

ACL Standing Committee on Peer Review (COPR) Report 1

Committee members:

- * Roy Schwartz (co-chair)
- * Anna Rogers (co-chair)
- * Steven Bethard
- * Roi Reichart
- * Yuki Arase
- * Vivian Chen
- * Barbara Plank
- * Tamar Solorio
- * Noah A. Smith

Abstract

The ACL peer-review committee has been working in the past few months to devise ideas on improving the peer-review process in *ACL conferences. We have started by running a survey on various ideas for the community.¹ We compiled the survey answers and based on them, our first actionable recommendation is moving from 6 ARR cycles per year to 5 cycles. We describe the motivation, the details, and a discussion of our proposal. We also discuss another recommendation whose exact implementation still needs to be worked out: allowing ARR to make Findings decisions.

Background and Motivation

ARR has converged over the past couple of years on 6 annual cycles, each starting on the 15th of even months. Each 2-month cycle includes multiple steps: *initial screening of the papers and assignment to reviewers; the actual review; chasing of late reviewers; a rebuttal phase; reviewer discussion; meta-review writing; and chasing late ACs*. At the meeting of ACL exec of July 17th 2024, the newly formed ACL Standing Committee on Peer Review (COPR) was asked to consider one of the common complaints about the current ARR workflow: the frustrating experience for the authors who dislike the decoupling of reviews and conference decisions. The problem was formulated as follows: **"ARR should make acceptance decisions together with PC chairs directly rather than running a separate commitment process."** We interpreted this problem as having two possible directions: somehow providing decisions in each ARR

¹ An analysis of the survey results is provided below.

cycle, or reducing the number of ARR cycles to the number of decision-making venues, which would remove the current decoupling.

COPR has discussed the problem at length over email and two meetings. The original plan was to produce a recommendation by ACL'24, but there was much disagreement in the committee on even the interpretation of the problem and the possible ways forward (6 options were discussed in total). The variety of perspectives and personal experience even within the committee was so great that we concluded that any such broad recommendations would need to be based on community input. We hence sought community input in a [survey](#) that was emailed at the end of October 2024, to the ACL membership and through ARR, as well as advertised on social media.

Besides the frustration with the decoupling process, originally posed to us, we considered the following input received from hands-on experience of the committee members with recent ARR cycles in various roles (EiC, SAC, AC):

- The increased volume of papers in recent cycles has made this job unsustainable. In particular, the current setup severely hurts editors and meta-reviewers (ACs and SACs), who need to chase authors/ACs in very short time windows, and are thus highly stressed out. It also hurts authors, who have little time for discussion with reviewers, and are thus often highly frustrated. On the EiC side, the job is becoming impossible, with the community burned out and too many reviews and meta-reviews coming in late.
- The different tasks described above start on random week days (as the cycle always starts on the 15th). This means that these tasks not only have too little time to complete, but also effectively *span weekends in many cases*.
- The back-to-back cycles do not allow for downtime for ARR that could be used for offline improvements.

The survey received 1600 responses, including 1040 from the current members of ACL. The analysis of the survey results is available in the appendix.

Proposed solutions

Recommendation 1: moving to 5 annual ARR cycles

As a final set of options for coupled vs decoupled process, we proposed in our survey an option to move to 4 cycles of 3 months each (which would fully remove the current decoupling of conferences and ARR). This proposal received support by reviewers, but authors prefer the current 6-cycle system.

As a result, we recommend a middle ground solution: **moving to 5 annual ARR cycles, each 10 weeks long**. This would increase the average cycle length by 10 days.

The extra days will be allocated by ARR so as to allow for more time in the current pain points, including chasing late reviewers, author-reviewer discussion, AC writing meta-reviews and SACs chasing late ACs. This solution also allows for starting each cycle, and each corresponding task **on Monday** to allow time to fulfill it during weekdays. Finally, it leaves two weeks of downtime (typically the final two weeks of the calendar year, during the December-January holiday) without any active ARR cycle, which can be used for improving ARR infrastructure.

A proposed schedule for 2025

We propose to start this change right after the Feb. 15 cycle (already dedicated to ACL 2025, and running on the original 2-month schedule). After that, ARR will run three more cycles in 2025: May 19th (dedicated to EMNLP); July 28th (dedicated to AACL); and Oct. 06 (TBD). The first 2026 cycle will start on Jan. 5th. **We have confirmed this schedule with EMNLP 2025 and AACL 2025 PCs, and they have accepted it.** The Oct. 2025 and Jan. 2026 cycles can be used for EACL and ACL 2026, but this remains to be discussed with the corresponding PCs. Below we outline the specific dates for these cycles, along with a tentative proposal on how to break down the first cycle into the different steps.

May 2025 Cycle: Mon 19 May 2025 through Sun 27 Jul 2025

Suggested schedule:

Mon 19 May 2025 through Sun 25 May 2025: PCs match SACs, ACs, and reviewers to papers

Mon 26 May 2025 through Sun 01 Jun 2025: ACs check papers and reviewer assignments

Mon 02 Jun 2025 through Fri 20 Jun 2025: Reviewers review papers

Mon 23 Jun 2025 through Sun 29 Jun 2025: ACs chase late and emergency reviews

Mon 30 Jun 2025 through Sun 06 Jul 2025: Author rebuttal

Mon 07 Jul 2025 through Fri 18 Jul 2025: ACs write meta-reviews

Mon 21 Jul 2025 through Sun 27 Jul 2025: SACs chase meta-reviews

July 2025 Cycle: Mon 28 Jul 2025 through Sun 05 Oct 2025

October 2025 Cycle: Mon 06 Oct 2025 through Sun 14 Dec 2025

Winter break (no active ARR cycle): Mon 15 Dec 2025 through Sun 04 Jan 2026

January 2026 Cycle: Mon 05 Jan 2026 Sun 15 Mar 2026

Discussion

Benefits

Our proposal will allow for more sustainable ARR cycles. We will have more time for ACs to chase authors and to write their meta-reviews. SACs will also have more time to chase late ACs. In addition, reducing the annual number of cycles by 17% (6 -> 5) will reduce the overall load on editorial members, who will be asked to serve in substantially fewer roles. Authors will benefit from this change indirectly (by having an overall improved review process), but also directly, as it will allow them longer discussion time with reviewers.

Downsides

The obvious downside of this proposal is having fewer options to submit within a given year. However, we believe this is a reasonable price to pay, as the current system is simply not sustainable.

Another minor downside to authors is the inconvenience of having to look up when the next submission deadline is, as opposed to "on the 15th every two months". To mitigate this minor inconvenience, the full schedule will be announced on ARR website and linked from ARR CFP.

Off-cycles

The proposed change will likely reduce the number of ARR off-cycles (i.e., those that are not linked to a specific conference). If all major NLP conferences (ACL, NAACL, EACL, AACL and EMNLP) adopt ARR, it will lead to having at most one off-cycle a year. While off-cycles are certainly valuable for some authors, who can submit their work early and enjoy another round of review, it is not an essential part of ARR and we believe it is a reasonable tradeoff to have fewer of them for the benefit of a more sustainable process. Reducing/removing off-cycles also directly addresses the problem this committee was formed to solve: the decoupling of ARR cycles and conferences.

Recommendation 2: Experiment with Making Findings Decisions in ARR

Having carefully considered multiple options for providing the authors with decisions at ARR level, given the still-decoupled 5-cycle process, the option that gained the most support in the committee was to allow ARR to make Findings decisions. The motivation comes from a frustration within the community of the discrepancy between ARR and the actual conferences. In particular, authors are surprised to receive a high score in their ARR meta review, but end up getting their paper rejected. To mitigate this, we propose shifting part of the conference decisions to ARR: **ARR will decide for each paper whether or not it is accepted to Findings**. In addition, only papers accepted to Findings can be then committed to a conference, which will make a binary decision: accept (to the conference) or reject (i.e., accept to Findings).

We sought feedback on this idea in the survey we released, and it received strong support from both authors and reviewers.

However, this idea is tricky to implement due to the following factors:

1. Currently, the only recommendation that happens at the end of an ARR cycle is the AC recommendation, i.e. the meta-review score. ARR meta-reviews are written by the ACs, who, despite ARR filters, are sometimes junior researchers (e.g., mid-term PhD students). Can we trust them with making such decisions? To better answer this question, we plan to run an experiment (possibly during the ACL or EMNLP 2025 cycles), in which ARR will explicitly ask ACs to make Findings recommendations, and then we will estimate the proportion of AC decisions overturned by the SACs. If this proportion is low enough, perhaps we could trust the ACs with making these decisions.
2. Following up on the previous item, currently about 35-40% of the papers in major conferences are accepted to either the main conference or Findings. If we allow ARR (and particularly ACs, who are responsible for a few dozen papers each) to make Findings decisions, we might find ourselves with much higher acceptance rates. We explicitly asked in our survey about the possibility of increasing the ARR acceptance rate as a result, and got mixed results. In particular, some of our survey respondents expressed concerns that a higher acceptance rate would create problems for them, even retroactively for older Findings publications, since their institutions rely on acceptance rates for assessing their performance.
3. Intuitively, this proposal would make the most sense if instead of Findings associated with different conferences (Findings of ACL/EACL/NAACL/AAACL), all these would become Findings of ACL. However, some survey respondents stated that Findings associated with conferences would be more beneficial for the reporting within their organizations, and also that it is important to offer Findings authors at least an option to somehow present their work at a conference (at least virtually, but onsite option would also be valuable for students, and, a TACL-style process would allow authors to choose the venue they can travel to). This requires further consideration and consultation with conference PCs.

Depending on the results of this experiment, this committee will consider the feasibility of this proposal and then provide a final recommendation on delegating the Findings decisions to ARR.

Summary of the October 2024 peer review survey

Roy Schwartz

Demographic questions

Overall, 82% of ~2000 responses came from current or past members of ACL, thus representing a mix of current community and people who would like to join it.² Most responses to [the survey](#) came from authors (90% submitted papers to *ACL in 2024), with 73% also serving as reviewers and 24% as ACs. 79% work in academia (including students), with 94% reporting that publications are an expected part of their job. 42% are students (PhD and below), 29% respondents in senior roles, and 28% in junior and postdoc roles. Demographically, most respondents are Asian (52%), of which 34% are east-Asians. Next are European (28%, of which 14% are Western European) and Americans (16%, almost all North American). 74% of respondents identify as men, 21% as women, 1% as non-binary and 4% preferred not to say.

Q1: Should ARR cycles and conferences remain decoupled (status quo)?

TL;DR: There is a clear preference for a decoupled process (status quo). It is mostly observed for authors, but also for reviewers. Meta-reviewers are mostly on the fence. It is also largely consistent across all dimensions we could consider (e.g., region, sector, venues, gender).

- **Authors** have a strong preference towards decoupling (60% vs. 33% for coupling, remaining 7% abstain). This is consistent across all dimensions we consider
- **Reviewers** have a weaker preference (43% vs 36%). Here the trend is not as consistent.
 - Reviewers with more publications (7+) have a preference towards decoupling, but there are naturally far fewer of those.
 - In general, senior people prefer decoupling. E.g., senior faculty; people who served in senior roles (AC, SAC, PC, workshop organizers),
 - Industry prefers decoupling
- **Meta reviewers** have a very weak preference for coupling (19% vs. 18%, 63% abstain). Here again we observe that senior people more strongly prefer decoupling.

Q2: What length and frequency of review cycles do you prefer?

TL;DR: a similar trend to Q1. Authors strongly prefer shorter cycles, but here reviewers prefer fewer, longer ones. Meta-reviewers mostly abstain.

- **Authors** strongly prefer the status quo (66% vs. 28%). This is consistent across all categories.

² See https://drive.google.com/file/d/1ZLinUkQCVAM3aC_id_VztqkbQ8FL9Jb-/view?usp=sharing for the raw survey results.

- **Reviewers** prefer longer cycles (46% vs. 37%). This is also mostly consistent, with a few exceptions:
 - Reviewers who review a lot (18+ papers per year) prefer the status quo.
 - Senior reviewers have an even stronger preference towards longer cycles, while students prefer the status quo.
- **Meta reviewers** mostly abstain (19% for shorter cycles vs. 18% for status quo, 63% abstain). This is mostly consistent, except senior people who, as before, have a strong preference towards longer cycles.

Q3: Should papers be considered by conferences in opt-in or opt-out mode?

TL;DR: there is a strong preference for the status-quo (Opt-in) across all categories (authors, reviewers, meta-reviewers), with increasing levels of abstention.

- **Authors** strongly prefer the status quo (59% vs. 32%). This is consistent across almost all categories, with people serving in senior roles showing more balanced opinions than those who haven't.
- **Reviewers** also prefer the status quo (46% vs. 25%). This is also mostly consistent, senior roles showing some preference towards opt-out.
- **Meta reviewers** mostly abstain (23% for opt-in vs. 13% for opt-out, 64% abstain). This is again mostly consistent.

Q4: Should ARR make Findings decisions?

TL;DR: there is a strong preference for accepting the proposal across all categories (authors, reviewers, meta-reviewers), with increasing levels of abstention.

- **Authors** strongly prefer the change (55% vs. 33%). This is mostly consistent across all categories.
- **Reviewers** also prefer the change (43% vs. 28%). This is again mostly consistent.
- **Meta reviewers** mostly abstain (21% for change vs. 14% for status-quo, 65% abstain). This is again mostly consistent, with people who actually meta-reviewed last year having more opinion, mostly towards accepting the proposal.

Q5: What kind of venue should ACL Findings be?

This question was hard to analyze given the multiple options to respond. Particularly, 7% marked both "I would publish in Findings in that situation, no matter what", and also another answer saying they wouldn't publish in Findings for some reason, which seems contradictory.

In general, it seems there is support for increasing the acceptance rate: 43% said they would publish in Findings no matter what, and only 33% said they wouldn't publish there if the rate was raised. Of the remaining options (that wouldn't publish in Findings any way), the leading reasons

were not being a first/second tier, lack of recognition by the institution, and lack of option to present at conferences (all marked by about 20% of the respondents).

Open comments:

We received 310 comments, which were all read by one of the COPR chairs. Below is Claude-assisted summary of the main points, chair-verified and edited. About 20% of Claude's points were irrelevant or incomprehensibly worded, and 15% more were added by the chair.

1. Review Quality Concerns

- Widespread concern about poor review quality and inconsistency
- Issues with reviewer expertise matching and qualifications
- Many reviewers not engaging during rebuttal/discussion periods
- Problems with emergency reviewers and late reviews
- Suggestions for reviewer training, recognition, and accountability
- Concerns about AI-generated reviews
- Concerns about low-quality reviews from conscripted authors
- Improving reviewer qualification criteria (currently it is possible to be a reviewer with 3 papers in ACL anthology, even if the reviewer is only a 3rd author on them)

2. ARR System Structure

- Many express a strong opposition to the decoupling system, finding that it adds too much complexity and not enough benefits
- Issues with timing between cycles (blackout periods, tight deadlines)
- Mixed opinions on 2-month cycle frequency (some prefer longer cycles)
- Workload concerns with minimum 4 reviews per cycle requirement
- Suggestions for more flexible reviewing options (1-2 papers instead of 4)
- Need for better reviewer-paper matching
- Need to consider reviews of resubmission towards the overall review load in the cycle
- Assumption of similar quality for papers associated with different conferences does not hold in practice, in some cycles it is easier to get high scores

3. Findings Venue Issues

- Varied institutional recognition and prestige concerns
- Requests for consistent presentation opportunities (virtual & in-person)
- Need for clearer differentiation between main conference and Findings
- Concerns about impact on workshops and other venues
- Mixed views on increasing acceptance rates to 50-60%

4. Process Transparency

- Unclear decision-making process for main conference vs. Findings acceptance
- Lack of transparency in meta-review decisions
- Need for clearer guidelines and expectations

5. Field-wide Concerns

- Growing submission volume creating strain on review system
- Need to maintain venue prestige while accommodating growth
- Impact on career advancement and institutional recognition
- Balance between fast science and quality

6. Suggestions for Improvement

- Consider desk rejections for clearly unsuitable papers
- Better mechanisms for handling resubmissions
- More recognition/incentives for reviewers
- Integration with anti-plagiarism systems