

Report on ACL Survey on Preprint Publishing and Reviewing

Jennifer Foster

Marti Hearst

Joakim Nivre

Shiqi Zhao

July 2017

Executive Summary

The ACL survey on preprint publishing and reviewing was conducted in June 2017 and received 623 complete responses. This section gives a high-level view of the main results and trends, based on the detailed results presented in subsequent sections.

The great majority of respondents (93%) are either current or previous members of ACL, and the sample represents about 20% of the total membership of the association. As far as we can tell from available membership statistics, the sample appears to be representative with respect to demographic factors like geographical affiliation, gender, role and academic background. (For more information, see Section 1.)

The first set of questions concerned the practice of uploading papers to preprint servers either before or after they have been submitted or accepted to conference or journals. The results show that just over half the respondents (53%) seldom or never use a preprint server to host their research papers, with just over a fifth (22%) doing so always or often. Authors who tend not to put their papers on a preprint server motivate this by not seeing the need when they intend to publish at a conference or in a journal and by wanting to preserve the integrity of double-blind reviewing. Of the authors who regularly put their research papers on a preprint server, 28% tend to upload them before notification and 43% wait until after the paper has been accepted. Those who upload before notification motivate this by wanting to publicize their research as soon as they think it is ready or by wanting to timestamp the ideas in the paper.

The next set of questions concerned the practice of reading and citing preprint papers. The results show that the great majority of survey respondents (86%) read preprint papers at least sometimes, while a smaller majority (54%) cite them at least sometimes. Authors who tend not to cite preprint papers almost unanimously motivate it by preferring to cite a

published version if available and express misgivings about the reliability of preprint papers due to the lack of peer review.

About 75% of the survey respondents had reviewed for ACL before. A large majority of them (95%) state that they search for related work as part of the review process at least some of the time, but only a minority (23%) have searched for a preprint version of a paper after having been informed by area chairs that one exists. A small minority of reviewers (14%) penalize a paper for not citing a very relevant preprint paper, and a slightly larger minority (20%) penalize a paper for not comparing to empirical results in a very relevant preprint paper. Around 40% in both cases make a comment but do not change the overall recommendation. The free text comments reveal that there are also reviewers who go in the opposite direction and penalize papers who cite preprint papers extensively.

The final set of questions concern future policies for reviewing and preprint papers at ACL conferences. When asked to weight the importance of double-blind reviewing and being able to submit preprint papers, the great majority of respondents (88%) consider double-blind reviewing at ACL conferences to be important. About 75% of those (65% of all respondents) consider double-blind reviewing more important than being able to submit preprints. Only a small minority (9%) consider preprint publishing more important than double-blind reviewing. Respondents who consider double-blind reviewing important motivate this primarily by saying that it helps maintain the quality of ACL papers and that it helps to prevent unconscious bias against authors from underrepresented groups.

When asked about their preferences for different reviewing models for ACL, the most popular choices are in support of double-blind reviewing, in combination with banning or discouraging preprint publication or keeping the status quo. Among those who advocate abandoning double-blind review, completely open reviewing has slightly stronger support than single-blind reviewing. In addition, the responses show support for lobbying preprint servers to allow (temporarily) anonymous publication (in order to preserve the integrity of double-blind reviewing), for providing guidelines for the citation of preprint papers (both for authors and reviewers), and for introducing journal-style reviewing with a rolling deadline. The idea of having a separate track for preprint papers received less support.

A final open question generated comments on the current situation, opinions on the suggestions for future policy provided above, new innovative suggestions for future policy, and comments on the survey itself. Taken together with the results presented above, the main views expressed can be summarized as follows:


- Many respondents express support for double-blind review as important for maintaining quality and preventing bias.
- Many respondents express support for allowing preprints to promote fast research dissemination and progress, and to prevent ACL alienating researchers from related fields if preprints are banned.
- Many respondents argue that the best way to resolve the conflict between preprint publishing and double-blind review is to convince preprint servers to allow temporary masking of author identity (and maintain the status quo in other respects).

- A substantial number of respondents highlight the need for guidelines (for both authors and reviewers) regarding the citation of preprints.
- A substantial number of respondents express concern about the current conference reviewing process with increasing reviewer load and decreasing review quality. Suggestions to come to terms with this problem include journal-style reviewing with rolling deadlines and making reviews public.

1. Demographics


In this section, we review the demographics of the survey respondents, covering region, gender, role, academic background and research area. Some of this information will be used as background variables in parts of the following analysis.

Region




The primary location of respondents was Europe/Africa/Middle East followed by the Americas and Asia. This can be compared to the percentage of members from the three regions in 2016: 39% in Europe/Africa/Middle East, 44% in North/Central/South America and 17% in Asia/Pacific.

Gender




The gender division was 22.3% female, 71.6% male, and 6% declining to state. We have no corresponding statistics for ACL members to compare with.

Role



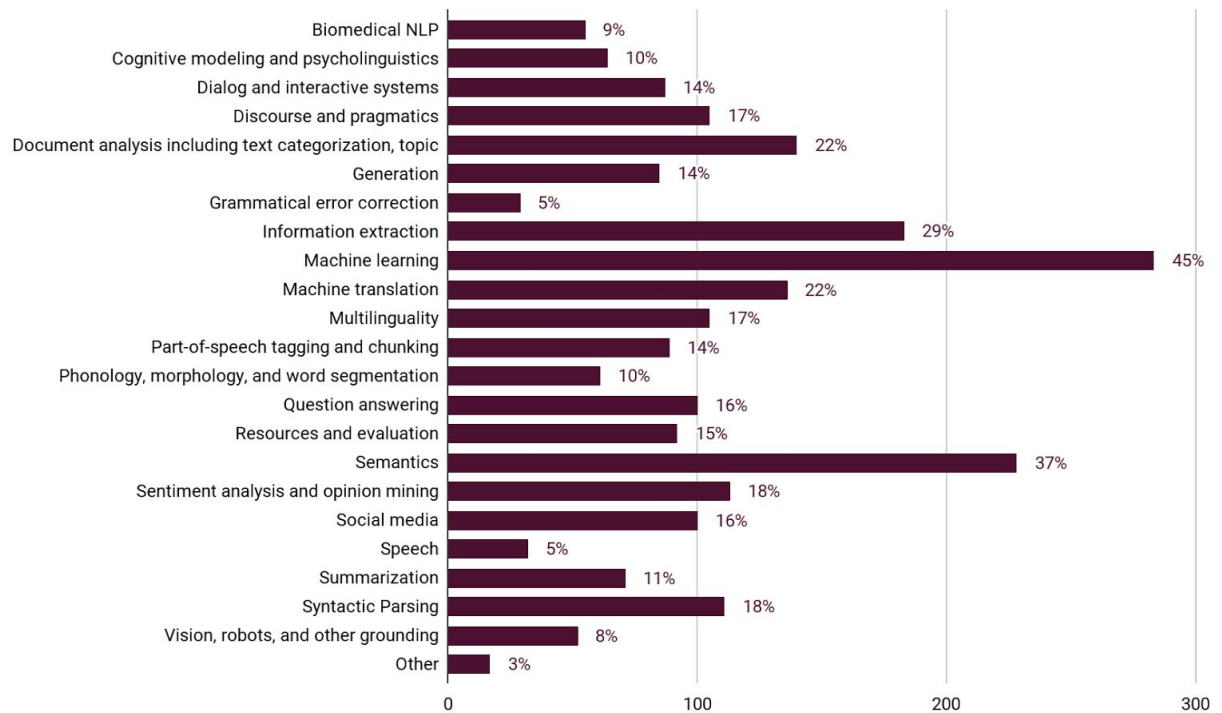
The top roles of respondents were graduate students and professors/lecturers, followed closely by researchers in academia. While academic researchers were in the majority, industrial researchers were also well represented.

What is your background?




Most respondents (88%) have a background in computer science, and almost a third of them (28%) have a background in linguistics. Backgrounds given under Other include Computational Linguistics (8), Cognitive Science (5) and Electrical Engineering (5).

What are your research areas within NLP?




Many different research areas are represented in the survey. The biggest areas are machine learning, semantics and information extraction, which mainly seems to reflect the popularity of these areas in the community.

How well known would your colleagues consider your research to be within your subfield?



The distribution of lesser vs. better known researchers follows an almost perfect normal distribution (with a slight skew towards lesser known).

How well known do you consider your institution to be?




By contrast, the distribution of lesser vs. better known institutions is heavily skewed towards better known. It thus seems that well-known institutions are overrepresented in the responses to the survey (unless there is a strong general tendency for people to overestimate the reputation of their own institution).

The great majority of respondents (93%) are either current or previous members of ACL.

2. Authors


This section contains the results for questions directed to authors, in particular their use of preprint servers like arXiv.

Do you upload research papers to preprint servers such as arXiv?




Just over half (53%) of respondents seldom or never use a preprint server to host their research papers, with just over a fifth (22%) doing so always or often. These results are broken down by gender, role, and geography in the following graphs.

Do you upload research papers to preprint servers such as arXiv?




16 female respondents said they always or often post to preprint servers compared to 112 male respondents. This is lower than the ratio of female to male survey respondents overall (139:446). By contrast, a third of all the respondents who never upload to preprint servers are female. Further analysis shows that the percentage of respondents who upload to preprint servers at least sometimes is 51% in the male subgroup and only 34% in the female subgroup.

Do you upload research papers to preprint servers such as arXiv?



Graduate students constitute two thirds of the group of respondents who always upload papers to preprints, despite being only about one third of the total number of respondents. Further analysis shows that, among graduate students, 56% upload their papers at least sometimes, which is significantly higher than for any other group. The group that is least likely to upload to preprint servers is that of researchers in academia, where only 36% do so at least sometimes. For all other groups, this percentage is 45%.


Do you upload research papers to preprint servers such as arXiv?



The same information by region shows that the responses pretty much align with the representation of each region in the survey response rate. In other words, geography does not seem to strongly predict whether or not someone uses preprint servers, although further analysis shows that the percentage of respondents who upload their papers at least sometimes is slightly higher for the Americas – 51% compared to 45% for Asia/Pacific and 44% for EMEA.


Do you upload research papers to preprint servers such as arXiv?

Self-reported Fame



There are no clear tendencies in this breakdown, except that researchers who do not consider themselves particularly well known are over-represented in the group that always upload their papers, which is consistent with many of them being graduate students. Interestingly, however, not particularly well known researchers are also over-represented in the group that never upload. Further analysis shows that the percentage of respondents who upload at least sometimes is slightly higher among researchers who consider themselves well known (but not very well known) – 51% compared to 41–46% in the other four groups.

Why do you tend not to upload your papers to a preprint server? [Multiple response]



Of the 333 respondents who tend not to put their papers on a preprint server, by far the most popular of the listed reasons were

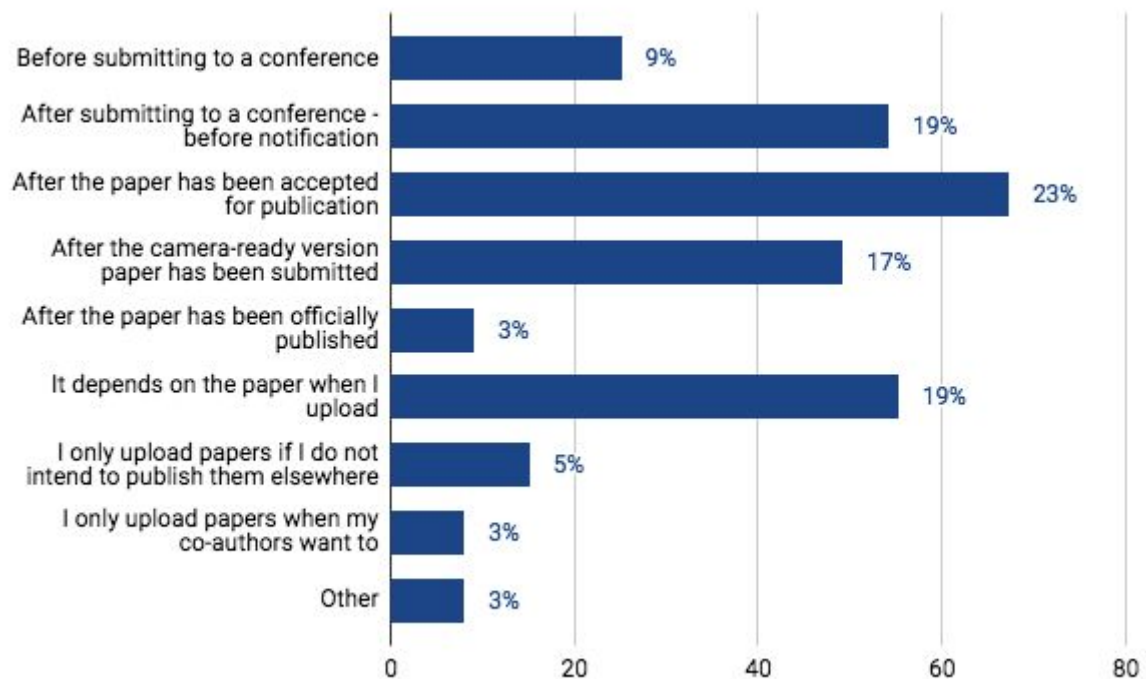
I do not see the need when I intend to publish my papers at a conference or in a journal

and

I want to preserve the integrity of double-blind reviewing


Reasons supplied by respondents in the *Other* category included not having had the chance yet to use a preprint server, not wanting to split citation counts between multiple versions of a paper, and not trusting the quality of preprint papers.

At what stage do you usually upload your papers?



Of the 290 respondents who put their research papers on a preprint server, 28% tend to upload them before notification and 43% wait until after the paper has been accepted. For just under one fifth of respondents (19%), the decision of when to upload a paper depends on the type of paper (a handful of respondents noted in the *Other* response category that they upload journal articles at submission time and conference papers after notification). Another recurring response in the *Other* category was the uploading of papers to a preprint server after they had been rejected from one or two conferences.

Why do you put a paper on a preprint server before it has been accepted for publication? [Multiple response]



Of the 134 respondents who sometimes put their papers on a preprint server before acceptance notification, 107 (80%) listed as a reason


To publicize my research as soon as I think it is ready

Another popular reason was

To timestamp the ideas in the paper


Reasons provided in the *Other* category include not wanting to wait until the next conference cycle after a conference rejection.

Do you update a paper's metadata once it has been accepted for publication?



Of the 134 respondents who sometimes put their papers on a preprint server before notification of acceptance, the majority update the papers' metadata once it has been accepted for publication at a conference or in a journal.


Why do you wait until after a paper has been accepted before putting it on a preprint server? [Multiple response]



Amongst the 125 respondents who wait until after a paper has been accepted for publication before uploading it to a preprint server, there was support for the three listed reasons, with the most popular being


I want to preserve the integrity of double-blind reviewing

Under what circumstances do you update the content of a preprint paper? [Multiple response]




Of the 290 respondents who put their research papers on a preprint server, 61 (21%) do not update the contents of their papers once uploaded. The most popular reason for updating the contents of a preprint paper is if an error is discovered in the paper. Most of the responses in the *Other* category refer to updating a preprint paper to match the camera-ready version of the paper.

Do you advertise your preprint papers on social media?




Of those respondents who regularly host their papers on a preprint server, 40% tend to use social media to advertise them, 40% do not, with the remaining 20% doing so sometimes.

Do you subscribe to a preprint news feed, e.g. the arXiv feed?




A small majority (57%) of survey respondents do not subscribe to a preprint server news feed. This is a similar proportion to those who don't upload papers to a preprint server.

How often do you read preprint papers?



The majority (86%) of survey respondents read preprint papers at least sometimes.


How often do you cite preprint papers?



Compared to the proportion of survey respondents who read preprint papers, a smaller majority (54%) cite them at least sometimes. 27% cite often or very often, while 45% cite seldom or never.

Further analysis shows that respondents who themselves upload papers to preprint servers at least sometimes are more likely to read and cite preprint papers than those who only upload seldom or not at all. In the former group, 74% often read preprint papers, 56% subscribe to a preprint news feed, and 47% often cite preprint papers. In the latter group, 41% often read preprint papers, 33% subscribe to a preprint news feed, and only 11% often cite preprint papers.

Why do you tend not to cite preprint papers? [Multiple response]



Of those 282 respondents who tend not to cite preprint papers, 93% mark as a reason


I would rather cite a published version if available

The majority of the 51 responses in the *Other* category state misgivings about the reliability of preprint papers due to the lack of peer review.

3. Reviewers

This section contains the results for questions directed to reviewers about how they handle preprint papers and citations of preprint papers.


Have you reviewed for an ACL conference or workshop?



Just over three quarters (76%) of survey respondents report that they have reviewed for an ACL conference or workshop.


All other questions in this section of the survey are shown only to those who have answered yes to this question.

When you are reviewing for ACL or related conferences/workshops, do you search for related work as part of the review process?




Of the 472 survey respondents who have reviewed for ACL before, a large majority (95%) search for related work as part of the review process at least some of the time.

Have you ever searched for a preprint version of a paper after an area chair has informed you that one is available online?




Of the 472 survey respondents who have reviewed for ACL before, the majority (77%) have never searched for a preprint version of a paper after having been informed by area chairs that one exists. 108 respondents (23%) have done so.

If a paper you are reviewing does not cite a very relevant preprint paper you are aware of, do you penalize the paper for this?



If a paper you are reviewing does not compare with the empirical results in a very relevant preprint paper you are aware of, do you penalize the paper for this?




14% of reviewers penalize a paper for not citing a very relevant preprint paper, whereas the majority (86%) do not. A slightly larger minority (20%) penalize a paper for not comparing to empirical results in a very relevant preprint paper.

4. Future Policy


This section contains the results for questions about future policies concerning preprint papers and reviewing.

Do you think that preprint servers will become a permanent feature of research within the ACL community?




More than two thirds of the respondents (69%) think that preprint servers with high probability will become a permanent feature of research within the ACL community. Further analysis shows that this view is shared by 84% of the respondents who themselves upload papers to preprint servers at least sometimes but only by 56% of the respondents who do so seldom or not at all.

Which is more important to you - being able to publish your research on a preprint server or having double-blind reviewing for ACL conferences?



The great majority of respondents (88%) consider double-blind reviewing at ACL conferences to be important. About 75% of those (65% of all respondents) consider double-blind reviewing more important than being able to publish preprints. Only a small minority (9%) consider preprint publishing more important than double-blind reviewing.

Why is double-blind reviewing important to you? [Multiple response]



Among those who consider double-blind reviewing important, the following four reasons are all given high weight (with most weight given to bias prevention and least weight to program diversity):

It helps to maintain the quality of ACL papers


It helps to maintain the diversity of ACL papers (keeps the program interesting)

It helps to prevent unconscious bias towards authors from underrepresented groups

I do not want my reputation (or lack thereof) to affect the review of my work (positively or negatively)


The free text comments emphasize the same points, in particular the prevention of undue bias and maintaining quality.

If publishing on a preprint server was banned before acceptance, would you still submit your work to ACL?



A clear majority (87%) would probably still submit their work to ACL if preprint publishing was banned for conference submissions. Less than 5% would probably stop submitting to ACL.

How would you like to see ACL's reviewing model working in the future? [Multiple response]




The most popular choices are in support of double-blind reviewing, combined with banning or discouraging preprint publication or keeping the status quo. Among those who support abandoning double-blind review, completely open reviewing has slightly stronger support than single-blind reviewing. The great majority of free text comments are in favor of having an option to conceal author identity for preprints while under review to preserve the integrity of double-blind reviewing (see next question). Other suggestions are to publish reviews (to improve quality and prevent bias) and to have continuous conference reviewing all year (to preempt the need for preprint publishing).

Further analysis shows that of the 240 respondents who selected *Ban preprint before paper acceptance*, 147 selected no other option. This group, which represents 24% of all respondents, can be seen as the strongest opponents of allowing preprint papers to be submitted to ACL conferences. Of the 383 respondents who did **not** select the banning alternative, 45% were respondents who themselves normally do not upload papers to preprint servers.

The graphs below show the responses to future reviewing models broken out by gender, role, and self-reported fame. The most notable aspect is that female respondents are less favorable than males, proportionally for the survey, of single blind reviewing or reviewing in which all authors are visible.


How would you like to see ACL's reviewing model working in the future?

Gender




How would you like to see ACL's reviewing model working in the future?

Role




How would you like to see ACL's reviewing model working in the future?

Self-reported Fame



The following are some actions suggested by members of the community. Please indicate which, if any, you would like to see implemented. [Multiple response]



There seems to be support for lobbying preprint servers to allow (temporarily) anonymous publication (see also comments to previous question), for providing guidelines for citing preprint papers, and for introducing journal-style reviewing in collaboration with CL and/or TACL. The idea of having a separate track for preprint papers appears to be less popular. Free text comments again advocated anonymous preprints (some suggesting that ACL should set up its own preprint server if arXiv does not provide this option). Some were in favor of more journal-style reviewing, and some specifically said that papers should not be penalized for not citing preprints.

The table below shows the responses broken down according to whether respondents themselves tend to upload papers to preprint servers. It may be noted that frequent users are more supportive of developing guidelines for citation of preprint papers, while infrequent users show relatively more support for the idea of having a separate track for preprint papers.

	A1	A2	A3	A4
All respondents	55.4%	50.4%	24.2%	45.3%
Respondents who always/often/sometimes upload papers to preprint servers	57.9%	59.0%	17.6%	45.5%
Respondents who seldom/never upload papers to preprint servers	53.2%	42.9%	30.0%	45.0%

A1 – Lobby preprint servers to allow papers to be anonymously uploaded

A2 – Make available author guidelines for citing preprint papers or referring to empirical results in preprint papers

A3 – Have a separate track at ACL for papers that have been uploaded to a preprint server before notification

A4 – Journal-style reviewing with a rolling deadline (in conjunction with CL and/or TACL)

If you have any additional comments or suggestions, please provide them here.

This section of the survey ended with a completely open question asking for additional comments and suggestions. This question was optional and was answered by 113 of the 623 respondents.

The responses encompass comments on the current situation, opinions on the suggestions for future policy provided in the previous two questions, new innovative suggestions for future policy, and comments on the survey itself. All sides of the debate are represented.¹ Below is a summary of the main views expressed.

Double-blind review

Double-blind reviewing is the cornerstone of unbiased reviews

The community should take a conservative approach to changing the current reviewing system, which is one of the best and fairest I know

I think it's important to keep the option of anonymity of authors at least for newcomers not to be at a disadvantage

Double-blind reviewing does not need to be perfect and pure to be useful

Double-blind reviews has the problem of unfair reviews that you can't deal with (even a rebuttal period hardly changes anything)

Many of the free-text comments express support for double-blind review as important for maintaining quality and preventing bias. Some suggest that authors should have the right to be reviewed anonymously even if not all submissions are reviewed in this way. Some maintain that, even in the absence of preprint papers, double-blind reviewing is not really blind. A small minority express support for abandoning double-blind review, while another group sees this as a possibility if the idea of anonymous preprints turns out not to be viable (see below).

Preprint publishing

Let's encourage the rapid and free exchange of ideas, data, and results by being open to new ways of publishing and reviewing.

¹ Of the 113 respondents who supplied comments, 47% seldom or never upload research papers to preprint servers, with 29% doing so often or always, and 24% doing so sometimes.

A community that encourages pre-publication penalizes researchers who value quality more than short-term hypes.

I see preprint servers as a convenient outlet for frustrations with the current state of scientific publishing. E.g., negative results and methodological papers are hard to publish, even if they are informative

[arXiv] decouples dissemination from conference attendance. It is well past time for our conferences to evolve in recognition of this. Rather than police the reviewing process even more tightly, ACL needs to recognise that most conference papers (and especially arxiv papers) are ephemeral, and not take them so freaking seriously

Many respondents express support for allowing preprints. Arguments in favor of this are based on promoting fast research dissemination and progress, but also on a concern that ACL may lose good papers and alienate researchers from related fields if preprints are banned. Some also argue that conference reviewing is conservative and of declining quality and that preprint publishing is needed to balance this. A few respondents expressed their support for banning preprints at conference submission time, while another group sees this as a possibility if the idea of anonymous preprints turns out not to be viable (see below). There were some also suggestions related to distinguishing preprint and non-preprint papers by, for example, not allowing preprint papers to be presented orally or to win a Best Paper award, or choosing a non-preprint paper in the event of a close acceptance decision. Some concern was also expressed about the use of preprint servers for “flag-planting” purposes.

Anonymous preprints

The best option would be to implement anonymous preprints, that can be de-anonymized once they are accepted.

Many respondents argue that the best way to resolve the conflict between preprint publishing and double-blind review is to convince preprint servers to allow temporary masking of author identity (and maintain the status quo in other respects). A handful of responses expressed concerns about the anonymous option (orphan citations, social media still being used to promote preprint papers). If the anonymous option turns out not to be feasible, opinions diverge about how to proceed, with four main positions:

- ACL should set up its own preprint server with this feature.
Why doesn't ACL start an anonymous arXiv server of its own instead of lobbying?
- ACL should ban preprints from submission.
Publications on preprint servers should be strictly forbidden before acceptance
- ACL should give up double-blind review.
Abandon double-blind reviewing. It makes things unnecessarily complicated, and does not provide any advantages.

- ACL should maintain status quo (aka double-blindish reviewing).
I think the current ACL policy is about the best we can do.

Citation of preprints

Papers should only be required to take into account related work that is peer-reviewed.

It is a real problem to be asked to cite preprint papers, that have not been validated by anyone! Some of these papers are just really bad.

I'm actively against according preprints any kind of "prior art" status in reviewing because that leads to a completely messed up incentive structure: publish any scrap of result on preprint servers to serve as prior art against other work in the area.

It is very hard to ignore a preprint version in literature review since many reviewers take them equal to a published paper.

A substantial number of respondents highlight the need for guidelines (for both authors and reviewers) regarding the citation of preprints. Most of them maintain that citation should be discouraged except in very limited circumstances, because preprints are not peer reviewed, and that papers should not be penalized for failing to cite preprints, but a few express the opposite view that everything counts as prior work.

Quality of reviews

Please have a survey on the problem of reviewer loads, reviewing processes and conference spamming next! The rate of increase of reviewing loads is not sustainable

I believe the lack of good reviewers is a much more severe problem that may ruin the standard of ACL in the near future

We reject a lot of good papers, and let in a lot of bad ones

I was really disappointed by the reviewing quality of ACL sometimes. Many reviewers are actually not qualified to review papers and many did not spend enough time on the papers.

A substantial number of respondents express concern about the current conference reviewing process with increasing reviewer load and decreasing review quality. Here are some of the recurrent suggestions:

- Rolling conference deadlines (for both submission and notification) may reduce the incentive for preprint publishing and improve reviewing quality. It may also be combined with more journal-style reviewing.
Please, please please give us the rolling deadline model that works so well in other fields

- Publishing reviews together with papers may improve reviewing quality. One suggestion is that reviews of preprint papers should be published regardless of whether they are accepted or not.

If a reviewer knows that his review will be published, probably he will more carefully choose the wording, remarks and grades and include better justifications.

- Preprint publishing and conference reviewing can be merged into a single process, where papers are submitted to preprint servers and accepted to a conference after they have received enough recommendations by reviewers.

Establishing a (parallel) workflow where pre-print papers can be peer reviewed and evaluated could help identify highly useful papers