ARR-Dv2: New Data from ACL Rolling Review and its Potential to Improve Reviewing in the ACL Community

Atnafu Lambebo Tonja[1,*], **LU Sheng**[2], **Nils Dycke**[2], **Dirk Hovy**[3]
Koen Dercksen[4], **Lizhen Qu**[5], **Xiaodan Zhu**[6],
Thamar Solorio[7], **Iryna Gurevych**[2], **Margot Mieskes**[8]

[1]UCL, [2]Ubiquitous Knowledge Processing Lab (UKP Lab), Technical University of Darmstadt,
[3]Bocconi University, [4]Radboud University Medical Center, [5]Monash University,
[6]Queen's University, [7]MBZUAI, [8]University of Applied Sciences, Darmstadt

Abstract

We present a new dataset to *support* AI-assisted and human studies of the peer review process, based on the ARR Open Review Platform. Unlike previous similar data collections, ours (ARR-Dv2) contains not only the **reviews** but also author-reviewer **discussions** and **meta-reviews**. We present statistical analyses of the data with respect to the overall score and the correlations among the overall score, confidence, and soundness scores. Additionally, we extend the task of *Review Score Prediction* to include rebuttal data and analyze its effect on score prediction. Finally, we introduce a new task: *Meta-Review Score Prediction*, which is based on a set of up to three reviews rather than a single review. Our initial results in a zero-shot setting indicate that the rebuttal data adds valuable information to the score prediction and enables reliable predictions.

1 Introduction

Peer review is the key mechanism for quality control of scientific manuscripts across academia (Birukou et al., 2011), playing a vital role in ensuring research integrity and reliable progress (Bornmann, 2011). Commonly, a group of researchers (the *reviewers*) scrutinizes the manuscript and each composes a *review report* that serves as the basis for a *meta-reviewer* to create a condensed overview of the individual reviews, known as *meta-review report*. Despite its importance to the scientific community with the growing number of publications in science (Landhuis, 2016) and Artificial Intelligence (AI) research in particular (Sculley et al., 2018), the reviewing load places significant burden on researchers. This amplifies existing reviewer biases (Lee et al., 2013), encourages outsourcing to junior researchers (Kuznetsov et al., 2024), or – with the rise of generative AI (GenAI) tools – to Large Language Models (LLMs) such as ChatGPT

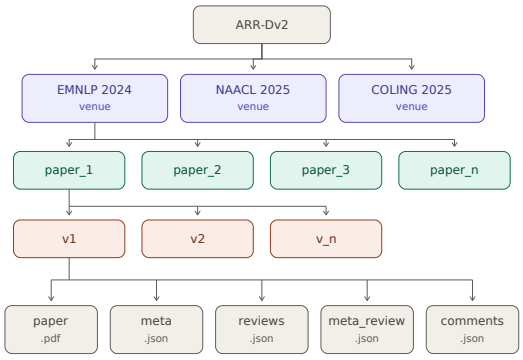


Figure 1: The structure of ARR-Dv2 dataset

(Achiam et al., 2023) and the like (Liang et al., 2024). This, in turn, introduces new biases and quality issues (Du et al., 2024). Putting the reliability of peer review at risk conceivably has severe consequences for the scientific discovery process as a whole (Ebadi et al., 2025; Hanafi et al., 2025). Human-AI collaboration and assistance for peer reviewers is a promising tool to mitigate these issues by speeding up the review process without sacrificing quality, and while maintaining full control for human reviewers (Kuznetsov et al., 2024). Training and evaluation data are essential for the development of such systems. While there is a plethora of peer review corpora (Yuan et al., 2022; Zhang et al., 2025), most of them come from AI research and open reviewing systems. However, review behavior varies between domains and reviewing systems, hereby limiting the generalization of findings to other domains (Kuznetsov et al., 2024). Additionally, using publicly available reviewing data for evaluation poses risks of contamination (Balloccu et al., 2024; Jiang et al., 2024; Ravaut et al., 2024). In this work, we introduce ARR-Dv2, a new peer review corpus collected from ACL Rolling Review (ARR) – a closed reviewing system that ensures its data has not been seen during training of current LLMs – to support the development of peer review

[*] Work done while the author was at MBZUAI

Proceedings of the 15th International Joint Conference on Natural Language Processing and the 5th Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics, pages 333–350
November 6-10, 2026 ©2026 Association for Computational Linguistics

assistance systems tailored to reviewing in natural language processing (NLP) research. Figure 1 shows the structure of ARR-Dv2.

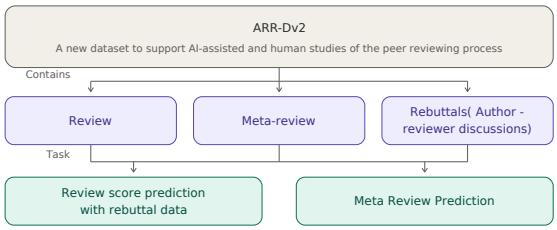


Figure 2: ARR-Dv2 Dataset

This new corpus complements previous work by Dycke et al. (2023) covering NLP reviewing data prior to 2023. However, NLPEER includes only review reports without author rebuttals and reviewer discussion or meta-reviews. ARR-Dv2 includes these important reviewing artifacts (see Figure 2) and is the largest and most recent peer review dataset from the ACL community so far, including the conference of Empirical Methods in Natural Language Processing (EMNLP) 2024, Nations of the Americas Chapter of the ACL (NAACL) 2025, and the International Conference on Computational Linguistics (COLING) 2025. We make the data easily accessible for further research. The dataset enables the study of the submitted and accepted versions of each paper, their reviews, and the author-reviewer discussions, along with meta-reviews, with open licenses and with ethical clearance.

Based on the newly collected ARR-Dv2, we quantify the reviewers' behavior across time and venues. We find that reviewing in NLP and AI across time does show a focus on similar aspects of papers, such as presentation and experimentation, but with varying importance. Additionally, we turn to a practical assistance task during peer review, *review score prediction*, aiming to support reviewers in choosing adequate scores for their textual report; this is particularly challenging for junior reviewers and reviewers from other fields, and may ease their adaptation to the NLP domain. We use the new author-reviewer discussion(rebuttals) and meta-review data from our ARR-Dv2, to investigate these artifacts' effect on score prediction performance. In summary, we contribute

- an extension to the ethically and legally aware collection protocol of Dycke et al. (2022),
- a new corpus of papers, reviews, meta-reviews and author-reviewer discussions from the NLP community from 2024-2025,
- a new perspective on rebuttal and meta-review data during review score prediction, laying the groundwork for the systematic investigation of rebuttals' effect on reviewers' scoring, and
- the exploration of the task of meta-review score prediction in the ACL community.

2 Related Work

Peer Review Corpora have been published over the past years. Most datasets originate from machine learning (ML) conferences such as ICLR and NeurIPS (Staudinger et al., 2024) crawled from open venues hosted on OpenReview[1] (Yuan et al., 2022; Kennard et al., 2022, i.a.) with the most recent and complete version presented by Zhang et al. (2025). Peer review data from other domains includes closed corpora, such as the Elsevier corpus (Willems, 2022), and open reviewing platforms, such as F1000Research[2], which are included in NLPEER (Dycke et al., 2023). Unlike most other datasets, ARR-Dv2 includes peer reviews from a closed-reviewing system, which is currently underrepresented despite its dominance in science (Dycke et al., 2023) and ensures contamination-free evaluation, as the reviews cannot have been part of the pretraining data of current LLMs yet. Only the Peer Review corpus (Kang et al., 2018) and the ARR and COLING subsets of NLPeer (Dycke et al., 2023) cover peer review data from the NLP community. Both corpora do not include meta-reviews or author responses, and they cover substantially fewer papers and reviews compared to ARR-Dv2. Table 1 summarizes the key differences between NLPeer and ARR-Dv2.

Review Score Prediction aims to infer the numerical review scores (e.g., Soundness, Excitement) typically accompanying review reports in the NLP community. These scores are critical to informing meta-reviewers and acceptance decisions. This task has gained increasing attention in recent years. Kang et al. (2018) are the first to study review score prediction, training a simple regression model. Ribeiro et al. (2021) investigate the feasibility of predicting review scores using sentiment labels of review texts and SVMs. Bharti et al. (2021) are the first to propose a neural model that combines both the paper and review text to predict

[1] https://openreview.net
[2] https://f1000research.com/



Figure 3: The procedure and checks we performed to ensure that the data collection was as ethically sound as possible.

review scores. Dycke et al. (2023) evaluate BERT-based models for review score prediction across domains. Fernandes and Vaz-de Melo (2024) conduct a large-scale evaluation of review score prediction models and identify specific inputs that challenge current models. Finally, review score prediction accuracy also has become an important metric for assessing the quality of automatic review generation systems (Yuan and Liu, 2022; Lu et al., 2024; Zhou et al., 2024; Idahl and Ahmadi, 2025; Zhu et al., 2025; Weng et al., 2025). With the newly collected ARR-Dv2 dataset, it is now possible to extend this task to meta-review score prediction using data from the ACL community, whereas prior work has primarily relied on review data from other communities, most notably the machine learning community (Bharti et al., 2021; Lu et al., 2024; Zhu et al., 2025).

3 ARR-Dv2

The following section presents details on the data collection process, the resulting dataset, and initial analyses.

3.1 Data Collection

Peer review in the ACL community is a highly confidential process, where great care is taken to keep the information about authors and reviewers secret. Dycke et al. (2022) presents a donation-based data collection protocol — called the Yes-Yes-Yes- or 3Y-workflow — designed for collecting data from sensitive domains with multiple stakeholders and applying it to peer reviews. Their protocol, developed in consultation with the ACL ethics committee, provides an approved framework for collecting and handling peer review data. Our work extends the existing protocol to include author-reviewer discussions and meta-reviews in the licensing workflow, thereby increasing coverage of the donated artifacts. All modifications have been approved and developed in coordination with the ACL Ethics Committee. Figure 3, summarizes the data collection procedure.

Extended Protocol The 3Y-Workflow is a default opt-out three-step workflow that first asks reviewers to donate their reports, then checks for acceptance of papers, and finally reaches out to authors of accepted papers for their approval of data release. The coverage of donated papers is low, which is arguably attributed to the long time period between submitting a paper and its acceptance, disincentivising authors to donate their review data and drafts. We modify the procedure by asking authors to make a donation decision at submission, with the option to revise it throughout the review process.

As a second major change, we extend the license transfer agreements of the 3Y-workflow to include author-reviewer discussions and rebuttals, requesting license transfers for these artifacts from authors and reviewers during donation. Analogous to reviewers, we ask meta-reviewers to donate their meta-reviews during the reviewing period if the authors approve.

All steps involve explicit disclaimers for all stakeholders. The donation of data is voluntary and can be revoked at any time. See Appendix A for more details on the collection workflow.

Collection We collected ARR-Dv2 in the time period from 2024 to 2025. In accordance with the ACL Rolling Review board, we retrieved the review data using the OpenReview API based on the codebase of Dycke et al. (2022). Compared to the earlier NLPeer dataset, which covers 476 ARR submissions and 684 reviews, ARR-Dv2 offers a substantial expansion: it includes 2838 papers and 1923 reviews, with 52.5% of papers accompanied by rebuttals.

Coverage Voluntary data donation results in subselection, which limits coverage relative to the full set of review data. We quantify coverage as the proportion of donated papers among all accepted papers at each venue. Higher coverage is desirable and helps reduce self-selection bias in the dataset (Dycke et al., 2022). Our dataset covers 23.8% of accepted papers for EMNLP 2024, 19.4% for NAACL 2025, and 22% for COLING 2025. This means that ARR-Dv2 covers, on average, roughly 20% of the accepted submissions, nearly double the rate of the previous collection (approx. 12%).

3.2 Metadata Analysis

We analyze the metadata of ARR-Dv2 collected from EMNLP 2024, COLING 2025, and NAACL

	ARR@NLPeer	ARR-Dv2
time	2021-2022	2024-2025
#papers	476	2838
% accepted	100	100
#reviews	684	1923
#papers w/ rebuttals	-	1403(52.5 %)
Avg. sent per rebuttal	-	15.0 ± 13.4
#reviews per paper	1.43 ± 1.1	0.68 ± 0.75
#sent per paper	220 ± 62.9	538.9 ±199.6
#sent per review	31 ± 27.0	22.9 ± 12.3
total #tok in reviews	266k	822k
total #tok in rebuttals	-	1.8M

Table 1: Comparison of NLPeer (Dycke et al., 2023) and ARR-Dv2in terms of papers, reviews, rebuttals, etc. We report the mean and standard deviation of per-paper and per-review statistics.

2025 to determine differences from the previously collected data from the ACL community, also with respect to self-selection bias.

Comparison to NLPeer As shown in Table 1, ARR-Dv2 covers more data compared to the ARR subset of NLPeer. The number of papers increased from 476 to 2838, and the number of reviews grew nearly threefold, from 684 to 1923. ARR-Dv2 includes rebuttals for more than half of the papers (52.5%), with an average of 15 sentences per rebuttal. Notably, the papers from 2024/25 are longer, while their reviews are shorter, compared to the data up to 2022. This may be because reviewers had more papers to review as the number of submissions increased, and also because the papers themselves often included longer appendices and reference lists.

Review Score Distribution Review scores are a useful proxy for reviewer sentiment. We estimate their distribution within the dataset to quantify the diversity of papers in ARR-Dv2 with respect to quality and review sentiment.
Figure 4 shows the distribution of reviewers overall scores (1-5) for each venue in ARR-Dv2. We observe that the distribution of scores is shifted toward lower values for COLING 2025, whereas NAACL 2025 and EMNLP 2024 are concentrated around 3 to 4, and the highest score (5) is rarely given to papers.

Meta-review Score Distribution A meta-review, usually written by the Area Chair (AC), summarizes the reviews and discussion for a submitted paper (Wu et al., 2022). Figure 5 shows their distribution in ARR-Dv2 the typical score, how spread out the scores are, and whether different venues

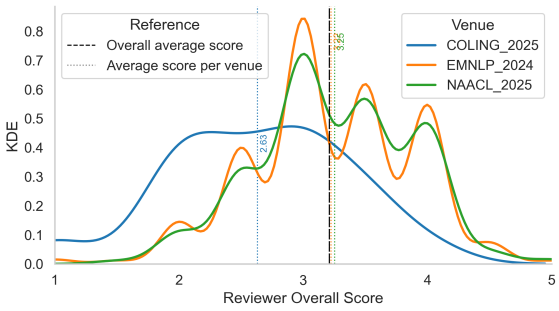


Figure 4: Distribution of review overall scores by venue. Kernel density estimates(KDE) curves show the score distribution for each venue. Colored dotted lines show the average review score per venue(COLING ≈ 2.63, EMNLP ≈ 3.23, NAACL ≈ 3.25), while the black dashed line indicates the overall average score across all reviews.

use the scoring scale differently. The meta-review

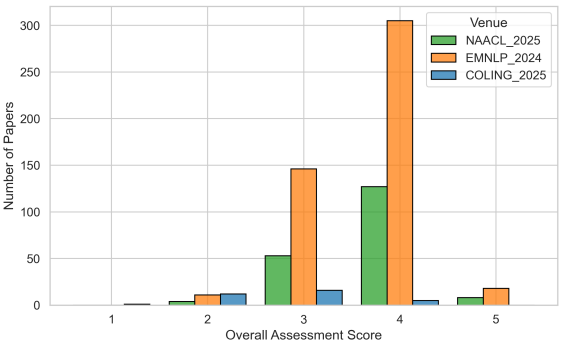


Figure 5: Number of papers by meta-review overall assessment score distribution for each venue.

scores for EMNLP 2024 and NAACL 2025 are primarily at 4, with some at 3, while COLING 2025 has more papers scoring between 2 and 3 and fewer at 4. We observe that some papers in the dataset have notably low meta-review overall assessment scores and were nevertheless accepted. The number of low-scoring papers accepted by each conference are shown in Table 2.

Scoring Behavior We are interested in scoring behavior to explore if evaluations across venues are comparable and to understand whether confidence and soundness relate to overall scores. We want to answer the following research question by exploring scoring behavior. **RQ**: How do soundness and reviewer confidence relate to the overall score, and does this relationship change by venue?
Figure 6 shows the correlation between the various scores. In all three venues, there is a strong positive correlation between overall score and soundness

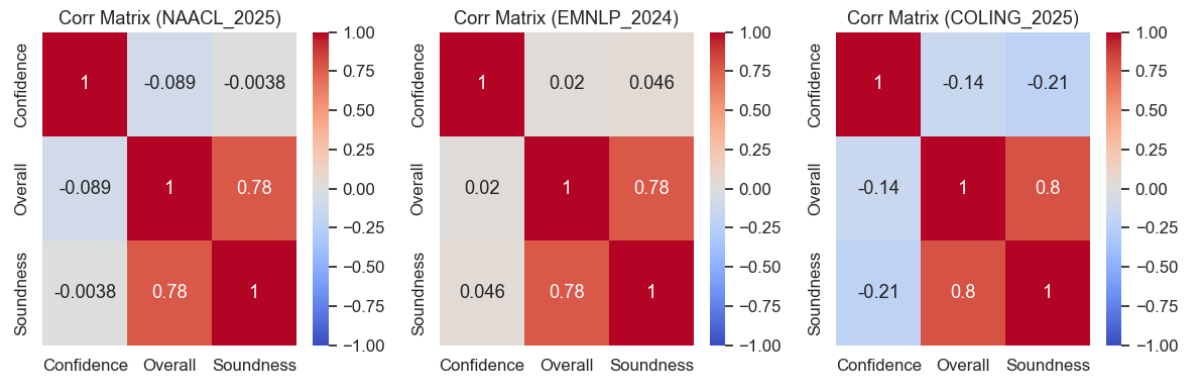


Figure 6: Review-level Pearson correlations between reviewer confidence, overall score, and soundness by venue.

venue	lowest score	#main	#findings
EMNLP 2024	2	2	9
COLING 2025	1	1	-
	2	12	-
NAACL 2025	2	1	3

Table 2: The lowest meta-review overall assessment scores in each venue (**lowest score**), the number of main conference acceptances (**#main**) and findings acceptances (**#findings**). COLING 2025 does not have findings proceedings.

(NAACL (0.78), EMNLP (0.78), COLING (0.80)), confirming that reviewers' overall assessments are closely tied to their evaluation of a paper's correctness. However, the relationship between reviewer confidence and scoring varies more across venues. At NAACL and EMNLP, the correlations between confidence and both overall and soundness scores are near zero (ranging from -0.09 to 0.046), indicating a weak relationship. In contrast, COLING shows a moderate negative correlation between confidence and soundness (-0.21) and overall score (-0.14), suggesting that more confident reviewers may tend to be slightly more critical. This variation at COLING could point to either differences in reviewer behavior or a possible mismatch between self-assessed confidence and actual scoring patterns in that venue.

3.3 Reviewer Focus Analysis

Review aspects are characteristics of a paper that a reviewer makes a judgment on when evaluating the paper (Lu et al., 2025). The distribution of review aspects measures the importance of different paper dimensions to the reviewers, useful for quantifying differences in review behavior across domains and time. In this section, we analyze reviewer focus using the NLPeer dataset, the ICLR dataset, and the newly collected ARR-Dv2 dataset. We present a cross-temporal and cross-venue comparison to show how reviewer focus differs.

Aspect Tagging We perform an aspect-based analysis of the newly collected data using the scheme and tagger by Lu et al. (2025). Their scheme includes two sets of labels based on granularity: the **COARSE** and **FINE** label sets. Table 7 in Appendix C gives examples of aspects for both. Table 3 shows the average number of aspects identified per review for each dataset. The number of aspects per review is generally similar for both the **COARSE** and **FINE** label sets.

dataset	COARSE	FINE
NLPeer ARR22	6.28	9.74
NLPeer COLING2020	7.62	13.42
ARR-Dv2 EMNLP 2024	6.35	10.71
ARR-Dv2 COLING 2025	6.31	11.12

Table 3: Average number of aspects identified per review for each dataset, using the **COARSE** and **FINE** label sets.

Figure 7 shows the five most frequent **COARSE** and **FINE** aspects identified in the reviews for each venue. We also include ICLR 2025 data[3] for comparison, which is not part of the ARR-Dv2 data. For **COARSE** aspects, older data (NLPeer ARR-22 and COLING 2020) show almost identical distributions to the new data. This suggests that the reviewer's focus remains largely stable over time. One notable exception is NLPeer COLING 2020, which includes the aspect DDDDEI (Definition, Description, Detail, Discussion, Explanation, Interpretation), suggesting an emphasis on explainability and interpretability. In all venues, the most frequent **COARSE** aspects are Methodology, Result, Data/Task, and Presentation. There is no major

[3]Publicly available through the OpenReview API.

difference between ACL venues and ICLR 2025, indicating similar reviewer focus between the NLP and ML communities at a COARSE level.

FINE aspects show results similar to those of the COARSE aspect: the reviewer's focus remains stable across time and venues. Common aspects such as Model, Data, Experiment, and Method are consistently mentioned across both earlier and more recent venues. Note that ICLR 2025 has a more balanced aspect distribution. This may suggest that ICLR reviewers place more uniform emphasis across different aspects.

Figure 8 shows the five most frequent **FINE** aspects related to Methodology in the reviews. Reviewers consistently focus on Model and Method, especially in EMNLP 2024 and NAACL 2025. Compared to ACL venues, ICLR 2025 shows a slightly lower emphasis on Model and Method.

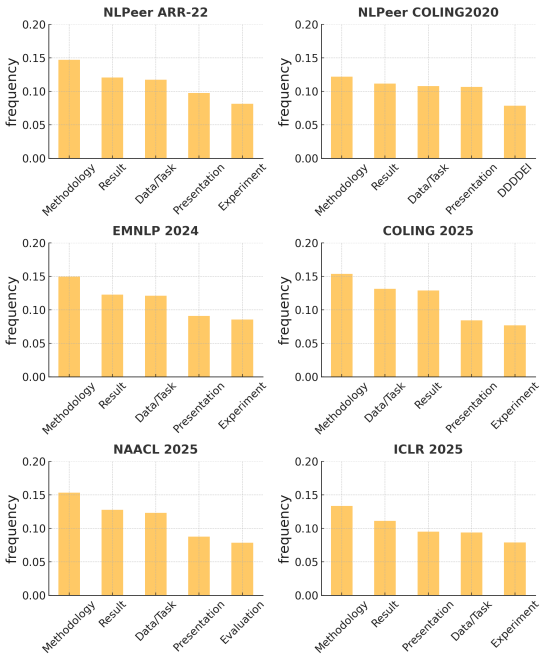
4 Review Score Prediction

Review score prediction is a well-established task in automated peer review (Kang et al., 2018), with the goal of predicting the numerical review score for a given review report, along with potential additional artifacts such as the review template and scoring rubrics. Solving this task requires modeling the relationship between the review text and the overall reviewer opinion, which may depend on the reviewing venue and reviewing habits. Automating this task is motivated by several practical needs: it can assist reviewers in making more consistent and calibrated scoring decisions, help authors identify weaknesses in their submissions prior to peer review, and enable program chairs to detect potential scoring biases or inconsistencies across reviewers (Gao et al., 2019).

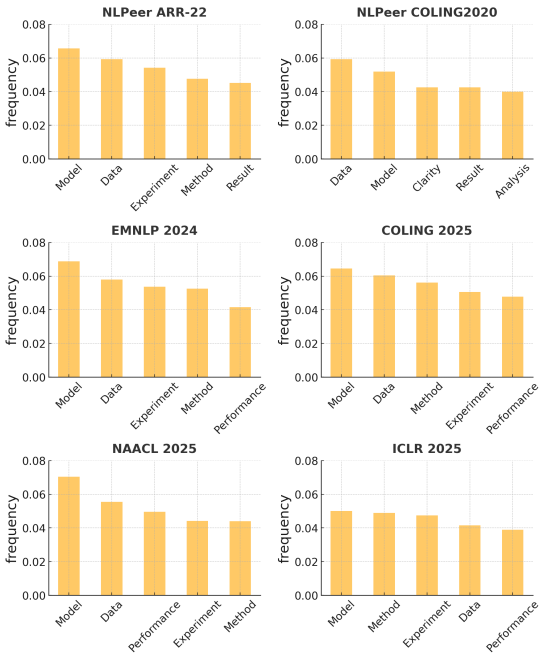
ARR-Dv2 allows us to extend previous work in two important ways: 1) Review score prediction incorporates discussion data (comments by authors and reviewers regarding the initial review and clarification of those comments). This allows us to explore whether incorporating rebuttal data improves the accuracy of review score prediction, which provides insights into how rebuttals influence reviewers' scoring. 2) Meta-review score prediction gives insight into the relationship between scores and text for meta-reviews, too.

4.1 Experimental Settings

For all experiments, we provide the review text together with venue-specific review rubrics in the prompt. In addition to predicting scores, we instruct the model to generate justifications for its score predictions to encourage consistency (see Tables 8 and 9 in the Appendix for the prompts). All experiments are conducted in a zero-shot setting using `Deepseek-R1-Distill-Qwen-14B` (DeepSeek-AI et al., 2025), `Qwen3-(14B & 32B)` (Yang et al., 2025), `LLaMA3-(3B & 8B)` (Grattafiori et al., 2024), and `GPT-OSS-20B` (OpenAI, 2025). We select these models as representative recent



(a) **COARSE**



(b) **FINE**

Figure 7: The 5 most frequent aspects in the reviews for each dataset using the **COARSE** and **FINE** label sets.

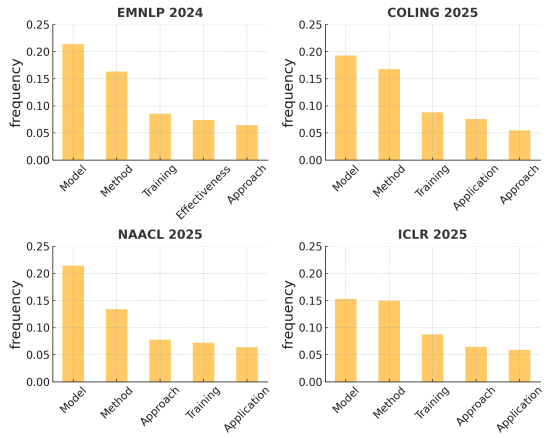


Figure 8: The 5 most frequent **FINE** aspects of Methodology in the reviews.

open-source models with strong reported performance.

Individual Review Score Prediction. We consider two input settings: the individual review alone (**review**) and the individual review plus its corresponding rebuttal (**review+rebuttal**).

Meta-Review Score Prediction. This setting is the major extension of this work, where we do not have any comparison data. We use the same setup as in individual review score prediction and use it as an initial test.

Evaluation. We evaluate model performance using exact match accuracy (**acc@0.0**) and the percentage of score predictions that differ from the human-assigned score by 0.5 points (**acc@0.5**). We also report **sum**, which denotes the sum of **acc@0.0** and **acc@0.5**. Together, these metrics offer an interpretable view of model performance: **acc@0.0** is the most strict metric, while **acc@0.5** allows for a small, reasonable margin of error, and **sum** provides a more balanced view. In addition, we report the root mean squared error (**RMSE**) and macro-F1 score. See Appendix D.1 for more details on the experimental settings.

4.2 Results

Table 4 shows the review score prediction results.

Individual Review Score Prediction. Overall, the models achieve reasonably strong agreement with human-assigned scores. A large number of predictions fall within 0.5 points of human-assigned scores (i.e., **sum** is above 0.8). Model size generally helps: larger models tend to yield better performance, and `Qwen3-32B` is the best-performing model in most cases. The comparison

input	field	model	acc@0.0	acc@0.5	sum	RMSE	f1
M-R	OA	majority	0.64	-	-	0.65	0.20
		random	0.21	-	-	1.70	0.14
		`R1-Qwen-14B`	0.70	-	-	0.58	0.33
		`Qwen3-14B`	**0.72**	-	-	0.55	**0.38**
		`Qwen3-32B`	**0.72**	-	-	**0.54**	0.36
		`LLaMA3-3B`	0.64	-	-	0.67	0.18
		`LLaMA3-8B`	0.68	-	-	0.59	0.31
		`GPT-OSS-20B`	0.68	-	-	0.58	0.34
R	S	majority	0.31	0.33	0.64	0.72	0.05
		random	0.10	0.20	0.30	1.75	0.07
		`R1-Qwen-14B`	0.30	0.46	0.76	0.65	0.10
		`Qwen3-14B`	**0.38**	0.47	**0.85**	**0.54**	**0.18**
		`Qwen3-32B`	0.36	0.46	0.82	0.58	0.13
		`LLaMA3-3B`	0.24	0.46	0.71	0.73	0.10
		`LLaMA3-8B`	0.35	0.47	0.82	0.57	0.13
		`GPT-OSS-20B`	0.28	**0.51**	0.79	0.65	0.11
	OA	majority	0.32	0.38	0.70	0.68	0.05
		random	0.12	0.18	0.30	1.72	0.08
		`R1-Qwen-14B`	0.25	0.49	0.73	0.72	0.10
		`Qwen3-14B`	**0.36**	0.45	**0.82**	**0.58**	**0.16**
		`Qwen3-32B`	**0.36**	0.44	0.81	**0.58**	0.14
		`LLaMA3-3B`	0.34	0.41	0.75	0.66	0.10
		`LLaMA3-8B`	0.30	**0.50**	0.80	0.62	0.11
		`GPT-OSS-20B`	0.29	0.46	0.75	0.68	0.11
R+R	S	majority	0.31	0.33	0.64	0.72	0.05
		random	0.10	0.20	0.30	1.75	0.07
		`R1-Qwen-14B`	0.28	0.38	0.66	0.76	0.09
		`Qwen3-14B`	0.33	0.42	0.74	0.66	0.11
		`Qwen3-32B`	**0.34**	0.42	**0.76**	**0.64**	0.11
		`LLaMA3-3B`	0.15	0.38	0.54	0.94	0.05
		`LLaMA3-8B`	**0.34**	0.34	0.68	0.71	**0.13**
		`GPT-OSS-20B`	0.28	**0.44**	0.72	0.72	0.11
	OA	majority	0.32	0.38	0.70	0.68	0.05
		random	0.12	0.18	0.30	1.72	0.08
		`R1-Qwen-14B`	0.20	0.39	0.58	0.92	0.08
		`Qwen3-14B`	**0.36**	0.39	0.75	0.66	0.14
		`Qwen3-32B`	**0.36**	0.40	**0.76**	**0.64**	**0.15**
		`LLaMA3-3B`	0.28	0.33	0.61	0.84	0.09
		`LLaMA3-8B`	0.33	0.36	0.69	0.75	0.13
		`GPT-OSS-20B`	0.28	**0.42**	0.70	0.75	**0.15**

Table 4: Results of review score prediction under different input settings: M-R refers to meta-review only and does not include rebuttals, review (R), review+rebuttal (R+R). **sum** denotes the sum of **acc@0.0** and **acc@0.5**. The meta-review rubrics do not have 0.5-point scoring, so **acc@0.5** and **sum** are marked as "-." **Bold** indicates the best performance and underline indicates the second best. "OA" is Overall Assessment and "S" is Soundness.

between `R1-Qwen-14B` and `Qwen3-14B` shows that the R1-distilled or reasoning-enhanced objective does not improve performance on this task.

Interestingly, incorporating the rebuttal into the input (**review+rebuttal**) does not seem to improve performance. One possible reason is that rebuttals tend to be positive, which may bias the model toward higher score predictions. To further investigate this, we conduct a more detailed analysis of the results, distinguishing between reviews that include reviewer responses (where the reviewer

replies, discusses with the author, or acknowledges the author's responses) and those without reviewer responses (among the 1923 reviews, 878 include reviewer responses). Table 5 shows results on the subset of reviews with reviewer responses. When reviewer responses are present, using **review+rebuttal** achieves higher **acc@0.0** on Overall Assessment score prediction than using the review alone. This suggests that reviewer responses help mitigate the positivity bias in the rebuttals.

field	model	acc@0.0	
		r	r+r
Soundness	R1-Qwen-14B	**0.32**	**0.32**
	Qwen3-14B	**0.40**	0.35
	Qwen3-32B	**0.37**	0.36
	LLaMA3-3B	**0.26**	0.17
	LLaMA3-8B	**0.35**	0.32
	GPT-OSS-20B	0.29	**0.32**
Overall Assessment	R1-Qwen-14B	**0.27**	0.26
	Qwen3-14B	0.38	**0.42**
	Qwen3-32B	0.39	**0.42**
	LLaMA3-3B	**0.34**	0.30
	LLaMA3-8B	0.32	**0.37**
	GPT-OSS-20B	0.29	**0.33**

Table 5: Results of review score prediction on reviews with reviewer responses.

Meta-Review Score Prediction. Overall, the models achieve good performance, with **acc@0.0** around 0.7 for most models. Larger models generally perform better, while reasoning-enhanced objective does not lead to improvements. Note that meta-review scores are integer values and do not have decimal-based scoring.

Positivity Bias. Results are shown in Table 6, which reports the percentage of cases where the model's predicted score is higher (**+%**) or lower (**-%**) than the human-assigned score. Positivity bias is much lower for meta-reviews than for other inputs. When rebuttal is included, models show a stronger positivity bias compared to using the review alone. This further supports our hypothesis that rebuttals introduce a positivity bias in model predictions.

We further observe that for COLING 2025, models show a strong positivity bias across all tasks and input settings (see Table 10 in the Appendix). One possible reason is incomplete data coverage[4]. Another possible explanation is that the review scores for COLING 2025 are relatively low compared to those of other venues (see Figures 4 and 5 in Section 3.2). As a result, models tend to overestimate scores, leading to a more noticeable positivity bias.

[4]COLING 2025 includes both direct and ARR submissions, while our data only covers ARR submissions (we only have access to ARR submissions).

input	field	model	+%	-%
M-R	OA	R1-Qwen-14B	21.56	8.83
		Qwen3-14B	14.16	14.31
		Qwen3-32B	17.14	11.05
		LLaMA3.2-3B	30.73	5.70
		LLaMA3.1-8B	25.70	6.70
		GPT-OSS-20B	19.41	13.03
R	S	R1-Qwen-14B	49.48	20.42
		Qwen3-14B	33.66	28.04
		Qwen3-32B	47.68	16.66
		LLaMA3.2-3B	63.02	12.57
		LLaMA3.1-8B	43.73	21.29
		GPT-OSS-20B	42.81	29.59
	OA	R1-Qwen-14B	54.32	21.09
		Qwen3-14B	34.13	29.44
		Qwen3-32B	40.29	23.27
		LLaMA3.2-3B	46.56	19.75
		LLaMA3.1-8B	50.57	19.39
		GPT-OSS-20B	43.28	27.65
R+R	S	R1-Qwen-14B	61.98	9.79
		Qwen3-14B	55.47	11.87
		Qwen3-32B	56.87	9.57
		LLaMA3.2-3B	80.10	4.59
		LLaMA3.1-8B	59.30	7.02
		GPT-OSS-20B	57.91	14.23
	OA	R1-Qwen-14B	71.56	8.75
		Qwen3-14B	47.50	16.04
		Qwen3-32B	52.91	11.34
		LLaMA3.2-3B	65.97	6.26
		LLaMA3.1-8B	61.40	5.79
		GPT-OSS-20B	59.81	12.27

Table 6: Proportions of outputs where the LLM predicted score is higher (**+%**) or lower (**-%**) than the human-assigned score.

Our analysis and interpretation assume that human-assigned scores are the gold standard. However, in some cases, human-assigned scores do not fully align with the corresponding review texts. From this perspective, the observed positivity bias in model predictions may also reflect mismatches between review texts and human-assigned scores. To further investigate this possibility, we conduct a small-scale qualitative analysis of cases with large discrepancies between human-assigned and LLM-predicted scores (see Table 11 in the Appendix). We find that while some discrepancies indeed stem from inconsistencies between the review text and the human-assigned scores, in most cases the model predictions tend to be overly positive.

5 Conclusion and Future Work

We introduced ARR-Dv2, a peer review dataset that considerably extends previous data collections in terms of quantity and artifacts, including data from the rebuttal phase and meta-reviews. We analyzed ARR-Dv2 on differences between venues, for example, regarding the link between the overall, confidence and soundness scores. Additionally, we compared the dataset to previous data collections regarding the aspect distribution in reviews. Further, we conducted experiments on review score prediction and observed that including rebuttal information improves prediction. Finally, we extended review score prediction to the task of meta-review score prediction, which shows promising initial results on which we can build in the future. As we have not done any fine-tuning in our experiments so far, this would be an obvious next step. Future work also includes the collection of unpublished data, as the current data collections show a strong bias towards accepted papers. Additionally, we are able to tackle various research questions and peer review tasks outlined by Kuznetsov et al. (2024), such as (meta-) review summarization and providing feedback to reviewers on their scores based on their written reviews.

Limitations

The data analysis and review score prediction experiments may be influenced by potential bias. The current dataset is not comprehensive, as a number of reviews are missing due to constraints on reviewer consent, and our dataset only includes reviews for accepted papers (see Section 3.1 for coverage statistics). As a result, the statistics may over-represent characteristics specific to accepted submissions, and the experimental results may not accurately reflect the models' true predictive capability across a broader range of reviews. Our study focuses on data analysis and score prediction tasks, which enable us to measure the representativeness of the data and the relevance of its constituent artifacts. We do not consider generation tasks, such as review or rebuttal generation, due to the lack of objective and standardized evaluation metrics for these scenarios. We leave the study of these tasks for future work.

Additionally, our main goal and focus in this work was to introduce and analyse the ARR-Dv2 dataset, rather than to optimise model performance. Our experiments are intentionally lightweight and use zero-shot inference as a baseline to illustrate how the new data, especially rebuttals and meta-reviews, can be used. Fine-tuning and architecture-specific optimization would shift the paper's scope from dataset and analysis to model development, which is outside this initial release. We also acknowledge, that our limitation to open-source models does not allow for conclusions to closed-source state-of-the-art models. Nevertheless, open-source models enables local deployment and full control, which are crucial for ensuring confidentiality and data protection in the peer-review context. We do not claim that open-weight models are inherently ethically or legally compliant; rather, open access makes it possible to meet compliance obligations by avoiding transmission of unpublished manuscripts or reviews to third-party servers. In this study, we used only accepted and consented data, but in real-world deployment, AI-assisted reviewer tools process unpublished, confidential submissions. In such settings, using closed-source, cloud-hosted APIs even if private would still need sharing sensitive materials externally, contrary to the confidentiality principles outlined in the ACL Publication Ethics. Open-weight models, by contrast, enable on-premise, auditable, and reproducible deployment that aligns with these ethical requirements.

Reproducing the data collection requires approval by the ACL Ethics Committee as well as support from the ACL Exec. Both can be readily obtained. In order to reproduce the experimental, technical aspect of our work, we refer to the Appendix, where the technical details (models, versions, parameters, and hardware specifications) are given.

Ethical Considerations

This work is associated with several ethical issues. First, the data collection itself and its publication come with their own ethical and legal challenges. We addressed those in the paper. Second, we are not promoting the automation of the peer review process. On the contrary, science requires feedback from peers, and the *peer* review process is at the core of this feedback mechanism. What we are suggesting is using technical *support* to reduce the workload on the reviewers while keeping the human reviewers in the loop.

References

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.

Simone Balloccu, Patrícia Schmidtová, Mateusz Lango, and Ondrej Dusek. 2024. Leak, cheat, repeat: Data contamination and evaluation malpractices in closed-source llms. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics, EACL 2024 - Volume 1: Long Papers, St. Julian's, Malta, March 17-22, 2024*, pages 67–93. Association for Computational Linguistics.

Prabhat Kumar Bharti, Shashi Ranjan, Tirthankar Ghosal, Mayank Agrawal, and Asif Ekbal. 2021. Peerassist: leveraging on paper-review interactions to predict peer review decisions. In *Towards Open and Trustworthy Digital Societies: 23rd International Conference on Asia-Pacific Digital Libraries, ICADL 2021, Virtual Event, December 1–3, 2021, Proceedings 23*, pages 421–435. Springer.

Aliaksandr Birukou, Joseph R. Wakeling, Claudio Bartolini, Fabio Casati, Maurizio Marchese, Katsiaryna Mirylenka, Nardine Osman, Azzurra Ragone, Carles Sierra, and Aalam Wassef. 2011. Alternatives to peer review: Novel approaches for research evaluation. *Frontiers in Computational Neuroscience*, volume 5 - 2011.

Lutz Bornmann. 2011. Scientific peer review. *Annual Review of Information Science and Technology*, 45(1):197–245.

DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Honghui Ding, Huajian Xin, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jiawei Wang, Jingchang Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, J. L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Liang Zhao, Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Meng Li, Miaojun Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang, Qiancheng Wang, Qinyu Chen, Qiushi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, R. J. Chen, R. L. Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shengfeng Ye, Shiyu Wang, Shuiping Yu, Shunfeng Zhou, Shuting Pan, S. S. Li, Shuang Zhou, Shaoqing Wu, Shengfeng Ye, Tao Yun, Tian Pei, Tianyu Sun, T. Wang, Wangding Zeng, Wanjia Zhao, Wen Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, W. L. Xiao, Wei An, Xiaodong Liu, Xiaohan Wang, Xiaokang Chen, Xiaotao Nie, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, X. Q. Li, Xiangyue Jin, Xiaojin Shen, Xiaosha Chen, Xiaowen Sun, Xiaoxiang Wang, Xinnan Song, Xinyi Zhou, Xianzu Wang, Xinxia Shan, Y. K. Li, Y. Q. Wang, Y. X. Wei, Yang Zhang, Yanhong Xu, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Wang, Yi Yu, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Yishi Piao, Yisong Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yuan Ou, Yuduan Wang, Yue Gong, Yuheng Zou, Yujia He, Yunfan Xiong, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Y. X. Zhu, Yanhong Xu, Yanping Huang, Yaohui Li, Yi Zheng, Yuchen Zhu, Yunxian Ma, Ying Tang, Yukun Zha, Yuting Yan, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhenda Xie, Zhengyan Zhang, Zhewen Hao, Zhicheng Ma, Zhigang Yan, Zhiyu Wu, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Zizheng Pan, Zhen Huang, Zhipeng Xu, Zhongyu Zhang, and Zhen Zhang. 2025. DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning. *arXiv preprint*. ArXiv:2501.12948 [cs].

Jiangshu Du, Yibo Wang, Wenting Zhao, Zhongfen Deng, Shuaiqi Liu, Renze Lou, Henry Peng Zou, Pranav Narayanan Venkit, Nan Zhang, Mukund Srinath, Haoran Ranran Zhang, Vipul Gupta, Yinghui Li, Tao Li, Fei Wang, Qin Liu, Tianlin Liu, Pengzhi Gao, Congying Xia, Chen Xing, Cheng Jiayang, Zhaowei Wang, Ying Su, Raj Sanjay Shah, Ruohao Guo, Jing Gu, Haoran Li, Kangda Wei, Zihao Wang, Lu Cheng, Surangika Ranathunga, Meng Fang, Jie Fu, Fei Liu, Ruihong Huang, Eduardo Blanco, Yixin Cao, Rui Zhang, Philip S. Yu, and Wenpeng Yin. 2024. LLMs assist NLP researchers: Critique paper (meta-)reviewing. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 5081–5099, Miami, Florida, USA. Association for Computational Linguistics.

Nils Dycke, Ilia Kuznetsov, and Iryna Gurevych. 2022. Yes-yes-yes: Proactive data collection for ACL rolling review and beyond. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 300–318, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Nils Dycke, Ilia Kuznetsov, and Iryna Gurevych. 2023. NLPeer: A unified resource for the computational study of peer review. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5049–5073, Toronto, Canada. Association for Computational Linguistics.

Saman Ebadi, Hassan Nejadghanbar, Ahmed Rawdhan Salman, and Hassan Khosravi. 2025. Exploring the impact of generative ai on peer review: Insights from journal reviewers. *Journal of Academic Ethics*, pages 1–15.

Gustavo Lúcius Fernandes and Pedro OS Vaz-de Melo. 2024. Enhancing the examination of obstacles in an automated peer review system. *International Journal on Digital Libraries*, 25(2):341–364.

Yang Gao, Steffen Eger, Ilia Kuznetsov, Iryna Gurevych, and Yusuke Miyao. 2019. Does my rebuttal matter? insights from a major NLP conference. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1274–1290, Minneapolis, Minnesota. Association for Computational Linguistics.

Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Roziere, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Allonsius, Daniel Song, Danielle Pintz, Danny Livshits, Danny Wyatt, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip Radenovic, Francisco Guzmán, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Govind Thattai, Graeme Nail, Gregoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel Kloumann, Ishan Misra, Ivan Evtimov, Jack Zhang, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Alwala, Karthik Prasad, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, Khalid El-Arini, Krithika Iyer, Kshitiz Malik, Kuenley Chiu, Kunal Bhalla, Kushal Lakhotia, Lauren Rantala-Yeary, Laurens van der Maaten, Lawrence Chen, Liang Tan, Liz Jenkins, Louis Martin, Lovish Madaan, Lubo Malo, Lukas Blecher, Lukas Landzaat, Luke de Oliveira, Madeline Muzzi, Mahesh Pasupuleti, Mannat Singh, Manohar Paluri, Marcin Kardas, Maria Tsimpoukelli, Mathew Oldham, Mathieu Rita, Maya Pavlova, Melanie Kambadur, Mike Lewis, Min Si, Mitesh Kumar Singh, Mona Hassan, Naman Goyal, Narjes Torabi, Nikolay Bashlykov, Nikolay Bogoychev, Niladri Chatterji, Ning Zhang, Olivier Duchenne, Onur Çelebi, Patrick Alrassy, Pengchuan Zhang, Pengwei Li, Petar Vasic, Peter Weng, Prajjwal Bhargava, Pratik Dubal, Praveen Krishnan, Punit Singh Koura, Puxin Xu, Qing He, Qingxiao Dong, Ragavan Srinivasan, Raj Ganapathy, Ramon Calderer, Ricardo Silveira Cabral, Robert Stojnic, Roberta Raileanu, Rohan Maheswari, Rohit Girdhar, Rohit Patel, Romain Sauvestre, Ronnie Polidoro, Roshan Sumbaly, Ross Taylor, Ruan Silva, Rui Hou, Rui Wang, Saghar Hosseini, Sahana Chennabasappa, Sanjay Singh, Sean Bell, Seohyun Sonia Kim, Sergey Edunov, Shaoliang Nie, Sharan Narang, Sharath Raparthy, Sheng Shen, Shengye Wan, Shruti Bhosale, Shun Zhang, Simon Vandenhende, Soumya Batra, Spencer Whitman, Sten Sootla, Stephane Collot, Suchin Gururangan, Sydney Borodinsky, Tamar Herman, Tara Fowler, Tarek Sheasha, Thomas Georgiou, Thomas Scialom, Tobias Speckbacher, Todor Mihaylov, Tong Xiao, Ujjwal Karn, Vedanuj Goswami, Vibhor Gupta, Vignesh Ramanathan, Viktor Kerkez, Vincent Gonguet, Virginie Do, Vish Vogeti, Vítor Albiero, Vladan Petrovic, Weiwei Chu, Wenhan Xiong, Wenyin Fu, Whitney Meers, Xavier Martinet, Xiaodong Wang, Xiaofang Wang, Xiaoqing Ellen Tan, Xide Xia, Xinfeng Xie, Xuchao Jia, Xuewei Wang, Yaelle Goldschlag, Yashesh Gaur, Yasmine Babaei, Yi Wen, Yiwen Song, Yuchen Zhang, Yue Li, Yuning Mao, Zacharie Delpierre Coudert, Zheng Yan, Zhengxing Chen, Zoe Papakipos, Aaditya Singh, Aayushi Srivastava, Abha Jain, Adam Kelsey, Adam Shajnfeld, Adithya Gangidi, Adolfo Victoria, Ahuva Goldstand, Ajay Menon, Ajay Sharma, Alex Boesenberg, Alexei Baevski, Allie Feinstein, Amanda Kallet, Amit Sangani, Amos Teo, Anam Yunus, Andrei Lupu, Andres Alvarado, Andrew Caples, Andrew Gu, Andrew Ho, Andrew Poulton, Andrew Ryan, Ankit Ramchandani, Annie Dong, Annie Franco, Anuj Goyal, Aparajita Saraf, Arkabandhu Chowdhury, Ashley Gabriel, Ashwin Bharambe, Assaf Eisenman, Azadeh Yazdan, Beau James, Ben Maurer, Benjamin Leonhardi, Bernie Huang, Beth Loyd, Beto De Paola, Bhargavi Paranjape, Bing Liu, Bo Wu, Boyu Ni, Braden Hancock, Bram Wasti, Brandon Spence, Brani Stojkovic, Brian Gamido, Britt Montalvo, Carl Parker, Carly Burton, Catalina Mejia, Ce Liu, Changhan Wang, Changkyu Kim, Chao Zhou, Chester Hu, Ching-Hsiang Chu, Chris Cai, Chris Tindal, Christoph Feichtenhofer, Cynthia Gao, Damon Civin, Dana Beaty, Daniel Kreymer, Daniel Li, David Adkins, David Xu, Davide Testuggine, Delia David, Devi Parikh, Diana Liskovich, Didem Foss, Dingkang Wang, Duc Le, Dustin Holland, Edward Dowling, Eissa Jamil, Elaine Montgomery, Eleonora Presani, Emily Hahn, Emily Wood, Eric-Tuan Le, Erik Brinkman, Esteban Arcaute, Evan Dunbar, Evan Smothers, Fei Sun, Felix Kreuk, Feng Tian, Filippos Kokkinos, Firat Ozgenel, Francesco Caggioni, Frank Kanayet, Frank Seide, Gabriela Medina Florez, Gabriella Schwarz, Gada Badeer, Georgia Swee, Gil Halpern, Grant Herman, Grigory Sizov, Guangyi, Zhang, Guna Lakshminarayanan, Hakan Inan, Hamid Shojanazeri, Han Zou, Hannah Wang, Hanwen Zha, Haroun Habeeb, Harrison Rudolph, Helen Suk, Henry Aspegren, Hunter Goldman, Hongyuan Zhan, Ibrahim Damlaj, Igor Molybog, Igor Tufanov, Ilias Leontiadis, Irina-Elena Veliche, Itai Gat, Jake Weissman, James Geboski, James Kohli, Janice Lam, Japhet Asher,

Jean-Baptiste Gaya, Jeff Marcus, Jeff Tang, Jennifer Chan, Jenny Zhen, Jeremy Reizenstein, Jeremy Teboul, Jessica Zhong, Jian Jin, Jingyi Yang, Joe Cummings, Jon Carvill, Jon Shepard, Jonathan McPhie, Jonathan Torres, Josh Ginsburg, Junjie Wang, Kai Wu, Kam Hou U, Karan Saxena, Kartikay Khandelwal, Katayoun Zand, Kathy Matosich, Kaushik Veeraraghavan, Kelly Michelena, Keqian Li, Kiran Jagadeesh, Kun Huang, Kunal Chawla, Kyle Huang, Lailin Chen, Lakshya Garg, Lavender A, Leandro Silva, Lee Bell, Lei Zhang, Liangpeng Guo, Licheng Yu, Liron Moshkovich, Luca Wehrstedt, Madian Khabsa, Manav Avalani, Manish Bhatt, Martynas Mankus, Matan Hasson, Matthew Lennie, Matthias Reso, Maxim Groshev, Maxim Naumov, Maya Lathi, Meghan Keneally, Miao Liu, Michael L. Seltzer, Michal Valko, Michelle Restrepo, Mihir Patel, Mik Vyatskov, Mikayel Samvelyan, Mike Clark, Mike Macey, Mike Wang, Miquel Jubert Hermoso, Mo Metanat, Mohammad Rastegari, Munish Bansal, Nandhini Santhanam, Natascha Parks, Natasha White, Navyata Bawa, Nayan Singhal, Nick Egebo, Nicolas Usunier, Nikhil Mehta, Nikolay Pavlovich Laptev, Ning Dong, Norman Cheng, Oleg Chernoguz, Olivia Hart, Omkar Salpekar, Ozlem Kalinli, Parkin Kent, Parth Parekh, Paul Saab, Pavan Balaji, Pedro Rittner, Philip Bontrager, Pierre Roux, Piotr Dollar, Polina Zvyagina, Prashant Ratanchandani, Pritish Yuvraj, Qian Liang, Rachad Alao, Rachel Rodriguez, Rafi Ayub, Raghotham Murthy, Raghu Nayani, Rahul Mitra, Rangaprabhu Parthasarathy, Raymond Li, Rebekkah Hogan, Robin Battey, Rocky Wang, Russ Howes, Ruty Rinott, Sachin Mehta, Sachin Siby, Sai Jayesh Bondu, Samyak Datta, Sara Chugh, Sara Hunt, Sargun Dhillon, Sasha Sidorov, Satadru Pan, Saurabh Mahajan, Saurabh Verma, Seiji Yamamoto, Sharadh Ramaswamy, Shaun Lindsay, Shaun Lindsay, Sheng Feng, Shenghao Lin, Shengxin Cindy Zha, Shishir Patil, Shiva Shankar, Shuqiang Zhang, Shuqiang Zhang, Sinong Wang, Sneha Agarwal, Soji Sajuyigbe, Soumith Chintala, Stephanie Max, Stephen Chen, Steve Kehoe, Steve Satterfield, Sudarshan Govindaprasad, Sumit Gupta, Summer Deng, Sungmin Cho, Sunny Virk, Suraj Subramanian, Sy Choudhury, Sydney Goldman, Tal Remez, Tamar Glaser, Tamara Best, Thilo Koehler, Thomas Robinson, Tianhe Li, Tianjun Zhang, Tim Matthews, Timothy Chou, Tzook Shaked, Varun Vontimitta, Victoria Ajayi, Victoria Montanez, Vijai Mohan, Vinay Satish Kumar, Vishal Mangla, Vlad Ionescu, Vlad Poenaru, Vlad Tiberiu Mihailescu, Vladimir Ivanov, Wei Li, Wenchen Wang, Wenwen Jiang, Wes Bouaziz, Will Constable, Xiaocheng Tang, Xiaojian Wu, Xiaolan Wang, Xilun Wu, Xinbo Gao, Yaniv Kleinman, Yanjun Chen, Ye Hu, Ye Jia, Ye Qi, Yenda Li, Yilin Zhang, Ying Zhang, Yossi Adi, Youngjin Nam, Yu, Wang, Yu Zhao, Yuchen Hao, Yundi Qian, Yunlu Li, Yuzi He, Zach Rait, Zachary DeVito, Zef Rosnbrick, Zhaoduo Wen, Zhenyu Yang, Zhiwei Zhao, and Zhiyu Ma. 2024. The Llama 3 Herd of Models. *arXiv preprint*. ArXiv:2407.21783 [cs].

Ahmed M Hanafi, Mohamed M Al-mansi, Omar A Al-Sharif, et al. 2025. Generative ai in academia: A comprehensive review of applications and implications for the research process. *International Journal of Engineering and Applied Sciences-October 6 University*, 2(1):91–110.

Maximilian Idahl and Zahra Ahmadi. 2025. OpenReviewer: A specialized large language model for generating critical scientific paper reviews. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (System Demonstrations)*, pages 550–562, Albuquerque, New Mexico. Association for Computational Linguistics.

Minhao Jiang, Ken Liu, Ming Zhong, Rylan Schaeffer, Siru Ouyang, Jiawei Han, and Sanmi Koyejo. 2024. Does data contamination make a difference? insights from intentionally contaminating pre-training data for language models. In *ICLR 2024 Workshop on Navigating and Addressing Data Problems for Foundation Models*.

Dongyeop Kang, Waleed Ammar, Bhavana Dalvi, Madeleine van Zuylen, Sebastian Kohlmeier, Eduard Hovy, and Roy Schwartz. 2018. A dataset of peer reviews (PeerRead): Collection, insights and NLP applications. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1647–1661, New Orleans, Louisiana. Association for Computational Linguistics.

Neha Kennard, Tim O'Gorman, Rajarshi Das, Akshay Sharma, Chhandak Bagchi, Matthew Clinton, Pranay Kumar Yelugam, Hamed Zamani, and Andrew McCallum. 2022. DISAPERE: A dataset for discourse structure in peer review discussions. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1234–1249, Seattle, United States. Association for Computational Linguistics.

Ilia Kuznetsov, Osama Mohammed Afzal, Koen Dercksen, Nils Dycke, Alexander Goldberg, Tom Hope, Dirk Hovy, Jonathan K Kummerfeld, Anne Lauscher, Kevin Leyton-Brown, et al. 2024. What can natural language processing do for peer review? *arXiv preprint arXiv:2405.06563*.

Esther Landhuis. 2016. Scientific literature: Information overload. *Nature*, 535(7612):457–458.

Carole J Lee, Cassidy R Sugimoto, Guo Zhang, and Blaise Cronin. 2013. Bias in peer review. *Journal of the American Society for information Science and Technology*, 64(1):2–17.

Weixin Liang, Zachary Izzo, Yaohui Zhang, Haley Lepp, Hancheng Cao, Xuandong Zhao, Lingjiao Chen, Haotian Ye, Sheng Liu, Zhi Huang, Daniel 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, volume 235 of *Proceedings of Machine Learning Research*, pages 29575–29620. PMLR.

Chris Lu, Cong Lu, Robert Tjarko Lange, Jakob N. Foerster, Jeff Clune, and David Ha. 2024. The AI scientist: Towards fully automated open-ended scientific discovery. *CoRR*, abs/2408.06292.

Sheng Lu, Ilia Kuznetsov, and Iryna Gurevych. 2025. Identifying aspects in peer reviews. *arXiv preprint arXiv:2504.06910*.

OpenAI. 2025. gpt-oss-120b gpt-oss-20b model card. *Preprint*, arXiv:2508.10925.

Mathieu Ravaut, Bosheng Ding, Fangkai Jiao, Hailin Chen, Xingxuan Li, Ruochen Zhao, Chengwei Qin, Caiming Xiong, and Shafiq Joty. 2024. How much are llms contaminated? A comprehensive survey and the llmsanitize library. *CoRR*, abs/2404.00699.

Ana Carolina Ribeiro, Amanda Sizo, Henrique Lopes Cardoso, and Luís Paulo Reis. 2021. Acceptance decision prediction in peer-review through sentiment analysis. In *EPIA Conference on Artificial Intelligence*, pages 766–777. Springer.

D Sculley, Jasper Snoek, and Alex Wiltschko. 2018. Avoiding a tragedy of the commons in the peer review process. *arXiv preprint arXiv:1901.06246*.

Moritz Staudinger, Wojciech Kusa, Florina Piroi, and Allan Hanbury. 2024. An analysis of tasks and datasets in peer reviewing. In *Proceedings of the Fourth Workshop on Scholarly Document Processing (SDP 2024)*, pages 257–268, Bangkok, Thailand. Association for Computational Linguistics.

Yixuan Weng, Minjun Zhu, Guangsheng Bao, Hongbo Zhang, Jindong Wang, Yue Zhang, and Linyi Yang. 2025. Cycleresearcher: Improving automated research via automated review. In *The Thirteenth International Conference on Learning Representations*.

Linda Willems. 2022. New dataset offers unique insights into peer review. Accessed: 2025-06-06.

Po-Cheng Wu, An-Zi Yen, Hen-Hsen Huang, and Hsin-Hsi Chen. 2022. Incorporating peer reviews and rebuttal counter-arguments for meta-review generation. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, pages 2189–2198.

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.

Weizhe Yuan and Pengfei Liu. 2022. Kid-review: knowledge-guided scientific review generation with oracle pre-training. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 11639–11647.

Weizhe Yuan, Pengfei Liu, and Graham Neubig. 2022. Can we automate scientific reviewing? *Journal of Artificial Intelligence Research*, 75:171–212.

Daoze Zhang, Zhijian Bao, Sihang Du, Zhiyi Zhao, Kuangling Zhang, Dezheng Bao, and Yang Yang. 2025. Re^2: A consistency-ensured dataset for full-stage peer review and multi-turn rebuttal discussions. *arXiv preprint arXiv:2505.07920*.

Ruiyang Zhou, Lu Chen, and Kai Yu. 2024. Is LLM a reliable reviewer? a comprehensive evaluation of LLM on automatic paper reviewing tasks. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 9340–9351, Torino, Italia. ELRA and ICCL.

Minjun Zhu, Yixuan Weng, Linyi Yang, and Yue Zhang. 2025. Deepreview: Improving llm-based paper review with human-like deep thinking process. *CoRR*, abs/2503.08569.

A More on Data Collection

We use the OpenReview Python package to collect the data. The data collection process adheres to the Yes-Yes-Yes protocol (Dycke et al., 2022), which involves two key license agreements: **Peer Reviewer Content License Agreement** (hereinafter referred to as the *reviewer agreement*) and **Blind Submission License Agreement** (hereinafter referred to as the *author agreement*). The main steps are as follows:

- **Obtain reviewer agreements**
  Use `or_client.get_all_notes(invitation=...)`, where `invitation` is the link to the reviewer agreement in OpenReview.
- **Retrieve accepted papers**
  Directly from the ACL Anthology Repository.
- **Collect submission details**
  Use `or_client.get_all_notes(invitation=...)`, where `invitation` is the link to the submissions in OpenReview, including reviews, rebuttals, and the author agreements.

- **Filter the submissions**
  Iterate through all submissions and select those that meet all of the criteria of the Yes-Yes-Yes protocol.

B Data analysis - soundness and overall score by confidence

Figure 9 shows an analysis of our data with respect to confidence and soundness. To maintain readability, the y-axis is skewed in each of these plots. But we can see that we have very few papers with a confidence of 1 or 2. The majority of the reviewers give their confidence as either 3 or 4, and a few reviewers claim a confidence of 5. We can also see that the distribution focuses around a soundness score of 3, indicating that the distribution approximately follows a Gaussian Distribution.

Figure 10 shows a similar distribution for confidence in relation to the overall score. Few reviewers give a low confidence rating of 1 or 2, while the majority note a confidence rating of 3 or 4, with a smaller group giving a confidence rating of 5. Again, the scores resemble a Gaussian distribution, with the majority of scores being 3.

Both Figures 9 and 10 show that contrary to the previous data collection our dataset also contains papers that would probably be rejected for the conference, as overall scores of 1 and 2 (a total of 80 papers for NAACL, 182 papers for EMNLP, and 31 papers for COLING) would probably be a clear reject, while papers with 3 (217 papers for NAACL, 657 papers for EMNLP and 26 papers for COLING) are borderline and papers with an overall score of 5 (3 papers for NAACL, 6 papers for EMNLP) are very likely to be accepted. Across all three NLP venues, both the soundness and overall score distributions per reviewer confidence show a strong central tendency in which the vast majority of reviews, regardless of venue or confidence level, assign moderate scores, with very few ratings for 1 or 5.

C More on Data Analysis

Table 7 shows examples of aspects derived by Lu et al. (2025).

COARSE	FINE
Contribution	Contribution
DDDDEI	Definition, Description, Detail, Discussion, Explanation, Interpretation
IJMV	Intuition, Justification, Motivation, Validation
Novelty	Innovation/Novelty/Originality
Presentation	Clarity, Figure, Grammar, Presentation, Typo
Related Work	Citation/Literature/Related Work
Significance	Impact, Importance, Significance
Ablation	Ablation
Analysis	Analysis
Comparison	Comparison
Data/Task	Annotation, Benchmark, Data, Task
Evaluation	Evaluation, Metric
Experiment	Experiment
Methodology	Algorithm, Implementation, Method
Theory	Theory
Result	Findings, Improvement, Performance, Result

Table 7: Examples of aspects in the COARSE and FINE label sets (Lu et al., 2025).

D More on Review Score Prediction

D.1 More on Experiment Settings

For all experiments, we set `temperature=0.5`, `max_tokens=1024`, and `seed=2266`. We used 4-bit quantization for model inference. All experiments were done on either an NVIDIA A100 or H100 GPU. Tables 8 and 9 show the prompts used in the experiments.

Given a paper meta-review and the corresponding meta-review form, recommend appropriate Overall Assessment score that the meta-reviewer might consider giving based on their review.

Instructions:
1. Read the review text carefully.
2. Analyze the review and recommend an Overall Assessment score.
3. Provide a brief explanation for the recommended score.
4. Output a json dictionary: `{"overall_assessment_score": ..., "overall_assessment_justification": ...}`
5. Output only the json dictionary and follow the json schema exactly, no extra keys in the output.

Review form:

Overall Assessment
5 = The paper is largely complete and there are no clear points of revision
4 = There are minor points that may be revised
3 = There are major points that may be revised
2 = The paper would need significant revisions to reach a publishable state
1 = Even after revisions, the paper is not likely to be publishable at an *ACL venue

Table 8: The prompt used for the meta-review score prediction task.

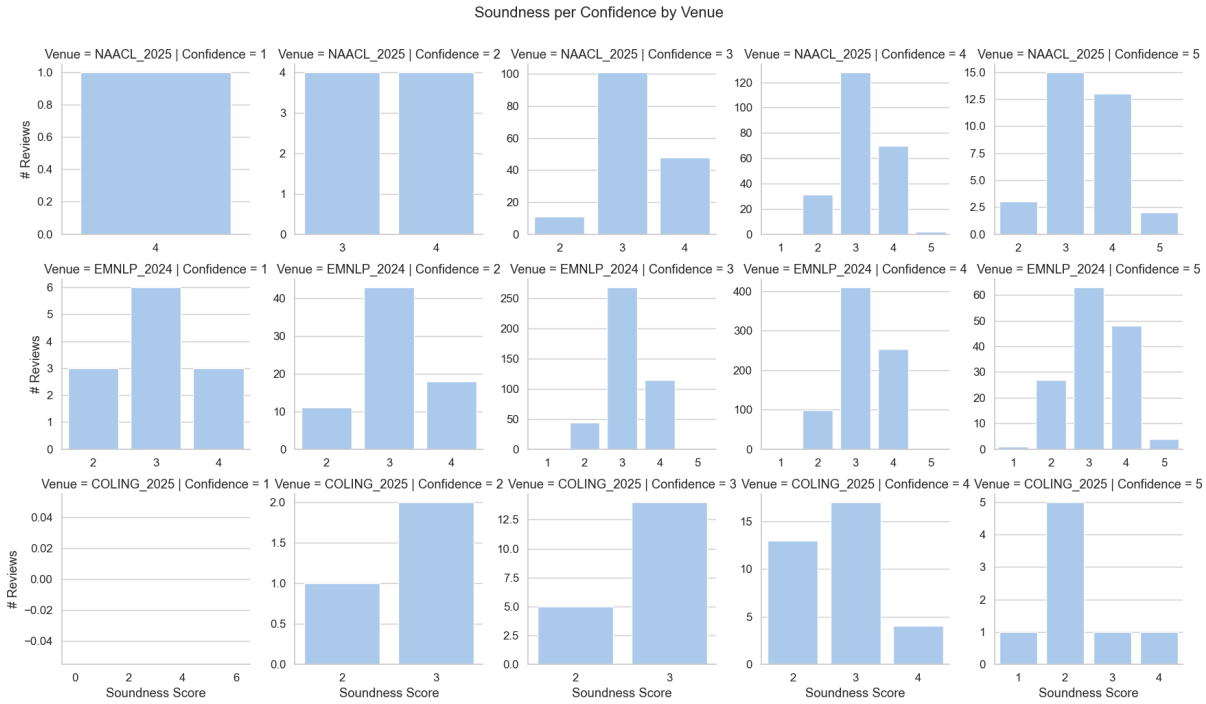


Figure 9: The bars show the average scores for *reviews*, with respect to *confidence* and *soundness* for the papers in our data collection.

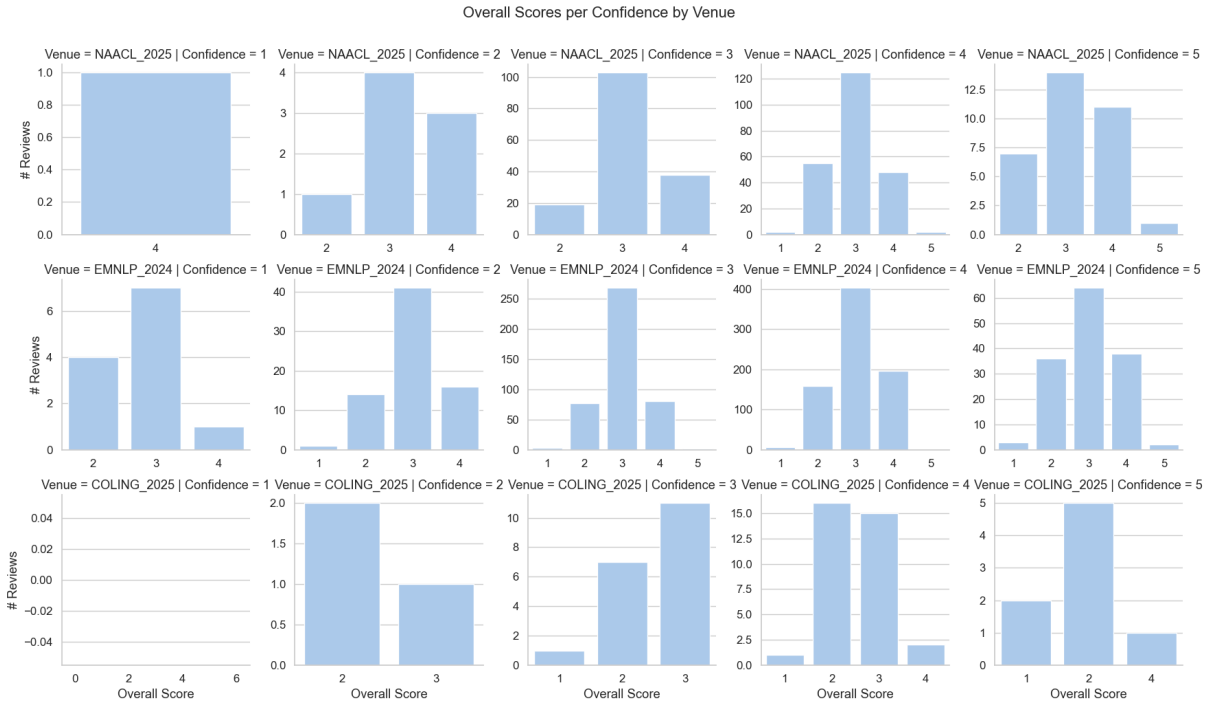


Figure 10: The bars show the average scores for *reviews*, with respect to *confidence* and *overall* score for the papers in our data collection.

D.2 More on Results

Table 10 shows the positivity bias results for COLING 2025. Table 11 shows examples of human-assigned and LLM-predicted review scores with large discrepancies.

Given a paper review and the corresponding review form, recommend appropriate Soundness and Overall Assessment scores that the reviewer might consider giving based on their review.

Instructions:
1. Read the review text carefully.
2. Analyze the review and recommend a Soundness score and an Overall Assessment score.
3. Provide a brief explanation for each recommended score.
4. Output a json dictionary: `{"soundness_score": ..., "soundness_justification": ..., "overall_assessment_score": ..., "overall_assessment_justification": ...}`
5. Output only the json dictionary and follow the json schema exactly, no extra keys in the output.

Review form:

Soundness
How sound and thorough is this study? Does the paper clearly state scientific claims and provide adequate support for them? For experimental papers: consider the depth and/or breadth of the research questions investigated, technical soundness of experiments, methodological validity of evaluation. For position papers, surveys: consider the current state of the field is adequately represented, and main counter-arguments acknowledged. For resource papers: consider the data collection methodology, resulting data & the difference from existing resources are described in sufficient detail. Please adjust your baseline to account for the length of the paper.

5 = Excellent: This study is one of the most thorough I have seen, given its type.
4.5
4 = Strong: This study provides sufficient support for all of its claims/arguments. Some extra experiments could be nice, but not essential.
3.5
3 = Acceptable: This study provides sufficient support for its major claims/arguments. Some minor points may need extra support or details.
2.5
2 = Poor: Some of the main claims/arguments are not sufficiently supported. There are major technical/methodological problems.
1.5
1 = Major Issues: This study is not yet sufficiently thorough to warrant publication or is not relevant to ACL.

Overall Assessment
Would you personally like to see this paper presented at an *ACL event that invites submissions on this topic? For example, you may feel that a paper should be presented if its contributions would be useful to its target audience, deepen the understanding of a given topic, or help establish cross-disciplinary connections. Note: Even high-scoring papers can be in need of minor changes (e.g. typos, non-core missing refs, etc.).

5 = Top-Notch: This is one of the best papers I read recently, of great interest for the (broad or narrow) sub-communities that might build on it
4.5
4 = This paper represents solid work, and is of significant interest for the (broad or narrow) sub-communities that might build on it
3.5
3 = Good: This paper makes a reasonable contribution, and might be of interest for some (broad or narrow) sub-communities, possibly with minor revisions
2.5
2 = Revisions Needed: This paper has some merit, but also significant flaws, and needs work before it would be of interest to the community
1.5
1 = Major Revisions Needed: This paper has significant flaws, and needs substantial work before it would be of interest to the community
0 = This paper is not relevant to the *ACL community (for example, is in no way related to natural language processing)

Table 9: The prompt used for the individual review score prediction.

input	field	model	+%	-%
meta-review	Overall Assessment	R1-Qwen-14B	61.76	5.88
		Qwen3-14B	47.06	2.94
		Qwen3-32B	50.00	2.94
		LLaMA3-3B	75.00	0.00
		LLaMA3-8B	65.38	3.85
review	Soundness	R1-Qwen-14B	73.44	7.81
		Qwen3-14B	48.44	17.19
		Qwen3-32B	62.50	6.25
		LLaMA3-3B	68.52	7.41
		LLaMA3-8B	60.00	0.00
	Overall Assessment	R1-Qwen-14B	73.44	6.25
		Qwen3-14B	59.38	14.06
		Qwen3-32B	64.06	15.62
		LLaMA3-3B	64.81	12.96
		LLaMA3-8B	80.00	20.00
review+rebuttal	Soundness	R1-Qwen-14B	84.38	0.00
		Qwen3-14B	75.00	3.12
		Qwen3-32B	68.75	3.12
		LLaMA3-3B	90.00	2.00
		LLaMA3-8B	81.25	6.25
	Overall Assessment	R1-Qwen-14B	87.50	0.00
		Qwen3-14B	71.88	14.06
		Qwen3-32B	70.31	10.94
		LLaMA3-3B	84.00	2.00
		LLaMA3-8B	87.50	6.25

Table 10: Proportions of outputs where the LLM predicted score is higher (**+%**) or lower (**-%**) than the human-assigned score for COLING 2025.

Human review	Human		LLM		LLM justifications
	S	OA	S	OA	
Strengths: The paper addresses formatting issues of web mathematics datasets and proposes a reasonable pipeline. **Weaknesses:** Missing details on key components and no validation across domains.	2	1	3	3	**Soundness:** Reasonable approach but lacks detail and cross-domain validation. **Overall:** Addresses a specific problem but limited generalizability; minor revisions could improve impact.
Strengths: Well-organized and smoothly written manuscript; dataset focuses on complex scenarios. **Weaknesses:** None.	2.5	3	5	5	**Soundness:** Thorough and sound research with clear writing. **Overall:** High-quality, valuable dataset of broad interest—top-notch contribution.
Strengths: K2Q increases diversity and complexity of VQA questions, improving robustness. Covers multiple domains and allows flexible expansion. **Weaknesses:** Some unnatural questions; domain coverage limitations.	3	2	4	4	**Soundness:** Technically sound dataset generation method with clear claims; some refinement needed. **Overall:** Significant contribution to dataset creation; limitations do not detract from value.
Strengths: Multidimensional dataset for summarization evaluation; detailed annotation; systematic LLM comparison. **Weaknesses:** Missing automatic metrics and slightly imbalanced data sources.	2	1.5	3.5	3.5	**Soundness:** Comprehensive description and evaluation, but missing key metrics. **Overall:** Useful dataset with room for improvement in size and balance.
Strengths: Clearly written; effective experiments. **Weaknesses:** Limited novelty; assumptions not validated; mechanism unclear.	2	1	3	3	**Soundness:** Clear problem and solid experiments, but lacks novelty and validation depth. **Overall:** Reasonable contribution needing minor revisions to strengthen claims.
Strengths: Systematic study on temperature parameter's effect. **Weaknesses:** Obvious findings; ignores other parameters (e.g., top-p).	1	1	3	3	**Soundness:** Supports main claims but lacks depth on related parameters. **Overall:** Reasonable contribution, though limited in novelty and breadth.
Strengths: Investigates stopword and lemmatization effects. **Weaknesses:** Lacks dataset details and regex explanation.	1	1	3	3	**Soundness:** Clear claim but insufficient methodological detail. **Overall:** Modest contribution; needs clarification for better reproducibility.
Strengths: Includes ablation and clear visualizations. **Weaknesses:** Lacks comparison with post-2022 work and more challenging datasets.	2.5	1	3.5	3	**Soundness:** Adequate support for claims, but limited scope and depth. **Overall:** Reasonable contribution, but relevance could be improved with newer benchmarks.

Table 11: Examples of human-assigned and LLM-predicted review scores with large discrepancies. S is for Soundness score, and OA is for Overall Assessment score.