Cover Page

RLC Review Instructions 2025

RLC is using the reviewing process introduced last year for RLC 2024 **(described here)**[1] with **senior area chairs** (SACs), **area chairs** (ACs), and two types of reviewers: **technical** and **senior**. The purpose of this document is to make each of these roles clear and to describe how to write each type of review. It is worth stressing that this will be notably different from the typical NeurIPS/ICLR/ICML review process. We are doing things differently because the RLC PC believes the usual process is not working well.

At a high level, let's distinguish the two types of base reviews. A **technical review** is more or less a list of technical errors in the paper, in the form of a checklist. There is no commentary on novelty, significance, or impact. This is not an oversight on our part. It is very intentional. A **senior review** is a typical review, but much closer to what is expected of a TMLR paper review: focus on claims matching evidence but also including comments on technical correctness.

Note: *RLC does not have an author response to reviews. Instead, area chairs may ask the author's specific questions via OpenReview. Reviews will not be visible to authors until final decisions are released*.

With terms set, let's get into the details of technical reviews. It is important that both reviewer types read and understand the following two sections to get the full picture of what we are trying to implement. You can navigate the instructions for each role using the document tabs (typically on the left) or by clicking the following links.

Technical Reviewer Instructions: Link.
Senior Reviewer Instructions: Link.
Area Chair Instructions: Link.
Senior Area Chair Instructions: Link.
High-level description of the RLC Review Process: Link.

For additional great in-depth resources on reviewing, see these resources:

- Daniel Dennet, Criticising with Kindness.
- Comprehensive advice: Mistakes Reviewers Make
- Views from multiple reviewers: Last minute reviewing advice
- Sample excellent reviews from ICLR 2020 reviewer instructions
- Rich Sutton's guide for writing good reviews
- The "Avoiding Common Errors" section of the JMLR instructions.
- Empirical Design in Reinforcement Learning

[1] Note last year Area Chairs were called SACs due to limitations in OpenReview. The SAC role this year is new to the RLC process and very similar to the SAC role at ICML/NeurIPs/ICLR.

Technical Reviewer Instructions

Technical Reviewer Instructions

The format of the review is simple: (1) fill out a simple checklist answering questions about minor and major errors**,** and (2) an optional free-form textbox to note any other thoughts on the technical excellence of the paper (or flaws). Note: **Major errors** are things that would jeopardize the paper's ability to support its claims. Every paper should be judged mainly on the contributions stated on the cover page alongside their contextualization. This is why it is also so important to explicitly list the results (plots, theorems, analyses) that support each contribution on the cover page when filling in the free-form textbook. Major errors include things such as (but not limited to):

- **Math errors:** vacuous theorem statements, flaws in a proof, etc
- **Overclaiming:** like claiming false things about a new method that is not supported by the theory or empirical evidence given
    - Example: claiming a new method is significantly better than another based on invalid statistics, such as learning curves averaged over 5 seeds with overlapping shaded regions (and there is no other statistically rigorous evidence provided)
    - Example: A new TD learning method is developed with the hopes that it leads to learning richer feature sets with neural networks. It is shown that it achieves higher performance, but there is no measure of what representations are being learned. Any claim regarding the method's representation learning ability would be unsupported by the experiments.
- **Missing literature:** that would have impacted the development of the ideas in the paper.
- Related, this **idea has been done before**, and the paper's treatment is nearly identical to prior work.
- **Collection of minor errors** that, when taken together, make it very difficult to understand the purpose of the paper: poor polish, organization, inconsistent notation, undefined terms, etc. Things that on their own are minor but when co-occurring in mass can make a paper difficult to follow.
- **Unclear problem statement:** we do not know what problem is being solved or the paper did not motivate the problem, and it's unclear why it matters.

Below, we list a few common flaws in empirical work. Many of these can be seen as **minor errors** but do **become major errors when combined with overclaiming**. However, since the empirical practices in RL are often substandard, we will specifically name a few common errors that should be noted in a technical review:

- **Missing information** like the number of trials (i.e., seeds), the definition of shaded regions / error bars, not describing how hyperparameters were chosen and why.
- **Using measures of variance as a measure of confidence**: plotting standard deviation is not the same as a confidence interval

- **Claims requiring statistical justification that is not provided**: When claiming that method A is significantly better than method B, such a claim must be paired with an appropriate statistical justification (e.g., based on confidence intervals). Some helpful questions to ask: Were any statistical tests performed? Are they the appropriate ones? Are they sound? **Note that 3 trials is often insufficient for statistical claims.**
- **Lack of justification for chosen statistical analysis**: using standard error without justification (assumes normality and known variance), using IQM (interquartile mean) without justifying why top and bottom percentiles should be excluded
- **Missing baselines:** relevant baselines that would help us understand the new method (e.g., if claiming a new continuous action method works well, SAC or PPO are likely a useful baseline). Avoid asking for all the baselines. This is not needed and often just makes things less clear
- **Misrepresenting baselines**: using default hyperparameters of agent X from the literature on an MDP that agent X has never been applied to. This is appropriate if just to demonstrate X must be tuned for this new MDP, but it is not evidence that X performs poorly on this new MDP.
- **Tuning the hyperparameters of the new method only**. All RL agents are sensitive to hypers. We know this is problem-dependent. Showing tuned performance is better than another untuned method is misleading.
- **Claiming SOTA.** This is almost impossible to establish. It requires more trials than you think, careful tuning of hypers (very expensive), valid confidence intervals (likely studentized bootstrap). Rarely informative. Rarely true.

The overall goal of the technical review is to identify errors that would suggest the claims of the paper don't match the evidence. To determine this you must weigh the evidence provided and determine if it is valid. One run (seed) could be enough to support the right claim. You must evaluate each experiment and the statistical evidence provided, not appeal to what other accepted papers do or do not do!

Let's move on to minor errors. These are things that can be corrected and do not jeopardize the main message of the paper:

- Grammar, spelling, and general polish
- Plots: axes properly labeled, the measure of spread properly specified, the number of independent runs stated?
- Missing citations that are not critical missing baselines or ideas that would change the work
- Undefined notation before it is used
- Minor math errors
- Missing details

Finally, it is worth listing a few criticisms that are overused in reviews and are largely invalid—**avoid putting these statements in your review!**

- “Not enough environments! The method may not be general”. True of all current RL algorithms. This request is largely nonsense. Maybe you mean you want to see if the method works with images? Did the authors state that as a goal? Why is that important?
- “Environments are too simple!” Is Mountain Car simple? Yes, but so is Atari. This is arbitrary taste-making. The paper should explain why each environment was used and what we learned
- “New method is not SOTA”. This is nearly impossible to show and only relevant if the paper states this as a goal
- “New method is too simple”. “New method only combines existing ideas”. Complexity does not equal contribution
- “No theoretical justification”. Theory is but one form of evidence and is not required in every paper.

The guidelines above are not strict rules, and it is possible that some accepted papers may violate these guidelines with good reason. As such, the primary requirement for a paper to be accepted is that it clearly synthesizes new knowledge and that knowledge is of sufficient interest to the RL community.

What to discuss in the free-form textbox?

- If the paper has any theoretical results, list the Theorems and Proofs you checked for correctness.
- List the main result(s) that back up each one of the main contributions on the cover page (a figure, a theorem, etc.)
- Are the assumptions in each proof clearly reflected in the theorem statement?
- Did the authors discuss whether the assumptions that proofs rely on are reasonable?

You can continue to navigate the instructions using the tabs (typically located on the left) or by returning to the cover page using this link.

Senior Reviewer Instructions

Senior Reviewer Instructions

The principles of the RLC should be reflected in the review process:

> "The RLC peer review process prioritizes rigorous methodology over perceived importance, aiming to foster scholarly discussions on both well-established and emerging topics in RL."

As such, the overall advice to senior reviewers is to **focus on three things:** (1) does the work provide evidence supporting the main contributions proposed by the authors—claims matched by evidence? **This includes technical correctness!!!** (2) Is there an audience for the work? (3) What knowledge gap did the paper address? The cover page in every RLC submission can be seen almost as a rubric for such an evaluation: is every claimed contribution properly contextualized, and does the work provide evidence to support each one of them?

To make this concise, we provide three main guidelines to remember when reviewing.

1. **Answer four key questions** for yourself, to make a decision to Accept or Reject:
    1. Does the paper clearly state its contributions?
    2. Is the approach well motivated, including being well-placed in the literature?
    3. Does the paper provide support for all of the claimed contributions? This includes determining if results, whether theoretical or empirical, are correct and if they are scientifically rigorous.
    4. Does the paper claim any minor contributions that are *not* listed on the cover page? If so, is support also provided for these claims? Major contributions claimed in the paper, must also be included on the cover page
2. **Organize your review as follows** (inspired by Sutton's guide for writing good reviews):
    1. Summarize what the paper claims to do/contribute. Be positive and generous.
    2. Clearly state your decision (accept or reject) with one or two key reasons for this choice.
    3. Provide supporting arguments for the reasons for the decision.
    4. Provide additional feedback to improve the paper. Make it clear that these points are here to help, and not necessarily part of your decision assessment.

Additional considerations:

- If the paper has theory and algorithms, is the theory related to the algorithms? Is it related to the actual implementation of the algorithm in the experiments? If not, is the paper abundantly clear about this?
- Claims of state-of-the-art across all environments are difficult to support (strong claims require strong evidence)

You can continue to navigate the instructions using the tabs (typically located on the left) or by returning to the cover page using this link.

Area Chair Instructions

Area Chair Instructions

The primary role of the ACs is to ensure high-quality reviews for their papers, find replacements as needed, start and facilitate discussion, and write the final decision and meta-review. All very similar to other ML conference

Primary Job: making sure reviewers submit their reviews on time

Review deadlines are not soft; the ACs big role is continually communicating with reviewers to make sure they do their jobs on time. You can use personal email, OpenReview messages, or whatever to communicate timelines and expectations with your assigned reviewers. In the event your reviewers don't submit reviews, they need to be replaced, as discussed below.

Dates and Deadlines:

- **Feb 21:** Submission deadline (anywhere on earth)
- **Feb 21 - Feb 26:** Paper bidding
- **Mar 10:** Finalize paper assignment
- **Mar 24:** Reviews due (2 weeks to review)
- **Mar 24 - Apr 7 (2 weeks):** [Review adjustment phase] AC asks SRs and TRs to adjust their reviews if needed; AC assigns emergency reviewers as needed
- **Apr 7 - Apr 14:** Reviewer discussion phase (authors may be contacted with questions and requests to update contributions listed on the cover page)
- **Apr 25:** Meta-reviews from ACs due
- **May 2:** Decision notification

Replacing Technical Reviewers

The TRs were given these instructions to check correctness. Look for signs that they did not read the instructions, for example, missing the common errors we detailed in that document.

If there are any issues, the first step is to message the TR (in Openreview) to improve the review. If they do not improve it sufficiently, then you will need to assign a new TR. You must also remove the old TR, since we do not want the authors to see these inadequate reviews, and they should not factor into the final decisions.

There is a pool of emergency technical reviewers (TRs). These TRs were not assigned papers in the first round, but they can be assigned papers in the next round. You can also request a TR

to be added to the system and/or to a paper by emailing Marlos C. Machado (machado@ualberta.ca or program-chairs@rl-conference.cc).

Replacing or Adding Senior Reviewers

You can follow the same process to add and remove SRs. To evaluate the quality of Senior Reviews, refer to the instructions for reviewers here.

Paper Discussion Phase

After you have one Technical Review and one Senior Review that you are happy with (after getting them improved or replaced as described above), the next step is reviewer discussion. Even if the reviewers agree (perhaps, several major errors in the paper), it is still good to discuss. Naturally, this step is critical where there is disagreement: perhaps you think the paper is a solid contribution based on an initial read, but the reviewers don't, or perhaps there is disagreement about the severity of the technical errors.

Use the discussion phase to seek alignment (though not always possible), but more importantly, you are looking to finalize the decision in your own mind and collect evidence for your meta-review.
Starting the discussion is the responsibility of the AC. Don't wait for the reviewers to do it for you!

Optional: Asking the authors questions

A paper might have Major Errors, such as an error in the theory, not reporting the number of trials for their results, etc. Some of these Major Errors might be resolved by asking the authors a clarifying question. It might be that a theoretical result just needed a few more details to verify correctness, which can be checked in a conversation with the authors. Experimental details can also be determined. In these cases, you can simply private message the authors using open review and ask clarifying questions. The authors have been told to expect this, but there is no traditional author response for RLC. Use this direct dialog with the authors at your discretion!

In some cases, however, a Major Error is not resolvable without more substantial updates. For example, if the authors used 3 trials, but realized that they needed more to make the claims they want to make, then this would require an updated paper with a re-review of the claims. We did explain to the authors to check these types of errors before submitting, so they were alerted to the importance of checking correctness. For these instances, the paper likely needs to be rejected.

Another reason for reaching out to the authors is the cover page. It might be that the reviewers consider the paper contributions to not match the contributions claimed on the cover page. In this context, the AC is encouraged to ask the authors to re-write (or re-contextualize) their contribution in light of the reviews they have received.

Writing the Meta-review and Finalizing the Decision

Unlike other conferences, we are trying a different review process described here (note last year Area Chairs were called SACs due to limitations in OpenReview). We are trying to have fewer people involved in the review process so it can scale as RLC grows. As a consequence, the AC role is slightly different from, say, ICML. **For RLC we expect ACs to read the paper and form an opinion on the paper.** The meta-review need not be a full-length review (you can do so if you want), but it should be more detailed and contentful than a typical ICML meta-review: don't just summarize what the SR and TR think; also add your conclusions about the work. Pay particular attention to the contributions claimed on the cover page, whether their scope is appropriate, and whether it is properly contextualized. This is something more senior researchers (like you!) are much better positioned to do.

Go through the following checklist in order to finalize your final decision and meta-review

1. Read the paper
2. Read the Technical Review(s) and Senior Review
3. If either review is problematic, discuss with them to improve the reviews or replace them
4. Once you have a solid TR and SR, start the discussion with them with the goal of reaching a final consensus on the paper. This is not always possible, and you can overrule your reviewers.
5. Ask the authors clarifying questions based on the above steps
6. Discuss any borderline or tricky papers with your assigned Senior Area Chair. Definitely discuss papers where you want to overrule the reviewers
7. Finalize decision and write a detailed meta-review

After you have submitted your meta-review, the SAC or the program chairs might ask you questions about a paper and, during final decisions, may ask you to update your meta-review.

Senior Area Chair Instructions

Senior Area Chair Instructions

The role of the SAC at RLC is exactly like other major ML venues. You want to support your ACs and make sure they are doing their jobs as described here. That means:

1. ensure ACs are checking review quality and dealing with bad reviews
2. ensure ACs are following the schedule and making their deadlines (see timeline below),
3. discussing with the ACs challenging cases where the AC is unsure what to do
4. liasoning with the Program Chairs
5. helping the ACs below you to troubleshoot and generally navigate OpenReview.

These are all the usual responsibilities of a Senior Area Chair. You have been selected for this role because you are personally known to at least one of the Program Chairs to have good judgement and be well aligned with the goals of RLC. Use that good judgment; make good decisions—we trust you!

Finally, as you well know, RLC has a different review process than other ML venues **(described here)**[2] . You are expected to understand and advocate for the RLC process and principles in your interactions with Area Chairs and Reviewers. You should also be aware that this year, RLC requires all papers to have a cover page that clearly states the paper contributions and properly contextualizes those. Two examples are provided (Policy Gradient Methods for RL with Function Approximation, and Human-level Control through Deep RL). Familiarize yourself with those examples and the instructions for the cover page provided in the submission template. Alignment is key, so **if you have questions and disagreements about the process please reach out to the Program Chairs ASAP**.

Dates and Deadlines:

- **Feb 21:** Submission deadline (anywhere on earth)
- **Feb 21 - Feb 26:** Paper bidding
- **Mar 10:** Finalize paper assignment
- **Mar 24:** Reviews due (2 weeks to review)
- **Mar 24 - Apr 7 (2 weeks):** [Review adjustment phase] AC asks SRs and TRs to adjust their reviews if needed; AC assigns emergency reviewers as needed
- **Apr 7 - Apr 14:** Reviewer discussion phase (authors may be contacted with questions and requests to update contributions listed on the cover page)
- **Apr 25:** Meta-reviews from ACs due
- **May 2:** Decision notification

[2] Note last year Area Chairs were called SACs due to limitations in OpenReview. The SAC role this year is new to the RLC process and very similar to the SAC role at ICML/NeurIPs/ICLR.

Review Process

Reinforcement Learning Conference: an exploration step in reviewing

ATTENTION! *This is the description of the 2024 RLC review process. This is provided to help 2025 reviewers better understand our design motivations. An important difference for the 2025 process is that SACs have been renamed Area Chairs and we have introduced a new Senior Area Chair role.*

RLC decisions came out recently; congratulations to all the accepted papers and thank you to everyone who submitted a paper. The aim of this post is to shed light on the new processes that we hope will result in a better research publishing and conference experience. Rather than choosing a particular target acceptance right, we chose to focus on technical rigor and checking “does this paper’s evidence match its claims". In addition, inspired by TMLR, senior area chairs were empowered to further evaluate if RLC represented an appropriate and non-trivial audience for the work (and thus judge novelty and positioning of the work as well). Finally, we aspired to uphold high scholarly standards, meaning clarity, writing, and polish were held to a high standard. Judgements on significance, focus on state-of-the-art claims and more generally “taste making” were discouraged.

How did this new approach work out? We are excited to report that we wound up with a 40% acceptance rate and 115 accepted papers, which is right in line with early versions of other ML conferences. Having said that, it’s a new process so we wanted to take this time to go over some of its interesting aspects, clear up some misconceptions, and discuss future improvements we are considering.

The Process

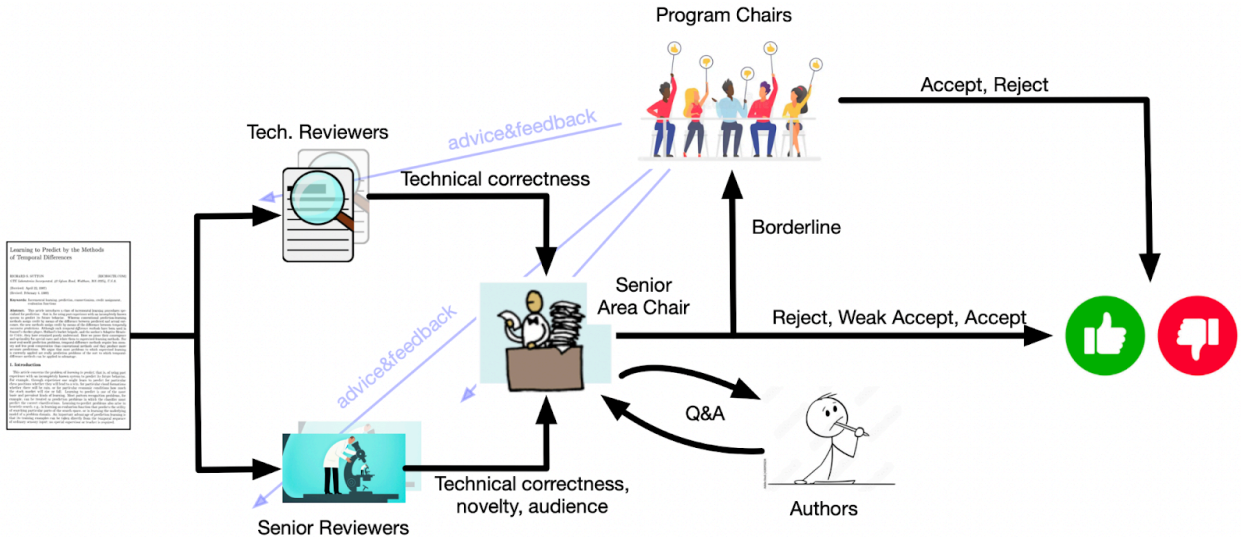

Figure 1: The purple arrows indicate the PC's influence over, but not control over, similar to the critic's feedback to an actor in [classic RL diagrams].

Our approach to improving reviewing was to reduce workload for reviewers, but ask for increased effort and quality from every reviewer per paper. In particular, we asked slightly less of our technical reviewers (who were typically more junior researchers, but not always), by having them only assess technical correctness, and relied on our senior reviewers to handle both technical correctness as well as more subtle factors like relevance to the field. The combination of these two reviews was then passed to a senior area chair (SAC) who synthesized the two, read and evaluated the paper (effectively a third reviewer), and then made a final call. These SACs we handpicked—senior folks with years of experience post PhD and known to be both generous to new ideas and have excellent judgment.

The program chairs were active at all stages of the process: giving comments and suggestions to improve both technical and senior reviews, discussing and troubleshooting issues with the SACs, and helping make final decisions about borderline papers. [In our system, if an SAC rated a paper as borderline they were communicating to the program chairs they needed help making the final call on the paper. All such cases were discussed by the PC.] Each PC member was assigned as a "buddy" to several SACs. The PC buddy provided a direct communication channel for questions, concerns and discussions about each paper, for each SAC. The full review process is summarized in the diagram above. Overall, each paper was read by four people and the decision was based on information from all four people.

The final decision process worked as follows. The SAC made the final decision based on (1) reading the paper, (2) the two reviews, (3) discussion with the reviewers and the PC, and (4) (optional) direct Q&A with the authors. If the decision was reject, weak accept, or accept that

decision was mapped to reject, accept, or accept respectively. If the SAC marked the paper as borderline, then the PC made the final decision. This is the process visualized in the figure above.

Note that at the scale we're operating at, we were able to handpick many of the reviewers, but still, many needed course correction. The program chairs, SACs, and many senior reviewers went to great lengths to identify and improve low-quality technical reviews. First, the technical reviewers were reminded of the reviewer guidelines and asked to improve their reviews. If the review was not updated satisfactorily, then a new technical reviewer was added. In several cases, the PC noted parts of technical reviews that were judged to be violating the spirit of the review process. We followed a similar process for attracting and improving senior reviews, including adding another senior reviewer in some cases. By large, this group was well known to the PC and organizers and the quality of these reviews was quite good. Like any conference, there were uncommunicative senior reviewers—and SACs—but nothing worth noting. Someone from the PC read every technical and senior review submitted to RLC.

This new process was an experiment—no system is perfect—but we believe this system is better for RLC than the standard approach. Our primary goals were to reduce workload, increase review quality, and focus on correctness and quality rather than taste. This was not easy and required a new process. Furthermore, this was a big change for our reviewers, who are accustomed to the typical review style that focuses primarily on the perception of significance and impact. It will take multiple iterations to fine-tune this system, but comments from the community suggest we are on the right path. We are excited about the papers that are accepted to RLC and looking forward to meeting all the authors. We will also continue to iterate and improve the system based on the results of this first experiment and feedback from the community.

What went well

We were happy with several aspects of this new review process. The feedback we received from folks in the community that our focus, on claims matching evidence and correctness, was seen as a breath of fresh air and a good way forward. Most importantly, we tried something very new and we now have data and comments to improve the process going forward.

The engagement and feedback provided to technical reviewers early on in the process by PCs, SACs, and SRs was effective. Many technical reviewers actually engaged and updated their reviews significantly. In cases where there was a lack of engagement, we assigned new reviewers from our pool of emergency reviewers, ensuring the review process continued smoothly. As mentioned, this required significant engagement from PCs throughout the process, which we will need to rethink to allow RLC to grow.

The detailed guidelines for TRs, SRs, and SACs were helpful to many. These guidelines provided instructions and expectations, helping to maintain consistency and quality across all

reviews. Naturally, these documents can be sharpened and consolidated, but they provided guidance for both reviewers and authors. Sharing these documents with authors before the submission deadline provided transparency and perhaps helped align submissions with the conference's goals.

The design of the review process was based on the principle of subtraction and appears to have had the desired effect. Typically, when things are not working well, people add new things: more reviewers, more tasks for reviewers and ACs, etc. This creates much more work, but often does not improve things. Subtraction, stripping things back to the barebones, can be a more effective alternative. This was the goal of introducing technical reviews, removing traditional author response, and reducing the number of reviewers overall. We believe the workload for technical reviewers and senior reviewers was much less than usual. For SACs, who typically must read reviews, discussions, author response, and long dialogs in OpenReview for up to five reviews per-paper, the workload was similar but now the SAC's time was put into reading the papers and writing more detailed meta reviews. Overall, we think workload was reduced and decision making accuracy was at least maintained and likely improved.

Finally, the introduction of PC buddies for SACs proved to be quite helpful. This system made it clear who the point of contact was for each SAC, facilitating better communication and support throughout the review process. We received constructive comments from several SACs on how to do this even better next year!

Looking forwards

There are a number of changes that we are contemplating in the next edition of RLC.

Continuing to improve reviewer guidelines

Many technical reviewers, instructed only to check for technical correctness, submitted almost one-line reviews of the form "this is correct." In many cases, the paper under review was excellent and this one-liner corresponded to an extremely thorough review! However, a paper author would not be able to tell if this review was thorough or not. In future years, we will likely have more detailed TR review instructions that make their work more visible.

Reviewing software and libraries

Software and libraries are a key part of RL in practice but require a slightly different review process. We're now starting the process of figuring out what a good, thorough review cycle would look like for software contributions.

Continuing to refine and grow our pool of reviewers

Initially, there were many low-quality technical reviews. We cannot be sure why. Some may have found the instructions confusing; some might not have dedicated enough time to the process; many clearly did not read the instructions and template matched to reviews they typically write. Others we know either lacked the ability to check correctness or refused to do so. Many reviewers who previously won reviewing rewards from major conferences produced some of the lowest quality reviews and never responded to comments or suggestions from the PC. Whatever the reason, we were quite surprised at the quality of some of the technical reviews; improving them required significant effort from the PC—and that will not scale as RLC grows.

We have to have an honest conversation with our community about mechanism design (rewarding good reviews) and training people on how to review papers. If folks do not take reviewing seriously and with an attitude of generosity towards the work, then the whole process breaks down. We can do better; we must do better!