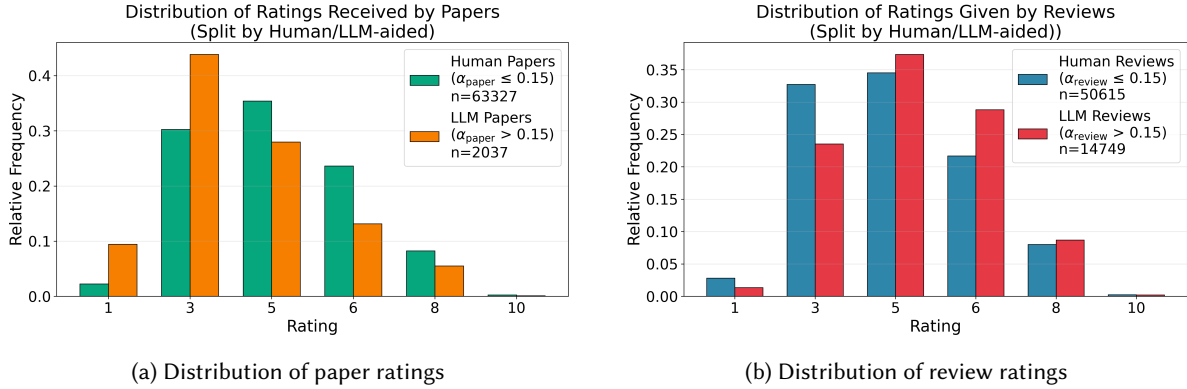


Fig. 9. (a-b) Pre-rebuttal Rating distributions for human vs. LLM-aided papers and reviews for ICLR 2024 and 2025 submissions. LLM-aided papers receive a higher distribution of lower scores while LLM-aided reviews more frequently give middling scores relative to their fully human counterparts. Here LLM-aided means the text showed sufficient similarity to LLM vocabulary ( $\alpha > 0.15$  in Section A.3).

Metric	CATE (LLM-aided Paper)	CATE (Human Paper)	Difference	p-value	Effect
Quality	0.636	0.519	+0.117	0.0468	Small
Clarity	0.391	0.272	+0.119	0.0474	Small
Significance	0.841	0.776	+0.065	0.2042	V. Small
Originality	0.879	0.831	+0.048	0.2634	V. Small
Rating (Overall)	0.408	0.383	+0.025	0.3616	V. Small
Confidence	0.311	0.430	-0.119	0.9619	Small

Table 35. Result of Mann-Whitney U test (one-tailed) on NeurIPS 2025 data with synthetic (fully LLM-generated) reviews. Kindness difference = LLM-generated score – average human score. Total papers: 375 (LLM aided papers: 75, Human papers: 300).

## A.9 Using prerebuttal scores for ICLR 2024-25

*A.9.1 Distribution of paper-wise and review-wise scores.* Figure 9 shows the distribution of scores (pre-rebuttal) received by papers, and scores given by reviews.

*A.9.2 Summary Statistics.* Table 36 shows summary statistics for each quadrant and rating, using pre-rebuttal scores from ICLR 2025. Meanwhile, Table 37 shows summary statistics for accepted papers from ICLR 2025 only. We see similar trends to the summary statistics based on post-rebuttal scores: the differences between quadrants significantly decrease when we restrict ourselves to accepted papers only.

*A.9.3 Full Regression Analyses.* The results of our regression analysis on ICLR Papers and Reviews from 2024 and 2025 are detailed in Table 38.

These results imply that LLM-aided reviewers are kinder than human reviewers to both human papers and LLM-aided papers, but their relative kindness, determined by  $\beta_3$ , is higher for LLM-aided papers than for human papers. This is very close to the observation we made in Table 36, where the differential kindness of LLM-aided reviews to LLM-assisted papers was  $0.63 - 0.37 = 0.26$ .

Again, we repeat the regression analyses on only accepted papers from ICLR 2024 and ICLR 2025 and present the results in Table 39.

Quadrant ( <i>Paper, Review</i> )	Rating	Confidence	Presentation	Contribution	Soundness
(Human, Human)	4.65 ± 1.60	3.68 ± 0.79	2.56 ± 0.77	2.28 ± 0.70	2.52 ± 0.71
(Human, LLM)	5.02 ± 1.47	3.62 ± 0.80	2.69 ± 0.67	2.42 ± 0.66	2.64 ± 0.63
(LLM, Human)	3.83 ± 1.69	3.83 ± 0.80	2.22 ± 0.83	1.99 ± 0.74	2.16 ± 0.75
(LLM, LLM)	4.46 ± 1.69	3.78 ± 0.79	2.43 ± 0.73	2.25 ± 0.73	2.45 ± 0.70

Table 36. Summary Statistics for (Paper, Review) pairs from ICLR 2025. (Paper, Review) pairs are divided into 4 quadrants based on the LLM status of the paper and the review.  $\alpha$  value of 0.15 is used as a threshold: papers and reviews are classified as LLM-assisted if they have  $\alpha > 0.15$ . Cell values are presented as Mean ± Std Dev. All scores are pre-rebuttal.

To understand the discrepancy between results on the accepted-only corpus and the all papers corpus, we examine whether LLM-aided reviews differ from human reviews across the quality spectrum. We stratify papers into four quality buckets (rating  $\leq 3$ ,  $= 5$ ,  $= 6$ ,  $\geq 8$ ) using a randomly selected held-out human review as the quality indicator, then compute average ratings from LLM-assisted and human reviews (excluding the quality-indicator review) within each bucket. Figure 10 reveals two key patterns. First, LLM-aided reviews provide substantially higher ratings than human reviews for low-quality papers, but this gap narrows and disappears as paper quality increases. For papers with quality indicator rating  $\leq 3$ , LLM-aided reviews average approximately 0.4-0.5 points higher than human reviews, shrinking to near-zero for medium and high-quality papers. Second, LLM-assisted papers are heavily overrepresented in the low-quality bucket - mirroring the rating distribution patterns in Figure 9a. The combination of these two factors could potentially be creating the appearance of differential kindness in aggregate analyses.

We probe this quality-dependent relationship further by running separate regressions within each quality bucket using all remaining reviews (excluding the quality-indicator review). Figure 11 plots the regression coefficients for each bucket. The leftmost panel shows that  $\beta_1$  (the effect of a paper being LLM-assisted on ratings from human reviewers) is strongly negative for low-quality papers ( $\approx -0.9$  for rating  $\leq 3$ ) but approaches zero for high-quality papers (rating  $\geq 8$ ,  $\approx 0.3$ )—human reviewers penalize low-quality LLM-aided papers but show minimal bias for high-quality LLM-aided papers. The center panel shows the interaction effect  $\beta_3$  follows an inverted-U pattern: it is small for the lowest-quality papers ( $\approx 0.1$ ), peaks around medium quality (rating 5, where  $\beta_3 \approx 0.35$ ), and decreases substantially for high-quality papers (rating  $\geq 8$ ,  $\approx -0.6$ ). Notably,  $\beta_3$  is positive for low- to medium-quality papers but becomes negative for high-quality papers, though zero continues to be a part of the 95% confidence interval throughout the quality range. The rightmost panel shows  $\beta_1 + \beta_3$  (the effect of a paper being LLM-aided on ratings from LLM-aided reviewers) is strongly negative for low-quality papers ( $\approx -0.8$ ) but increases slightly for medium-quality papers ( $\approx -0.1$  for rating 5-6) before decreasing again for

Quadrant ( <i>Paper, Review</i> )	Rating	Confidence	Presentation	Contribution	Soundness
(Human, Human)	5.71 ± 1.44	3.55 ± 0.78	2.86 ± 0.70	2.65 ± 0.64	2.87 ± 0.62
(Human, LLM)	5.84 ± 1.29	3.53 ± 0.80	2.92 ± 0.61	2.71 ± 0.62	2.89 ± 0.57
(LLM, Human)	5.60 ± 1.34	3.60 ± 0.74	2.80 ± 0.75	2.64 ± 0.64	2.74 ± 0.63
(LLM, LLM)	5.86 ± 1.33	3.63 ± 0.88	2.85 ± 0.59	2.71 ± 0.68	2.82 ± 0.61

Table 37. Summary Statistics for (Paper, Review) pairs from ICLR 2025 (accepted papers only). (Paper, Review) pairs are divided into 4 quadrants based on the LLM status of the paper and the review.  $\alpha$  value of 0.15 is used as a threshold: Papers and Reviews are classified as LLM-aided if they have  $\alpha > 0.15$  and Human otherwise. Cell values are presented as Mean ± Std Dev

Quantity	Interpretation	Estimate	95% CI	p-value
$\beta_0$	Baseline: Human paper, Human review	4.731	[4.617, 4.846]	<0.001
$\beta_2$	CATE: LLM-aided review effect on Human papers	+0.337	[0.307, 0.367]	<0.001
$\beta_2 + \beta_3$	CATE: LLM-aided review effect on LLM papers	+0.603	[0.420, 0.786]	<0.001
$\beta_3$	Differential: Extra kindness to LLM-aided papers	+0.266	[0.113, 0.420]	<0.001

Table 38. Regression Analysis for determining LLM Review Kindness using the model from Eq. (2) on ICLR 2024+ICLR 2025 combined.

Quantity	Interpretation	Estimate	95% CI	p-value
$\beta_0$	Baseline: Human paper, Human review	5.849	[5.658, 6.039]	<0.001
$\beta_2$	CATE: LLM-aided review effect on Human papers	+0.117	[0.068, 0.166]	<0.001
$\beta_2 + \beta_3$	CATE: LLM-aided review effect on LLM-aided papers	+0.264	[-0.047, +0.575]	0.096
$\beta_3$	Differential: Extra kindness to LLM-aided papers	+0.147	[-0.164, +0.458]	0.354

Table 39. Regression Analysis for ICLR accepted papers only (2024+2025 Combined). Model: Rating =  $\beta_0 + \beta_1 X_i^{LLM} + \beta_2 T_{ij} + \beta_3 (X_i^{LLM} \times T_{ij}) + \gamma' Z_i$ . Area fixed effects included. The interaction effect  $\beta_3$  is not statistically significant, in contrast to the full dataset (Table 38).




Fig. 10. Average ratings from LLM-aided reviews vs. human reviews across quality buckets. Shaded area is  $1.96 \times \text{std-error}$ .




Fig. 11. Regression coefficients by quality bucket. Each point represents a separate regression run on papers within that quality bucket. Shaded area is 95% CI. The dashed red line shows the estimate from the full regression for comparison. Zero is indicated with a black dotted line.

high-quality papers ( $\approx -0.3$ )—LLM-aided reviewers penalize low-quality LLM-aided papers (but to a lesser extent than human reviewers) and show minimal differential treatment across the quality spectrum.

These patterns reconcile our earlier findings: the positive  $\beta_3$  estimated by regression over the complete paper corpus reflects the low-middle quality range where the effect is strongest. Restricting to the accepted-papers corpus is similar to restricting to the higher-quality range where zero is in the middle of the shaded confidence interval.

**A.9.4 Within Paper Paired Analysis.** Table 40 shows the results. The mean treatment effect is 0.367 for LLM-assisted papers versus 0.347 for human papers (difference: +0.020), and the p-value of 0.437 makes us reject the null hypothesis. Comparing this number to  $\beta_3$  (0.266) from Table 38 and the estimate from Table 36 (0.26), we find the effect size to be non-existent upon controlling for paper-level confounding factors.

**A.9.5 How influential are LLM-aided reviews in final decisions?** Repeating the analysis from Section 4.1.3 on pre-rebuttal scores, we obtain the results in Table 41. The results are consistent with those for post-rebuttal scores

Metric	CATE (LLM-aided papers)	CATE (Human papers)	Difference	p-value
Rating	0.367	0.347	+0.020	0.437

Table 40. Within-paper treatment effects (ICLR 2024-2025, using pre-rebuttal scores). Mann-Whitney U test (one-tailed) for distributional difference.

(Table 4): Human reviews consistently show higher alignment rates than LLM-aided reviews across most quality levels, with alignment scores between 64% – 90% compared to 59% – 81% for LLM-aided reviews.

Paper Quality	Human Align %	LLM Align %
1	89.7	81.2
3	80.4	72.4
5	74.6	71.3
6	68.0	66.7
8	64.7	66.2
10	65.8	59.3
Overall	74.6	70.6

Table 41. Review alignment rates with final decisions by paper quality for ICLR 2024-2025. Alignment defined as: rating > 5 for accepted papers, rating ≤ 5 for rejected papers. Quality determined by hold-out human review.

#### A.10 Rating Changes During the Rebuttal Period

This appendix documents how individual review ratings evolved between the initial submission of reviews and the final post-rebuttal snapshot. We match reviews across the two snapshots by only considering matches classified as high confidence by Kargaran et al. [5], yielding  $N = 65,670$  paired ratings on the ICLR scale {1, 3, 5, 6, 8, 10}. Each review is additionally labeled by whether the paper is LLM-generated and whether the review itself is LLM-generated (classifier threshold  $\alpha = 0.15$ ).

*A.10.1 How many reviews changed?* Table 42 summarizes the aggregate change. The large majority of ratings (80.5%) are identical across the two snapshots. Among the 19.5% that do change, revisions are overwhelmingly upward: increases outnumber decreases by roughly 17:1. Conditional on a change, the mean signed revision is +1.45 points and the mean absolute revision is 1.65 points; the magnitudes of increases and decreases are comparable (+1.65 vs. -1.74), so the net positive drift is driven by frequency, not size.

Outcome	Count	% of all	Mean $\Delta$
Unchanged	52,867	80.5	0
Increased	12,078	18.4	+1.65
Decreased	725	1.1	-1.74
Any change	12,803	19.5	+1.45

Table 42. Aggregate rating changes between initial and final review snapshots ( $N = 65,670$ ).

The full transition structure is given in Table 43. Movement concentrates near the diagonal: the single most common revision is  $5 \rightarrow 6$  (5,816 reviews, 45.4% of all changes), and transitions spanning more than two scale points are rare (< 1% of matched reviews).

From \ To	1	3	5	6	8	10	Total
1	1,501	48	5	1	0	0	1,555
3	101	16,976	195	38	2	0	17,312
5	20	2,260	16,607	272	22	0	19,181
6	6	806	5,816	12,489	135	2	19,254
8	3	65	453	2,487	5,141	56	8,154
10	0	1	1	3	56	153	214
Total	1,631	20,156	23,077	15,290	5,356	160	65,670

Table 43. Rating transition matrix. Rows index the initial rating; columns index the final rating. Diagonal entries (unchanged ratings) account for 80.5% of the total mass.

*A.10.2 Were changes concentrated in LLM-written reviews?* Table 44 splits revision behavior by whether the review itself is classified as LLM-generated. Human reviewers revise their ratings more often than LLM reviewers (20.1% vs. 17.4% of reviews change), and show a slightly larger mean signed change (+0.30 vs. +0.23). The *size* of each revision is similar across the two groups – the difference is in propensity to revise, not magnitude. Among the 12,750 reviews that do change, the split (79.9% human, 20.1% LLM) closely tracks the base rate in the matched corpus, so the aggregate upward drift is not disproportionately driven by either authorship type.

Reviewer	$N$	Mean $\Delta$	Mean $ \Delta $	% Inc.	Mean Inc.	% Dec.	Mean Dec.	% Unch.
LLM	14,749	+0.23	0.27	16.1	+1.56	1.3	-1.65	82.6
Human	50,615	+0.30	0.34	19.1	+1.67	1.0	-1.77	79.9
All	65,364	+0.28	0.32	18.4	+1.65	1.1	-1.74	80.5

Table 44. Rating changes by reviewer type ( $\alpha = 0.15$ ).

*A.10.3 Were changes concentrated on LLM-written papers?* Table 45 further crosses reviewer type with paper type. Reviews of LLM-authored papers are more stable than reviews of human-authored papers: roughly 87% of ratings on LLM papers are unchanged, compared to  $\approx 80\%$  on human papers. The reviewer-type gap observed in Table 44 persists within each paper type – human reviewers revise more often than LLM reviewers regardless of paper authorship. Of the 12,750 reviews that change, only 2.0% pertain to LLM-authored papers, consistent with the low base rate of LLM papers in the corpus (3.0% of matched reviews).

Paper	Review	$N$	Mean $\Delta$	Mean $ \Delta $	% Inc.	Mean Inc.	% Dec.	Mean Dec.	% Unch.
Human	Human	49,203	+0.30	0.34	19.3	+1.67	1.0	-1.76	79.6
Human	LLM	14,124	+0.23	0.27	16.2	+1.55	1.3	-1.66	82.5
LLM	Human	1,412	+0.17	0.21	11.5	+1.67	0.9	-2.23	87.6
LLM	LLM	625	+0.21	0.22	13.0	+1.68	0.5	-1.33	86.6

Table 45. Rating changes by authorship quadrant.

### A.11 Analysis using Pangram 3.0

We use Pangram 3.0 (Thai et al. [16]), developed by Pangram Labs to detect the presence of LLM written or assisted text in papers and reviews from ICLR 2026. We use the data provided by Pangram from their own analysis on ICLR 2026 data: we scrape this data from their hosted site (<https://iclr.pangram.com/>) and obtain an LLM classification score for every paper and review from ICLR 2026. Our intention is to replicate the experiments from the paper for ICLR 2026, but this time using Pangram’s AI detector, and compare the core results. By using a different detection method, and a different dataset, we aim to simultaneously vary two axes that add robustness to our results.

*Preprocessing note:* Pangram breaks down LLM-assistance percentage for papers into the following buckets: **0-10, 10-30, 30-50, 50-70, 70-90, 90-100**. We consider the first two buckets to be “Human” and the remaining buckets to be “LLM-assisted”. For reviews, Pangram provides tags from **{Fully AI-generated, Heavily AI-edited, Moderately AI-edited, Lightly AI-edited, Fully human-written}**. We consider the last 2 tags to be “Human” and the remaining to be LLM-assisted. The results are as shown below.

Again, when running regression (without area correction) for the whole paper corpus, we seem to see a positive  $\beta_3$ , denoting the existence of a positive interaction effect.

Quantity	Interpretation	Estimate	95% CI	p-value
$\beta_0$	Baseline: Human paper, Human review	4.238	[4.220, 4.256]	<0.001
$\beta_2$	CATE: LLM-aided review effect on Human papers	+0.197	[0.166, 0.228]	<0.001
$\beta_2 + \beta_3$	CATE: LLM-aided review effect on LLM papers	+0.360	[0.297, 0.424]	<0.001
$\beta_3$	Differential: Extra kindness to LLM-aided papers	+0.164	[0.093, 0.234]	<0.001

Table 46. Regression Analysis using Pangram-based LLM detection on ICLR 2026. Model:  $Y_{ij} = \beta_0 + \beta_1 X_i^{\text{LLM}} + \beta_2 T_{ij} + \beta_3 (X_i^{\text{LLM}} \times T_{ij}) + \epsilon_{ij}$ .

This is confirmed by the within-paper paired analysis.

Table 47. Kindness differential: LLM vs. Human reviewers on LLM-authored papers

Group	Mean	Median	<i>n</i>
LLM-aided papers	0.3988	0.5000	2,763
Human papers	0.2632	0.0000	11,793
Mean difference (LLM – Human): <b>0.1356</b>			
Mann–Whitney $U = 16,995,398.50$ , $p = 0.0002$ ***			
Effect size: small			

**Conclusion.** The difference is statistically significant: LLM reviewers rate LLM-authored papers more favorably than human reviewers do ( $\Delta = 0.136$ ,  $p < 0.001$ ).

However, restricting ourselves to high quality papers (avg rating > 5.39 (average rating of an accepted paper in ICLR 2026)) in Table 48, we find that the interaction term is now close to zero, and includes zero in its confidence interval, like the original result from the paper.

These results echo the results from the main paper.

### A.12 Generating Fully LLM Reviews

The prompts used to generate reviews for each of the three conferences are provided below (ICLR: Figure 12, ICML: Figure 13, NeurIPS: Figure 14). The reviewer guidelines for each conference are taken directly from the official

Quantity	Interpretation	Estimate	95% CI	p-value
$\beta_0$	Baseline: Human paper, Human review	6.000	[5.965, 6.035]	<0.001
$\beta_2$	CATE: LLM-aided review effect on Human papers	-0.129	[-0.192, -0.066]	<0.001
$\beta_2 + \beta_3$	CATE: LLM-aided review effect on LLM papers	-0.063	[-0.248, +0.123]	0.509
$\beta_3$	Differential: Extra kindness to LLM-aided papers	+0.066	[-0.130, +0.262]	0.509

Table 48. Regression Analysis using Pangram-based LLM detection for ICLR 2026 accepted papers only. Model:  $Y_{ij} = \beta_0 + \beta_1 X_i^{LLM} + \beta_2 T_{ij} + \beta_3 (X_i^{LLM} \times T_{ij}) + \epsilon_{ij}$ . The interaction effect  $\beta_3$  is not statistically significant, consistent with the  $\alpha$ -based results (Table 6).

conference websites (see International Conference on Learning Representations [3], International Conference on Machine Learning [4], Neural Information Processing Systems [13]).

You are an expert reviewer for a top machine learning conference.  
 A PDF research paper has been attached above. Read the PDF file completely and provide your assessment following the guidelines below. This is a one-shot task and the user will not be able to provide you with any more information.  
 {iclr\_reviewer\_guidelines}  
 Review Template:  
 Provide the review in JSON format with the following fields in order:

1. Summary of the Paper
2. Strengths of the Paper
3. Weaknesses of the Paper
4. Questions for Authors
5. Flag for ethics review: (Yes or no)
6. Presentation: (a score from 1 to 4)
7. Contribution: (a score from 1 to 4)
8. Soundness: (a score from 1 to 4)
9. Rating: A final rating for the paper in accordance with the following rating scale:  
 {iclr\_rating\_scale}
10. Confidence: Your confidence in the rating you provided in accordance with the following confidence scale: {iclr\_confidence\_scale}

This JSON will be automatically parsed, so ensure the format is precise. Do not include numerals in the JSON keys.

Fig. 12. Prompt used to generate ICLR reviews

The prompt used to generate synthetic metareviews is provided in Figure 15. Here, `iclr_rating_scale` and `iclr_confidence_scale` refer to the set of possible ratings and confidence scores reviewers can assign and their given interpretations (for example: 5: Weak Accept). These are provided in detail in Figure 16 and Figure 17.

You are an expert reviewer for a top machine learning conference.  
A PDF research paper has been attached above. Please read the PDF file completely and provide your assessment following the guidelines below. This is a one-shot task and the user will not be able to provide you with any more information.  
{icml\_reviewer\_guidelines}  
Review Template:  
Provide the review in JSON format with the following fields in order, as described in the reviewer guidelines above:

1. Summary
2. Claims and Evidence
3. Methods and Evaluation Criteria
4. Theoretical Claims
5. Experimental Designs or Analyses
6. Supplementary Material
7. Relation To Broader Scientific Literature
8. Essential References Not Discussed
9. Other Strengths and Weaknesses
10. Other Comments or Suggestions
11. Questions for Authors
12. Code of conduct acknowledgement
13. Overall Recommendation (use JSON field name "Rating")

This JSON will be automatically parsed, so ensure the format is precise. Do not include numerals in the JSON keys.

Fig. 13. Prompt used to generate ICML reviews

```

You are an expert reviewer for a top machine learning conference.
A PDF research paper has been attached above. Read the PDF file completely and provide your
assessment following the guidelines below. This is a one-shot task and the user will not be
able to provide you with any more information.
{neurips_reviewer_guidelines}
Review Template:
Provide the review in JSON format with the following fields in order, as described in the
reviewer guidelines above:
1. Summary
2. Strengths and Weaknesses
3. Quality
4. Clarity
5. Significance
6. Originality
7. Questions
8. Limitations
9. Overall
10. Confidence
11. Ethical concerns
12. Code of conduct acknowledgement
This JSON will be automatically parsed, so ensure the format is precise. Do not include
numerals in the JSON keys.

```

Fig. 14. Prompt used to generate NeurIPS reviews

```

You are an area chair for a top machine learning conference.
You have been provided with all the reviews for a paper submission. Your task is to make a
final accept/reject decision based on these reviews.
## Rating Scale for Reference: {iclr_rating_scale}
## Confidence Scale for Reference: {iclr_confidence_scale}
Presentation, Contribution, and Soundness are all scores from 1 to 4.
## Reviews: {{reviews}}
## Your Task: Based on the reviews above, make a final decision for this paper.
Provide your response in JSON format with exactly one field: 1. "decision": Either 1 (Accept)
or 0 (Reject) - use integer values only
Ensure the JSON is valid and parseable. Do not include any text outside the JSON object.

```

Fig. 15. Prompt used to generate ICLR decisions (metareview)

**ICLR Rating Scale**

- **1:** Strong Reject
- **3:** Reject, not good enough
- **5:** Marginally below the acceptance threshold
- **6:** Marginally above the acceptance threshold
- **8:** Accept, good paper
- **10:** Strong accept, should be highlighted at the conference

Fig. 16. ICLR rating scale used in synthetic review generation

**ICLR Confidence Scale**

Confidence values (reviewer's certainty about assessment):

- **1:** You are not confident in your assessment and your review should be taken lightly.
- **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.

Fig. 17. ICLR confidence scale used in synthetic review generation