

Submit to journals

Katherine Lee

April 23, 2018

Abstract

Conference attendance and submissions have grown exponentially in deep learning, while journal submissions have not kept pace. Conferences' growing influence and competition in the field drive researchers to publish at each major conference. While conferences are designed to facilitate in-person interactions and arXiv is designed to disseminate research, journals are designed to publish completed work. As a field, we should increase the impact of open-access rolling journals by publishing in them. This will allow us more control over research timelines and diversity in types of articles while maintaining the peer review process and the defining openness of our culture.

Submitting to conferences by default sets the research timeline to three months¹. This is unsustainable and does not encourage diverse forms of research. As deep learning becomes more mature and public, we must think long-term, maintain investments made in young researchers, and incentivize seminal work. We, as a field, will alleviate pressure, promote influential work, and maintain a high bar of research by incorporating rolling deadlines while maintaining peer review. Thankfully, we already have open-access publications that incorporate both rolling deadlines and peer review: journals like Journal of Machine Learning Research (JMLR), Journal of Artificial Intelligence Research (JAIR), Public Library of Science (PLOS), Distill, and more. Let's submit to them more.

Rolling deadlines provide relief and allow diverse questions

The three-month timeline creates resource crunches, is stressful, and incentivizes work that can be completed in three months. Rolling deadlines help alleviate these problems and promote higher quality of work by forcing researchers to decide when they are done.

During the weeks leading up to a major conference deadline, researchers work into the night and shared computational resources become strained. Days before the deadline, researchers

¹Conference deadlines are not exactly three months apart. In the past year, the three largest conferences' submission dates were: ICLR October 27, ICML February 9, NIPS May 18. Interwoven between conference deadlines are workshop submissions and specialized conferences deadlines (ICLR 2018; ICML 2018; NIPS 2018).

wonder if the science in their project will be finished. When everyone submits to the same deadlines, researchers are unable to help peers review work. The expectation that researchers submit to multiple conferences a year means that stressful periods before deadlines repeat continuously. Spreading work throughout the year through rolling deadlines allows better utilization and allocation of human and computational resources. When researchers can get feedback before submitting papers, they can iterate and arrive at more complex insights. Additionally, I have observed weeks of low computational resource utilization following major conference deadlines. Spreading experiments more throughout the year can make resource sharing more equitable and efficient.

A project that is not yet finished during the week leading up to the deadline could either be rushed to completion, or submitted at the next conference. The former incentivizes researchers to submit work before they are confident, leaving writing the paper to the last hours before the deadline. Pressure to submit can cause researchers to neglect performing critical experiments, decreasing the strength of technical arguments in the paper. With rolling deadlines, researchers can take time to be sure the project is complete and submit it for review immediately afterwards. Now, it could be that deep learning is so competitive that the three-month conference timeline extends how long researchers spend on a project. However, given resources crunches and the number of researchers in the office past midnight, it seems unlikely that the same types of questions research asks today could be done faster.

The push to publish encourages researchers to start projects with the goal of submitting to a particular conference, setting a cap on the amount of time available. While some questions can be furthered in three months, many projects, not building off of larger work, will be incremental. The difference between starting a six month project and letting a three month project roll to the next conference deadline is the types of questions researchers ask. The three-month timeline means we tend not to explore questions that take both shorter and longer to consider. As a new and continually evolving field, we need research that lays groundwork, such as: developing metrics, doing theoretical research, defining terms (what does it mean for a model to be interpretable?), debating opinions (to what extent should researchers join policy discussions on health-care?), and surveying the state of a subfield. We need to incentivize a greater diversity of work and, in particular, seminal work.

Peer review is valued

If rolling deadlines are the key, why not submit to arXiv? Though arXiv allows researchers to publish anytime, it does not review individual articles² (arXiv 2018). Currently, researchers concurrently submit to arXiv and high-impact conferences like NIPS and ICML, or journals like Nature³. Researchers continue to use peer-reviewed publications even when presented

²arXiv asks authors to be endorsed by an already endorsed author. An endorsed author can submit anytime.

³From Nature's preprint policy: "The original submitted version of the manuscript (the version that has not undergone peer review) may be posted at any time" (Nature 2018).

with other options. This signals that peer reviews are still valuable as an indication of legitimacy.

A culture of peer review extends beyond the absolute number of papers reviewed. Currently, the number of submitted publications is overwhelming and gives the impression that deep learning is moving forward fast. The community uses tools built on top of arXiv, such as arXiv Sanity which sorts arXiv articles by hype (the number of shares on Twitter), and social media like Twitter and Reddit to stay informed of relevant work. Reputation for strong peer reviews encourages the community to submit good work and may help decrease the number of submissions. Additionally, unlike social media mentions, peer review explicitly filters for conflicts of interest which helps level the playing field.

We should specifically use journal style reviews with options to accept, revise and resubmit, and reject. Recommendations to revise and resubmit will help researchers incorporate reviews and iteratively refine their ideas. Unlike conference review policies, which do not ask reviewers to verify that papers incorporated suggestions they made, journal-style reviews encourages authors to more seriously consider suggestions. Revise and resubmit is sometimes mimicked by revising and submitting to the next conference. However, since this does not ask the same set of people to review the paper again⁴, it is more akin to re-rolling the peer review dice⁵. Additionally, the distinction between revise and resubmit and reject provides researchers with more feedback on how to progress their work.

Today's journals and conferences

But if journals incorporate rolling submissions and a preferable peer review system, why have deep learning researchers turned to conferences instead? In the past, researchers frequently published in JMLR, giving it an impact factor of 7.48 in 2005 (ResearchGate 2018). Common wisdom now states that a good conference paper is on par with a journal article. This shift is reflected in JMLR's decreased impact factor along with NIPS, ICML, and arXiv's increased impact (ResearchGate 2018; Google Scholar 2018)⁶⁷. This may be because

⁴TACL (Transactions of the Association for Computational Linguistics), a journal in an adjacent field has options to accept, conditionally accept, reject but revise and resubmit, and reject. Conditional acceptances at TACL will be reviewed by the same reviewers. Reject but revise and resubmit may or may not have the same reviewers (Transactions of the Association for Computational Linguistics 2018).

⁵Conference acceptances are thought to have some randomness. The NIPS experiment in 2014 found that two independent program committees disagreed on 25.9% of the papers reviewed by both committees (Price and 2018).

⁶This ranking is done by h-5. From the Google Scholar page: "h5-index is the h-index for articles published in the last 5 complete years. It is the largest number h such that h articles published in 2012-2016 have at least h citations each" (Google Scholar 2018).

⁷In 2013, JMLR saw 764 submissions (Journal of Machine Learning Research 2018). In 2014, NIPS received 1660 submissions. There are no statistics on JMLR's website for submissions past 2013, but NIPS 2017 saw 3590 submissions (NIPS 2017). To be absolutely correct, we should look at deep learning submissions at JMLR and NIPS; however, I was unable to find statistics for this. Anecdotally, researchers in 2018 prefer to submit

researchers believe publishing in journals is slower for the same impact as conferences, conferences provide a necessary social component, and younger researchers may not have heard of journals at all.

Going through the journal review process can be slow. JMLR, deep learning’s flagship journal, has not had a deadline for reviewers and may have given the impression that all journal review processes are slow. But journal reviews can be done quickly. JAIR claims to have a publishing decision in two to three months (Journal of Artificial Intelligence Research 2018) and ACL (Transactions of the Association for Computational Linguistics) gives reviewers three weeks to write a review and publishes the paper when it passes review⁸. On the contrary, the NIPS review process in 2018 from submission to the date of the conference was seven months and four months from submission to receiving a decision from area chairs (NIPS 2018). Before NIPS 2017, I heard the following refrains: “My work was published months ago and everyone’s read it already, but I still have to poster for it, maybe I should talk about what I’m doing now instead,” “I’m going to NIPS to see people, but I think I’ve seen all the interesting papers already.” If rolling-submission journals prioritized quick reviews, they would be better equipped than conferences to turn around reviews as submissions are spaced throughout the year. As an example, ACL has a maximum of one paper per reviewer and a three-week window to review papers.

Publication impact is not a clear institutional change like prioritizing quick reviews. The prestige of a journal is a measure of how much social capital and trust the community gives it. Like how rigorous reviews attract better submissions, prestigious journals attract good work enabling them to become more prestigious and drive up metrics like h-5 and its impact factor. Prestige and impact can be seeded. All it takes is for a few influential papers to be published in the venue to attract more. If we submit more to journals, their impact factors and prestige will rise.

This is not to say researchers should not submit to conferences. Use conferences for what they are uniquely good at: fostering in person research discussions on in-progress research or new and controversial research⁹. By having conversations earlier on in the research process, we can iterate and refine research ideas with more feedback and move further as a field¹⁰. Instead of using conference deadlines as a forcing function to start and finish work, use conference deadlines as motivation to write up and set goals for larger forms of work. Submit

conferences.

⁸Conditional acceptances have two months to complete the revisions and then another round of reviews (three weeks to review) for an estimated four months for two decisions and a revision.

⁹Twitter, while an amazing tool and democratizer in the research community, does not replace in-person interactions. Both forms of interaction, online and in-person, pose their own problems and have their own advantages—the exact nature of which has been extensively studied and would be a much longer paper to discuss. For now, I will simply state that the two are not the same.

¹⁰Currently, journals allow for work that has been previously published in conference proceedings. From JMLR’s author info: “Submissions to JMLR cannot have been published previously in any other journal. We will consider submissions that have been published at workshops or conferences. In these cases, we expect the JMLR submission to cite the prior work, go into much greater depth and to extend the published results in a substantive way” (JMLR 2018).

in-progress work to conferences and get feedback from peers. The key is: conferences are not the only submission venue.

Additionally, journals need not worry too much about fast turnaround time. While journal reviews can be at pace or faster than conference reviews, submitting to preprint servers, like arXiv, will always be fastest since they specialize in dissemination. On the other hand, journals specialize in evaluation, and conferences are uniquely good at community building, but our conferences currently conflate evaluation, community building, and, to an extent, dissemination. We may do better by having specialized venues. In my ideal world, we would separate the goals of dissemination, evaluation, and community building. We would use arXiv for dissemination, journals for evaluation, and conferences for community building. Researchers would send work to arXiv the same day they submit to a rolling journal for peer review. Work accepted to the journal could be presented in a poster session at the next affiliated conference. This frees conferences to focus on in-person interactions and prioritize workshops discussing narrow realms of in-progress research and controversial discussions. In my ideal world, we would not submit directly to conferences. I am far from the first person to propose divorcing dissemination, evaluation, and community building. This proposal is heavily based off how the journal TACL interacts with ACL conferences. Additionally, events like Unconference and Dali successfully prompted deeper discussion into the state of the field and controversial topics.

Conclusion

We can reap most benefits of the ideal: decreased pressure and promoting seminal research by submitting to JMLR, JAIR, PLOS, Distill, and other open-access journals.

To researchers who have been in the field for longer, perhaps the roles of journals and conferences are clear. I have observed that many researchers new to deep learning are less familiar with journals (“Oh yeah, journals, you mean conference publications?”, “JMLR, is that like JSTOR?”) and are under the impression that we should submit to conferences, or upon missing the deadline, put it on arXiv. Ask your students and mentees if they are familiar, and if not, please share with your students.

I acknowledge the irony that I am submitting this piece to a conference workshop. However, I am hopeful that this is the beginning of the discussion on conference timelines and journals, and not the end. I look forward to hearing your perspectives on how to reduce stress and promote diversity of research and a hearty round of debate at ICML.

Acknowledgments

Thank you to Ishaan Gulrajani, David Dohan, Doug Eck, James Bradbury, Samy Bengio, Ian Goodfellow, Catherine Olsson, Chris Olah, Ludwig Schubert, Colin Raffel, Kevin Murphy, Maithra Raghu, Tom Brown, and many others at Google Brain for insightful discussions and feedback. I've learned a lot through this process, and hope we can continue the conversation.

References

- arXiv. 2018. "General Information About ArXiv." <https://arxiv.org/help/general>.
- Google Scholar. 2018. "General Information About ArXiv." https://scholar.google.com/citations?view_op=top_venues&hl=en&vq=eng_artificialintelligence.
- ICLR. 2018. "ICLR 2018 Meeting Dates." <https://iclr.cc/Conferences/2018/Dates>.
- ICML. 2018. "ICML 2018 Meeting Dates." <https://icml.cc/Conferences/2018/Dates>.
- JMLR. 2018. "Journal of Machine Learning Research (J MACH LEARN RES)." <http://www.jmlr.org/author-info.html>.
- Journal of Artificial Intelligence Research. 2018. "Submissions." <https://www.jair.org/index.php/jair/about/submissions#submission>.
- Journal of Machine Learning Research. 2018. "Journal of Machine Learning Research Statistics." <http://www.jmlr.org/stats.html>.
- Nature. 2018. "Preprints." <https://www.nature.com/authors/policies/preprints.html>.
- NIPS. 2017. "The Conference on Neural Information Processing Systems 2017 Sets Record for Attendance, Introduces New Tracks." <https://nips.cc/Conferences/2018/Press?year=2017#demo-1853>.
- . 2018. "NIPS Meeting Dates 2018." <https://nips.cc/Conferences/2018/Dates>.
- Price, and Eric. 2018. "The NIPS Experiment." <http://blog.mrtz.org/2014/12/15/the-nips-experiment.html>.
- ResearchGate. 2018. "Journal of Machine Learning Research (J MACH LEARN RES)." https://www.researchgate.net/journal/1532-4435_Journal_of_Machine_Learning_Research.
- Transactions of the Association for Computational Linguistics. 2018. "Peer Review Process." <https://www.transacl.org/ojs/index.php/tacl/about/editorialPolicies#peerReviewProcess>.