



BLOG @ CACM

The *Communications* website, <http://cacm.acm.org>, features more than a dozen bloggers in the BLOG@CACM community. In each issue of *Communications*, we'll publish selected posts or excerpts.



Follow us on Twitter at <http://twitter.com/blogCACM>

DOI:10.1145/3510545

<http://cacm.acm.org/blogs/blog-cacm>

Picking Publication Targets

Carlos Baquero offers guidance on how to decide where to publish one's paper.



Carlos Baquero
Publishing, The Choice and The Luck
<https://bit.ly/3czMr6h>
November 22, 2021

The Context

Selecting publishing targets is a process with no simple solutions.

The best target for a given paper draft will depend on multiple factors. **The quality** and length of the work—all results are different, and not all results, presentation quality, and breadth have the same potential; **The moment**, having in mind approaching deadlines and contemporary practices of the community—with time conferences can lose or gain impact, new journals are created; **The career** phase and project commitments—early-stage researchers might not afford to wait years for a high-impact publication, and the same might happen for fulfilment of project deliverables in time for reporting.

Even taking all these factors into account, the outcome is still highly unpredictable, especially in highly competitive conferences. In 2014,

NeurIPS, a top tier conference in Machine Learning, ran an experiment (<https://arxiv.org/pdf/2109.09774.pdf>) with its review process to find out that acceptance decisions are highly inconsistent among reviewer subsets and, maybe not surprisingly, that the randomness in the decision increases as acceptance rates decrease. Reviewers are more consistent in detecting weaker papers than in agreeing on the best ones. Another important finding was that perception of paper clarity was the best predictor of future paper impact (in number of future citations), while the overall ranking score of accepted papers had almost no correlation with future impact.

While journal publishing will likely suffer from the same assessment inconsistencies, by keeping a dialogue with the same set of reviewers, the chance of successfully improving the paper to meet review goals is more predictable. However, the time span until publishing can be considerable in the more established computer science journals.

Some conferences, such as VLDB (<https://www.vldb.org/>), have migrated to a mixed journal/conference

model, with frequent deadlines across the year and quick turnaround time for the first feedback on submissions. When a submission succeeds, it is published in the journal and assigned a presentation slot in the upcoming conference.

Speed and predictable review cycles, one of the advantages of conferences, are also pursued by diamond open access journals. These journals, like *The Programming Journal* (<https://programming-journal.org/>), are not-for-profit approaches to reviewing and publishing, supported by sponsoring academic institutions. This model, if proven successful at scale, can address the shortcomings of open access. The migration to open access transferred the publication cost from (library) subscribers to authors, and the economic incentive to accept publications can often lead to predatory practices, including oversized editorial boards and pressure for special issues with invited submissions and no bounds on number of publications per issue.

Another important component is the role of early dissemination of ideas, networking, and nurturing of

new niches. For early dissemination and self-publishing, ArXiv (<https://arxiv.org/>) is now mostly accepted and compatible with future submissions (with some exceptions, so authors should always double-check). Workshops in established conferences are also usually a good option for early stage ideas, can drive networking with a highly committed audience, and often count on keynotes from leaders in the field.

The spectrum of possible publishing modalities is quite vast, and obviously not limited to the cases covered above. When a given modality is chosen, others are not and, given the limitations in time and resources, there is always an opportunity cost with the choice. In the end, the decision is with the authors, given the concrete circumstances.

A Balanced Portfolio

One could consider two opposing publishing strategies:

- A **maximalist** strategy would focus on producing as many papers as possible with minimal investment in each result, and full use of conferences and journals with high acceptance rates. Such a strategy, with time and resources, would lead to a respectable h-index. However, it does not seem very likely to produce high-quality results with worldwide impact.
- A **perfectionist** strategy would strive to achieve the perfect scientific result and only publish high-potential results in the most demanding venues. While this strategy could succeed in producing high-impact results (for example, Wiles's proof of Fermat's Last Theorem, which took seven years to complete (<https://bit.ly/3daOePv>)), it is obviously very risky. Also, often authors need frequent feedback to adjust course and find the most fruitful research lines.

The following ideas are meant as possible directions for authors that strive to balance the bets and provide a chance for high impact. All this, while managing a more steady flow of production and evolution of ideas. The suggestion is for each author to find a mix among the established mo-

dalities that offsets the risks of betting on a single modality, but still allows for the level of commitment to results that is essential for high-impact research.

Top-Tier Conferences—Conferences such as those identified in the CS-Rankings site (<http://csrankings.org/>) or as CORE A* (<https://www.core.edu.au/conference-portal>), have been identified by the scientific community as the most prestigious venues in each field. The acceptance rate is low, and perfectly good papers might not be accepted, both due to the high bar and inevitable random effects that come with it. However, an accepted paper will be highly visible, hard to ignore in the related work of future papers in the area, and can be on the right track to high impact. Even for rejected papers, the quality of review feedback is usually high, and a common strategy is to keep improving the work for a couple of re-submissions before targeting second-tier venues. Some authors even keep this improvement cycle along three or more years, never downgrading targets, and eventually succeed (notice that in the presence of random effects, estimated as high as 50%, resubmission also amortizes that risk and is more justifiable).

Top-Tier Journals—Here, top-tier does not mean a low quantile or high journal h-index, since a journal that publishes many articles per year can have a high h-index and a low citation rate per article (for example, Springer LNCS (<https://www.springer.com/gp/computer-science/lncs>), a series that serves as a vehicle for conference proceedings publication, has the second-highest h-index in SCImago, but each article has less than two citations on average). Established top journals have a small and recognizable editorial board, that often vets submissions before review to match scope criteria, have limited slots per year, still support printed editions, might offer open access as an option but do publish without author fees and are often associated to ACM, IEEE, and major publishers. Due to the usually lengthy review time, some authors adopt the strategy of publishing first a shorter version in a conference and then submitting the complete version to a journal (such as a version with all the for-

mal proofs of correction or the full experimental results). Another option is to place a version in a preprint repository at the time of submission, allowing the work to be known and timestamped during review.

Workshops, Seminars, and Schools—Workshops in top-tier conferences are usually an excellent opportunity for networking and for the presentation of short results or initial research in order to receive useful feedback. The co-location with a reference conference draws many world-class researchers and usually enforces a level of quality control on hosted workshops and offers (short-format) publication mechanisms. Some workshops allow presentation only and do not require publication. Dedicated seminar centers, such as Dagstuhl (<https://www.dagstuhl.de/>), also have a rigorous screening process for hosted events and support post-event reports, however these events are often by invitation only. Topic-centered Summer/Winter schools are other potential venues for networking and collaboration. Due to the usual high quality of attendance, all these events are a better investment of time than many third- and fourth-tier conferences.

The Choice

Even with a perfect strategy and efficient use of time, few world-class researchers are able to consistently produce high-quality work that later proves to have a lasting impact in the field. While some of their papers will have an impact (in terms of influencing research and industry use), others will not. Accordingly, the number of papers produced is still important, but what seems to be the best predictor of future impact is still paper quality and clarity. That might be the lucky choice.

Acknowledgments

I would like to thank Rui Abreu and João Cardoso for feedback on this text.

Carlos Baquero is a professor in the Department of Informatics Engineering within the Faculty of Engineering at Portugal's Porto University, and also is affiliated with INESC TEC. His research is focused on distributed systems and algorithms.