


119 [responses](#)


Summary [See complete responses](#)

I am a




Student	46	39%
Professor	42	35%
Researcher (industry)	15	13%
Other	16	13%

Are NLP conferences (ACL, NAACL, EACL, and so forth) the main venues where you present your work?



Yes	103	87%
No	14	12%

What do you see as the primary goal of reviewing?




Identify any and all flaws in the experiments and interpretation of results. Papers with flaws should not appear in the conference.

Ensure that all recent related work is properly acknowledged and discussed, and compared with. No paper that seems similar to recently published work should appear in the conference.

Identify papers that contain interesting ideas and experiments that give clear analysis of a particular relevant research problem, and support their publication. Correction of any minor errors.

Other

When should reviewers suggest further work?



When the presented experiments are too preliminary to draw any interesting conclusions

98 82%

When there are obvious directions for further experimentation


6 5%

When the paper did not compare results with work the reviewer is familiar with

13 11%

Should the reviewer take into account the amount of work and effort that has gone into the experiments reported in the paper?

est their effort. [67]



Yes. Even if there :

Yes. Even if there are minor flaws and the results are not hugely superior to the state of the art, their effort should be considered.


48 40%

No. The authors should have chosen more wisely where to invest their effort.

67 56%

Is it part of a reviewer's job to make certain that portions of the work have not been previously published?

p if aware of it. [61]



Yes, reviewers sho

No, authors know ti

Yes, reviewers should make every possible effort to prevent this from happening.

46 39%

It is ok to point out the overlap if aware of it.


61 51%

No, authors know they are solicited to submit unpublished work and it is their decision to judge if the overlap will hurt their reputation.

9 8%

What is a reasonable reviewing load for conference reviewing?

r assignments. [68]



3-1 It is not good

9-6. This will ens

9-6. This will ensure that the same standards are applied to a larger portion of the papers. Reviewers can always select secondary reviewers if they need help with this load.

9

5-4 This will give reviewers a good sense of the general quality of the submissions but each reviewer can do all of his/her assignments.

61

3-1 It is not good to burden reviewers with too much papers; the lighter their load the better.

41

Should we look for ways to identify and reward good reviewers?


Yes, and good rev...

Yes, and good rev...

No, this is too s...

No, there are too...

No, writing good ...



Yes, and good reviewers should be identified by feedback from area chairs

65 55%

Yes, and good reviewers should be identified by feedback from the authors

25 21%

No, this is too subjective a question

11 9%

No, there are too many good reviewers to acknowledge

3 3%

No, writing good reviews is expected behavior, there is no need for special acknowledgements

13 11%


To what extent should different ACL conferences coordinate the review of YOUR papers that were previously rejected?

No coordination necessary; we will fix the issues that were mentioned. Having access to prior reviews will bias and confuse the reviewers.

53 45%

No coordination necessary. We need a fresh set of opinions; mistakes do happen.


41 34%



I would like to get the same reviewers as in the previous conference.
I will certainly hope that I get different reviewers.

8 7%
13 11%


Should area chairs be expected to read all the papers in their area and write a meta-review summarizing all reviews and discussions that have led to the final decision on the paper?



Yes, this would ensure that the chairs for each area have read all papers and can mitigate problems with reviewers' strictness or leniency.
No, this would be too much work.

47 39%
68 57%


How should reviewers be recruited?



Always provide area chairs with the complete list of reviewers from previous year. They may want to use it to balance geographic areas or expertise.
Recruit reviewers for the conference and use them across areas as needed, no need to commit to a particular area.

75 63%
37 31%

Should there be a minimum length of reviews?



Yes, the review should make it clear why a paper is good or bad, not just state the reviewer's opinion.
No, some papers are so clearly good or bad that not much needs to be said.


70 59%
45 38%

Should reviews contain explicit suggestions about what needs to be done in order for the work to become acceptable for publication?

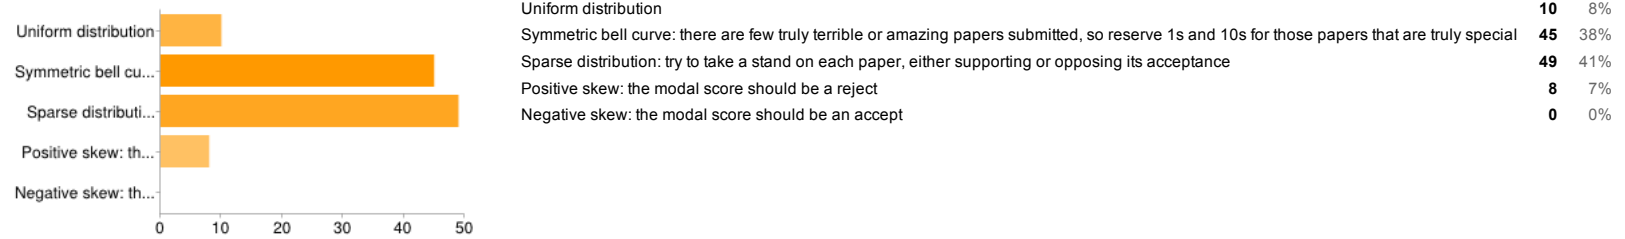


Yes, the purpose of reviewing is to ensure that all reasonable work is eventually published
No, it is the authors' problem to figure that out

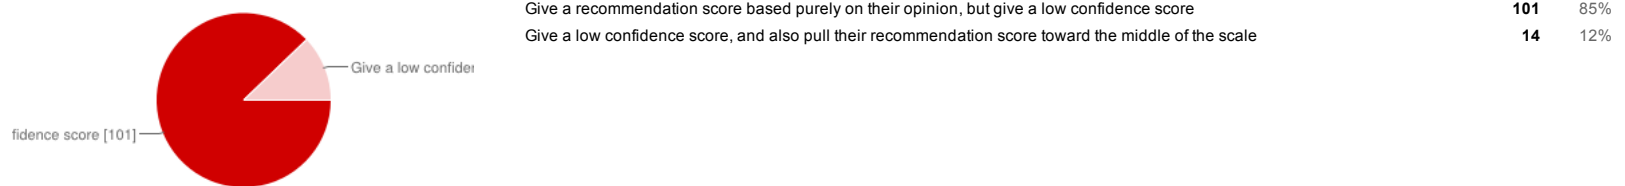
100 84%
14 12%



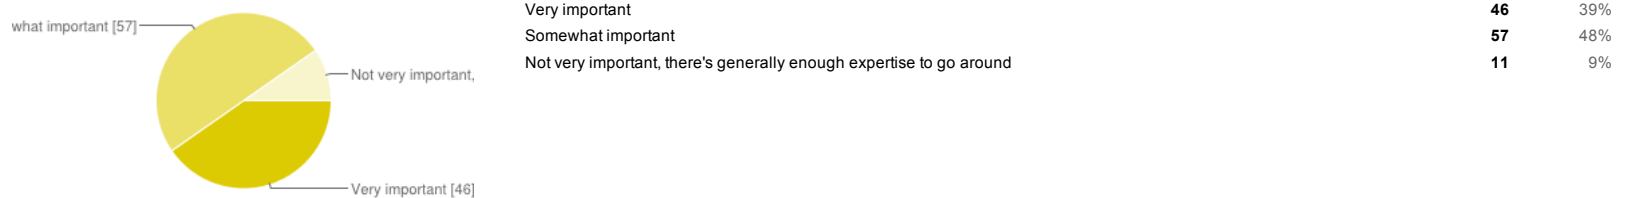
What overall distribution of recommendation scores (assuming a 1-10 scale) should reviewers aim for?



When a reviewer leans strongly towards either acceptance or rejection, but also feels that the paper is a bit out of their central expertise, which of the following courses of action is better?




How important is it for area chairs to easily be able to recruit reviewers outside of the general conference pool to handle specific papers based on expertise?




Author feedback periods have become more popular in the past few years in CL conferences. How useful do you feel that these are?

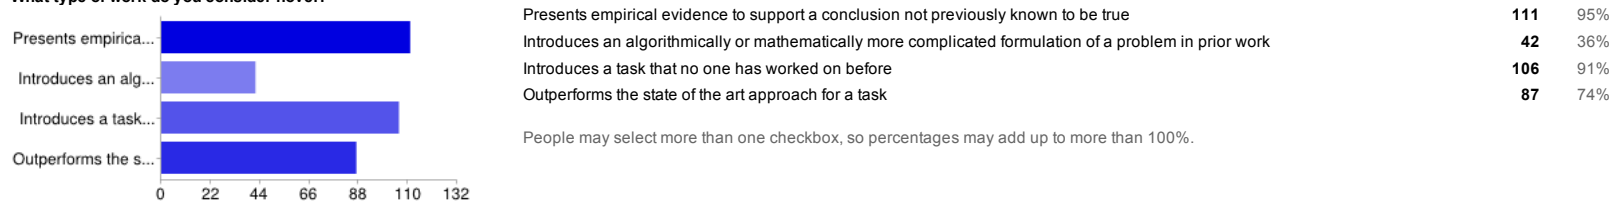
Usefulness	Count	Percentage
Very useful and worth the time: I have seen opinions swing from accept to reject and/or the reverse on the basis of author feedback, and the decisions are better for it	47	39%
Somewhat useful, but not worth the time	39	33%
Not useful: I have seen nothing that affects accept/reject decisions, or I have but the decisions were not any better overall	27	23%



Journal-style reviewing is different from conference-style reviewing in a number of respects, perhaps most notably the lack of a cap on the number of back-and-forth iterations between authors and reviewers. Compared with journal-style reviewing, how do you feel conference-style reviewing stacks up?



What type of work do you consider novel?



People may select more than one checkbox, so percentages may add up to more than 100%.

For submissions where empirical results are presented, how much value should be ascribed to those that include code and/or data, compared to submissions that do not:

