

# How to Write Good Reviews for CVPR

Based on material by ICCV 2019 / CVPR 2019 Program Chairs

# Thank you for serving as a reviewer!

## We are all counting on you:

- **Area chairs** for clearly justified guidance for paper accept/reject decisions.
- **Authors** for fair consideration and constructive feedback.
- **Community** for ensuring that every conference paper teaches something worthwhile.

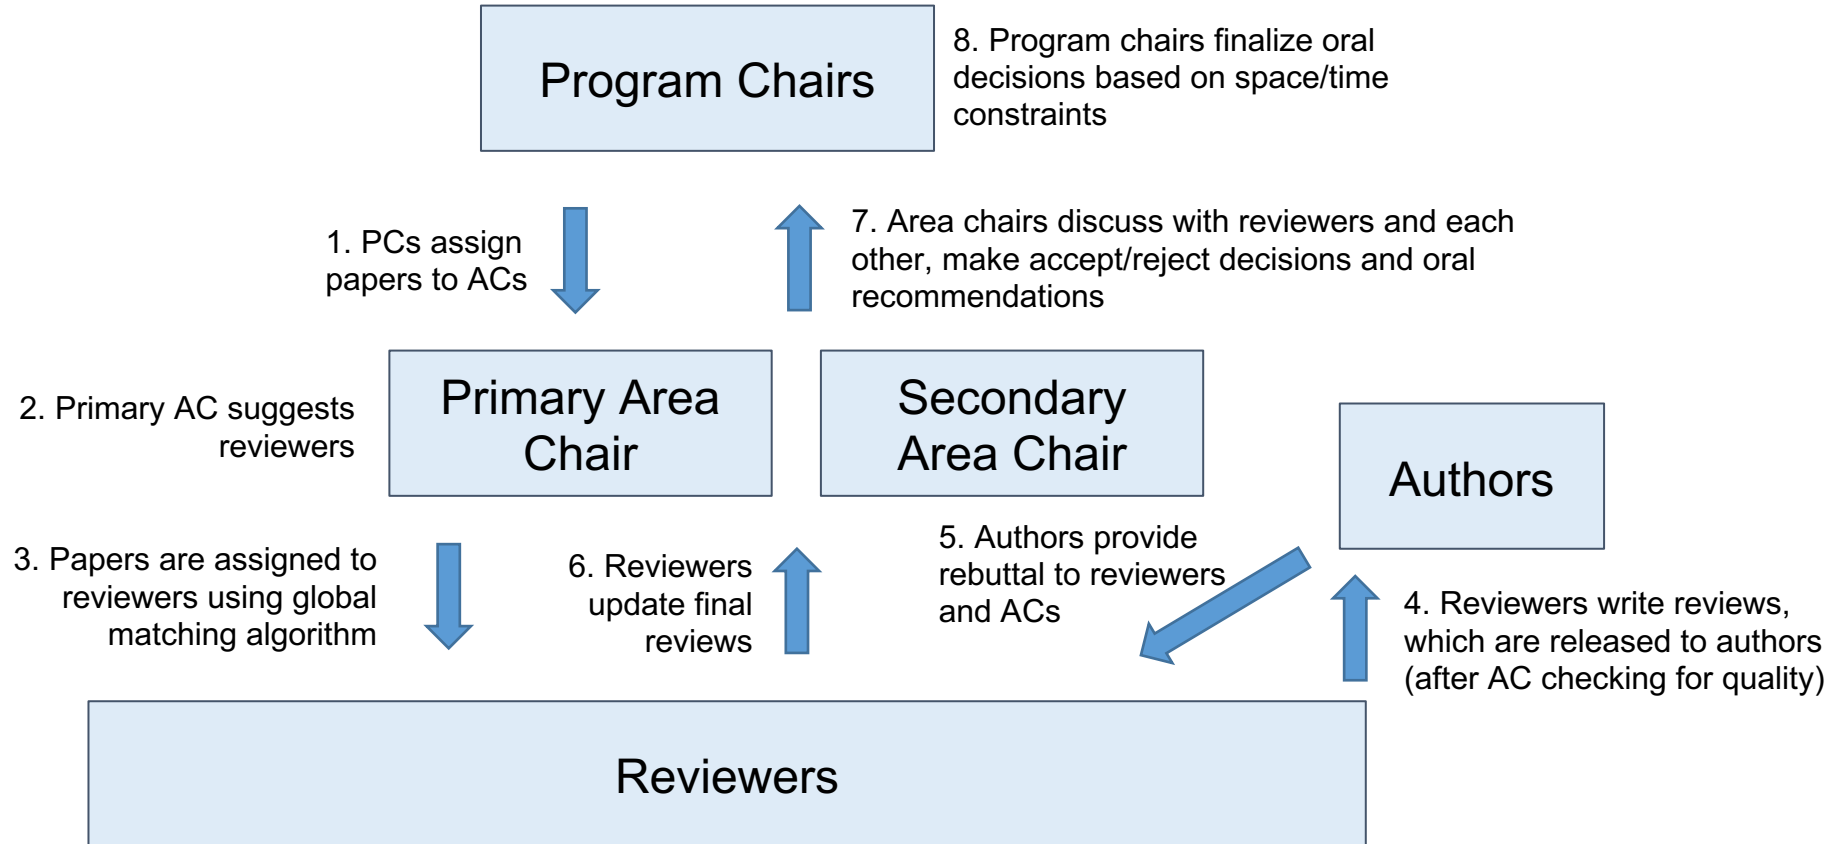
# **If you write bad, poorly justified, ill-considered, or unfair reviews.....**

- Area and Program Chairs, who may greatly influence your career advancement, may remember that you let them down.
- Authors may feel unwelcome or mistreated by the review process.
- A reader may waste time on a flawed or uninformative paper that was accepted, or may waste time in research because a valuable paper was rejected.

# **If you write good, insightful, well-justified, constructive reviews....**

- Area and Program Chairs will love you because you will make the paper decision much easier.
- The authors' faith in the vision community will increase, and, even if they need to resubmit, they will know what needs to improve.
- Researchers will continue to flock to vision conferences for the latest and greatest in computer vision ideas and techniques.

# The Decision Process: Overview



# The Decision Process: In Detail

1. **Program chairs (PCs) assign papers to area chairs (ACs)**, usually not more than 40 papers per AC.
2. **ACs suggest 10 reviewers per paper**, with help from Toronto Paper Matching Software (TPMS).
3. **Papers are assigned to reviewers** (3 per paper) using an optimization algorithm that takes into account AC suggestions, paper load and conflict constraints, and the prayers of PCs that nothing goes wrong.
4. **Reviewers submit initial reviews**, typically handling 6-10 papers each. ACs check quality of reviews, chase late reviewers, and assign emergency reviewers as necessary.
5. **Authors receive reviews** with a mixture of gasps, grimaces, grumbles, and the occasional grin. After much thought and re-reading of paper and reviews, authors submit rebuttals.
6. **Discussion ensues among reviewers and AC**, based on all reviews, rebuttal, and paper. Reviewers update their ratings and justification.
7. **ACs make decisions and write meta-reviews**. The decision and meta-review are recorded by the primary AC for each paper and checked/approved by the secondary AC. Primary and secondary ACs discuss borderline papers. Additional opinions may be sought from other expert ACs after checking for conflicts. In addition to accept/reject decisions, AC pairs provide a roughly ranked list of oral/spotlight nominations to the PCs.
8. **PCs make final determination of poster/spotlight/oral** for accepted papers, almost entirely based on the recommendations of the ACs but taking into account time and space constraints and topic diversity.

# Acceptance criteria

Your job as a reviewer is to provide well-reasoned recommendations to Area Chairs to enable them to make final decisions on all papers:

- **Award:** major advances that will heavily impact the field; will be used by many people, create new capabilities, etc.
  - E.g., ResNet (CVPR 2016 Best Paper), Mask R-CNN (ICCV 2017 Best Paper)
- **Oral:** potential to be very significant; worthwhile for the whole community to hear about.
- **Poster:** incremental steps that expand the sum of the community's knowledge or add bricks to the cathedral of knowledge; papers introducing useful tools; papers of interest to a subcommunity.
  - Also, creative ideas that are hard to judge but could be promising -- no one knows the future, so we should give the benefit of the doubt to *plausible* ideas.
- **Reject:** unlikely to be significant.

# Why not accept everything?

Papers can have a *negative* impact:

- Wrong or fraudulent results mislead the field and damage the reputation of the conference.
- Misleading evaluation makes it hard to compare with, kills follow-up.
- Creates bad precedent (weak paper X got in, so this one should too).
- Fatigue/overload of too many papers, wastes everyone's time.

Each weak or mediocre paper we accept hurts the conference a little (though not as much as rejecting a good paper).



# Review form outline

- **Summary:** Explain the key ideas, contributions, and their significance. This is your abstract of the paper. The summary helps the AC and the authors understand the rest of your review and be confident that *you* understand the paper.
- **Strengths:** What about the paper provides value -- interesting ideas that are experimentally validated, an insightful organization of related work, new tools, impressive results, something else? Most importantly, what can someone interested in the topic learn from the paper?
- **Weaknesses:** What detracts from the contributions? Does the paper lack controlled experiments to validate the contributions? Are there misleading claims or technical errors? Is it possible to understand (and ideally reproduce) the method and experimental setups by reading the paper?
- **Rating and Justification:** Carefully explain why the paper should be accepted or not. This section should make clear which of the strengths and weaknesses you consider most significant.
- **Additional comments:** minor suggestions, questions, corrections, etc. that can help the authors improve the paper, but are not crucial for the overall recommendation.

# New this year: code submission

- To improve reproducibility in AI research, we asked the authors to voluntarily submit their code as part of supplementary material
- We encourage (but do not require) you to check this code to ensure the paper's results are reproducible and trustworthy
- Use the [Reproducibility Checklist](#) as a guide for assessing whether a paper is reproducible or not.
- All code/data should be reviewed confidentially, kept private and deleted after the review process is complete.

# Guidelines

- **Take the time to do a good review**

- Many experienced reviewers take 2-4 hours per paper. If you're fairly new to reviewing (e.g. grad student), plan on at least 4 hours per paper and take the time to read the paper twice, consider related work, look up unfamiliar techniques, etc.

- **Be impartial**

- **Judge each paper on its own merits.** There is no global quota on the number of papers the conference can accept, and no requirement that the acceptance rate in your pile should match the acceptance rate of the conference.
- **Be aware of your own bias.** We all tend to assign more value to papers that are relevant to our own research. Try to ignore “interestingness of topic” or “fit to the conference” and focus on whether the paper can teach something new to an interested reader.
- **Try to discount the identity of the authors** if you happen to know it (e.g., through arXiv). If you *do not* already know who the authors are, *do not* attempt to discover them by searching arXiv.

# Guidelines (cont.)

- **Be specific and detailed**

- Your comments will be much more helpful to the ACs and the authors than your scores
- Do not simply give summary judgments (“not novel”, “unclear”, “incorrect”) – justify them in detail!
- This is particularly important for prior work. It is *not OK* to simply say “this has been done before”: you need to give specific references!

- **Be professional and courteous**

- Belittling, sarcastic, or overly harsh remarks have no place in the reviewing process.
- Avoid referring to the authors in the second person ("you"). Instead, use the third person ("the authors" or "the paper"). Referring to the authors as "you" can be perceived as being confrontational, even though you may not mean it this way.
- Do not give away your identity by asking the authors to cite several of your own papers.
- Proofread and spellcheck your reviews.

# Guidelines (cont.)

- **Be aware that different kinds of papers require different levels of evaluation**

- Potentially transformative idea: basic proof-of-concept.
- Established problem, plausible idea: benchmark results.
- Weird, overly complex, implausible, and/or seemingly incremental: extraordinary results (which need to be scrutinized carefully).
- Position piece or theory paper: no experiments.

# Ethics

- **Do not post anything online**

- Posting *any* information about the papers you are reviewing will result in severe consequences, e.g., revocation of submission privileges

- **Avoid conflicts of interest**

- Contact the Program Chairs if you suspect you may be conflicted with one of the authors (refer to Author Guidelines for detailed definition of conflicts).

- **Protect the authors' ideas**

- Do not show submissions to anyone else, including colleagues or students, unless you have asked them to write a review, or to help with your review.
- Do not use ideas/code from submissions you review to develop your own ideas.
- After the review process, destroy all copies of papers and supplementary material and erase any code you downloaded or wrote to evaluate the ideas in the papers.

# Examples of reviews

- The following examples are from ICLR, which published reviews in the public domain
- For ICLR, the review is written as a single narrative, rather than broken into sections as for CVPR/ICCV, but the same criteria apply
- Here we consider the **quality of the form**, rather than the accuracy of the content, of the review.

**Review quality: Good.** Though missing a summary of contribution, the review clearly explains why the paper should be accepted

(Note: this was a late-added review, which may account for brevity)

Rating: 9: Top 15% of accepted papers, strong accept

Review: First off, this paper was a delight to read. The authors develop an (actually) novel scheme for representing spherical data from the ground up, and test it on three wildly different empirical tasks: Spherical MNIST, 3D-object recognition, and atomization energies from molecular geometries. They achieve near state-of-the-art performance against other special-purpose networks that aren't nearly as general as their new framework. The paper was also exceptionally clear and well written.

The only con (which is more a suggestion than anything)--it would be nice if the authors compared the training time/# of parameters of their model versus the closest competitors for the latter two empirical examples. This can sometimes be an apples-to-oranges comparison, but it's nice to fully contextualize the comparative advantage of this new scheme over others. That is, does it perform as well and train just as fast? Does it need fewer parameters? etc.

I strongly endorse acceptance.

+ Clearly explains why the paper should be accepted  
– Does not contain many details about the contribution or why it is novel, so relies on the AC trusting the reviewer's judgment on these points

Note: though the proposed method does not achieve the best results (according to the review), the paper is highly valued for proposing a more general framework. Achieving best results is not necessary to validate the key idea (e.g., generality by testing with diverse datasets, or including an ablation study that isolates the impact of the key idea).

+ Indicates that the reviewer tried to think of weaknesses but could not come up with anything that should negatively impact the paper rating  
+ Constructive feedback for the authors



**Review Quality: OK but not great.** Makes general factors in decision clear and provides detailed feedback to authors, but does not provide adequate explanation for strengths and weaknesses

Rating: 8: Top 50% of accepted papers, clear accept

The paper proposes a framework for constructing spherical convolutional networks (ConvNets) based on a novel synthesis of several existing concepts. The goal is to detect patterns in spherical signals irrespective of how they are rotated on the sphere. The key is to make the convolutional architecture rotation equivariant.

Pros:

- + novel/original proposal justified both theoretically and empirically
- + well written, easy to follow
- + limited evaluation on a classification and regression task is suggestive of the proposed approach's potential
- + efficient implementation

Cons:

- related work, in particular the first paragraph, should compare and contrast with the closest extant work rather than merely list them
- evaluation is limited; granted this is the nature of the target domain

Presentation:

- \* While the paper is generally written well, the paper appears to conflate the definition of the convolutional and correlation operators? This point should be clarified in a revised manuscript.
- \* In Section 5 (Experiments), there are several references to  $S^2$ CNN. This naming of the proposed approach should be made clear earlier in the manuscript. As an aside, this appears a little confusing since convolution is performed first on  $S^2$  and then  $SO(3)$ .

Evaluation:

- \* What are the timings of the forward/backward pass and space considerations for the Spherical ConvNets presented in the evaluation section? Please provide specific numbers for the various tasks presented.
- \* How many layers (parameters) are used in the baselines in Table 2? If indeed there are much less parameters used in the proposed approach, this would strengthen the argument for the approach. On the other hand, was there an attempt to add additional layers to the proposed approach for the shape recognition experiment in Sec. 5.3 to improve performance?

Minor Points:

- some references are missing their source, e.g., Maslen 1998 and Kostolec, Rockmore, 2007, and Ravanbakhsh, et al. 2016.
- .... [abridged minor points due to lack of space in this slide]
- Figure 5, caption: "The red dot correspond to" -> "The red dot corresponds to"

Final remarks:

Based on the novelty of the approach, and the sufficient evaluation, I recommend the paper be accepted.

+ Highlights key ideas and contributions.

- The summary should also include one sentence on experimental setup

- Summary should include one sentence on significance of the contribution

+ Itemizes strengths and weaknesses

- Does not provide enough detail. E.g., what is original about the paper? How is the evaluation limited?

+ Includes clarifications questions and constructive feedback for authors

+ Makes it clear that "Minor Points" are not an important factor in decision

+ Identifies key positive factors in rating

- Would have been better to say why the weaknesses are given less weight



## Review quality: Bad

. The review lists only weaknesses and requests for clarification, omitting a summary and justification for decision. Thus, it is unclear to author or AC which of these points are the primary basis for the rating.

Rating: 4: Ok but not good enough - rejection

Review: 1. The idea of multi-level binarization is not new. The author may have a check at Section "Multiple binarizations" in [a] and Section 3.1 in [b].

The author should also have a discussion on these works.

2. For the second contribution, the authors claim "Temperature Adjustment" significantly improves the convergence speed. This argument is not well supported by the experiments.

I prefer to see two plots: one for BinaryNet and one for the proposed method. In these plot, testing accuracy v.s. the number of epoch (or time) should be shown. The total number of epochs in Table 2 does not tell anything.

3. Confusing in Table 2. In ResBinNet, why 1-, 2- and 3- level have the same size? Should more bits required by using higher level?

4. While the performance of the 1-bit system is not good, we can get very good results with 2 bits [a, c]. So, please also include [c] in the experimental comparison.

5. The proposed method can be trained end-to-end. However, a comparison with [b], which is a post-processing method, is still needed (see Question 1).

6. Could the authors also validate their proposed method on ImageNet? It is better to include GoogleNet and ResNet as well.

7. Could the authors make tables and figures in the experiment section large? It is hard to read in current size.

Reference

[a] How to Train a Compact Binary Neural Network with High Accuracy. AAAI 2017

[b] Network Sketching: Exploiting Binary Structure in Deep CNNs. CVPR 2017

[c] Trained Ternary Quantization. ICLR 2017

+ Cites papers that make the idea "not new"

- Does not say how these methods relate, so it is not clear if they are very similar techniques

- Because it is not tested by experiments, or that the convergence speed is not different?

- The remaining points may help authors improve the paper, but it is not clear if they are a significant factor in the rating to reject

### Big problems:

- AC can't make good use of the review without reading the paper, due to lack of summary/justification.
- No strengths listed, which may indicate that reviewer is just looking for reasons to reject.
- Author and AC don't know which of the listed points are important for reject rating.

# Take-away points

- **Respect authors and protect their ideas**
- **Take the time to do a good review**
- **Clearly justify your ratings**
- **Be constructive**
- **Do your work on time!**