



# Reflections on reciprocity in research

Peter A. Flach<sup>1</sup>

© The Author(s), under exclusive licence to Springer Science+Business Media LLC, part of Springer Nature 2020

This is a time of reflection—for more reasons than one. I am writing this piece in the eleventh week of ‘working from home’ which has rapidly become the ‘new normal’. But even without the Covid-19 pandemic I would have been writing an editorial for *Machine Learning* at this point in time—albeit perhaps not in my back garden!

After having had the privilege to serve the international machine learning community as Editor-in-Chief of the *Machine Learning* journal since July 2010, it is now the moment for me to step down and hand over the reigns. The previous decades have seen tremendous change in the practice and perception of machine learning research, accelerating in the last ten years in particular. When I started out as a young researcher the question whether computers could learn or think was confined to nerdy newsnet groups. Today, the Turing test might make an appearance as a plot device in mainstream movies, and machine learning and AI are seen as key technologies and even marketing narratives. Clearly, the research landscape has changed considerably. The question I want to consider here is: how do we as machine learning researchers respond to these changes?

## 1 The changing research landscape: scale or substance?

As a practicing machine learning researcher you would be forgiven to think that, on the face of it, the changes are predominantly ones of scale rather than substance. Most of us will be aware of the dramatic increase in N(eur)IPS attendance, which nearly tripled over the last four instalments to some 13,000 in 2019. But, while quietly suffering the longer queues in the coffee breaks, most of us still go about our research in the same way: driven by the annual conference cycle, we submit our papers to the next available conference or journal, and wait for the peer reviews to come in.

However, when digging a little deeper it is not hard to see a range of recent changes in the way research is done and reported. For one thing, the increased scale of today’s conferences has actually *reduced* the opportunities for scientific debate, which arguably was the main reason scientists started having such meetings in the first place. Consequently, debate is finding new outlets on blogs and social networks, which hasn’t exactly improved the quality of the arguments (just imagine Niels Bohr and Albert Einstein battling it out on Twitter regarding the ‘right’ interpretation of quantum mechanics...).

---

✉ Peter A. Flach  
peter.flach@bristol.ac.uk

<sup>1</sup> Intelligent Systems Laboratory, University of Bristol, Bristol, UK

Another notable change is that much of machine learning research is now done in industry. Today, all large technology companies have machine learning research groups, as do a lot of the medium-sized ones. To some extent one could argue that this reflects a natural progression from fundamental research to R&D closer to market, as has happened in many other, once distinctly academic research areas such as processor design and computer graphics. However, the difference in machine learning is that many of these industrial research groups are doing *fundamental* research, on a par with research undertaken at universities. A large share of the research papers at NeurIPS and other machine learning conferences are now produced by Google, Microsoft, and Amazon, among others. Researchers at these companies are often working under much more favourable conditions compared to their academic colleagues, and are major players in setting the research agenda and driving it forward.

The steadily increasing appetite for alternate dissemination channels is also worth mentioning. Questions about the role and profits of commercial publishers are partly driving this. *Machine Learning* is no stranger to this discussion, having given rise to the establishment of the *Journal of Machine Learning Research* nearly 20 years ago. This is not the place for me to comment on that debate: suffice it to say that the machine learning community is fortunate in being healthy enough to sustain two top-tier journals; and that Springer, the publisher of *Machine Learning* (now Springer-Nature), has understood the need to reciprocate and make contributions to the community, as I will briefly explicate later on in this article.

We furthermore see a steep increase in dissemination of work that hasn't (yet) been peer-reviewed through online repositories. The most popular repository, arXiv, started life as a pre-print server for physics research. In physics, as in other more traditional scientific disciplines, academic journals are the main vehicles for peer-reviewed publication. In such disciplines pre-prints are a useful way to speed up dissemination and establish priority for new results. It is interesting to note that computer science has now also embraced the use of online repositories, even though the need for rapid dissemination is less strongly felt in the presence of a wide range of high-quality, annual conferences. There could be several reasons for this: e.g., it may be felt that the top conferences and journals are overly selective (roughly four in five papers are sent back); but it may also be the case that peer review is now seen as less important or productive for the academic endeavour than it was in the past.

It is worth expanding a bit on the role and status of peer review in contemporary research and publishing, traditionally seen as a cornerstone of the scientific method. Much like parliamentary democracy, it may be far from perfect but is still considered the best way known to achieve minimum standards of transparency, reliability, and accountability. At least, this is the theory—but is the peer review system really still fit for purpose? Allow me to make a few observations from my ten years' experience as Editor-in-Chief:

- on the whole, editors need to ask more people (sometimes ten or more) before they find three people willing to review the paper;
- while the average quality of reviews remains high, there is more variability and a longer tail of less useful (e.g., overly succinct) reviews;
- often reviews take longer to materialise, and some never materialise at all.

Without offering a detailed analysis, the least I conclude from these observations is that peer review is under considerable and increasing pressure. This is not, I hasten to add, because researchers have become less well qualified for their job or are taking it less

seriously. Rather, people experience many pressures in their day job, some of which inevitably propagate to the peer review process. Given the central role of peer review in the scientific process this should be of concern to us all. So, what to do?

## 2 Reciprocity as a guiding principle in research

Reviewing someone else's paper may, on the face of it, be perceived as one of the less interesting or productive aspects of being a researcher; something that comes with the territory, but takes time away from the main part of the job; in short, something of a chore. While it *can* be a way of keeping up with the field and learning something new if the paper is good, for the average conference or journal submission the odds are stacked against this. Sure, the authors may benefit from getting some feedback on their work, but the reviewers don't get much reward for their time and effort. Or do they?

In essence, scientific publishing is a voluntary system that works insofar it adheres to the social norm of *reciprocity*: people should roughly put as much into the system as they take out of it. Submitting a paper to a conference or a journal means some people have to review it. Put differently, every submitted paper incurs a reviewing debt—and every review yields submission credits. For such a system to be sustainable, as a community we need to balance the books. To understand what this means for the average participant, it is illuminating to do a back-of-the-envelope calculation. Let

- $A$  be the average number of authors per paper (say 2.5);
- $P$  be the expected acceptance rate (say 0.2);
- $R$  be the typical number of reviews per paper (say 3);

then

- each submitted paper incurs a reviewing debt of  $R$  (3) per paper or  $R/A$  (1.2) per author;
- each *published* paper incurs a reviewing debt of  $R/P$  (15) per paper or  $R/PA$  (6) per author.<sup>1</sup>

The latter number (6) is roughly equal to the reviewing load for an individual program committee member of a conference. In other words, *for each published conference paper an author should review for one conference in a year*. Similarly, for each published journal paper one should expect to review three or four journal papers (assuming that such a review is 50–100% more work compared to a conference review). This is roughly equal to the number of reviews an editorial board member is expected to do annually.

These very rough calculations are just given to make a point, and require much further reflection and elaboration. For one thing, reviewing activity is not expected to be uniformly distributed over the academic community: there are many variabilities, e.g., career stage. However, the underlying principle of reciprocity is an important one that deserves full recognition as an important factor in the sustainability of a research community.

<sup>1</sup> This assumes that rejected papers get resubmitted until eventual acceptance, which on average takes  $1/P$  submissions.

In fact, the social norm of reciprocity in the research process is not limited to peer review. I already mentioned that it would be reasonable for commercial academic publishers to reciprocate in recognition of the—often unpaid—work that academic editors and reviewers do for a journal. I think that here *Machine Learning* is doing more than some, in that

- from one year after the issue publication date, the published article is freely available online;
- authors keep copyright of their work;
- the journal sponsors best student paper awards at many machine learning conferences.

Having said that, perhaps more can be done by publishers, for example in providing copy-editing services for accepted papers in order to make sure the published versions achieve their full potential.

There are several other areas where reciprocity could lead to a rethinking of the *status quo*, for example in relation to the increased involvement of industry I mentioned earlier. Just to illustrate:

- academic conferences bring additional value for companies in terms of visibility and recruitment opportunities, so reciprocity suggests that those companies provide financial support, e.g. through higher registration fees for company delegates;
- recruiting graduates fresh out of university could be seen as creating a ‘PhD debt’ as most of those recruits will not undertake a PhD, so reciprocity suggests that industry compensate by co-designing and sponsoring relevant doctoral training opportunities with academia.<sup>2</sup>

Considering the full potential of reciprocity as a guiding principle will help us build and maintain a productive and sustainable ecosystem for research.

### 3 Welcome to the new Editor-in-Chief

As I wrote at the beginning of this article, having had the opportunity to steer the journal through a decade of change has been a privilege. Attaching my name to the list of my illustrious predecessors gives me a great sense of pride:

- Pat Langley (1986–1989)
- Jaime Carbonell<sup>3</sup> (1989–1992)
- Tom Dietterich (1993–1997)
- Rob Holte (1998–2003)
- Foster Provost (2004–2010)
- Peter Flach (2010–2020)

---

<sup>2</sup> In the United Kingdom this is currently happening through Centres for Doctoral Training, see <https://www.ukri.org/research/themes-and-programmes/ukri-cdts-in-artificial-intelligence/>.

<sup>3</sup> Jaime Carbonell sadly passed away early this year.

It is now my pleasure to introduce my successor as Editor-in-Chief and latest addition to this list: Hendrik Blockeel.

Hendrik is a full professor at the University of Leuven in Belgium and a highly respected machine learning researcher. Hendrik has made many research contributions in machine learning, for example in broadening the scope of relational learning to include decision trees, among others. Also worth mentioning is that he has a particular interest in academic publishing and pioneered the journal track of the European machine learning and data mining conference (ECML-PKDD). Since 2013 the conference has a journal track alongside the regular conference track, with papers submitted to the former reviewed to journal standards by a dedicated editorial board and accepted papers appearing in an annual special issue of or *Data Mining and Knowledge Discovery*, for data mining papers). The model has been very successful and beneficial to both the conference and the journal; similar journal tracks have since been introduced for other conferences including the *Asian Conference on Machine Learning* (ACML). Hendrik will take up the editorship from July 1st this year. I am sure the journal will be in very good hands for years to come<sup>4</sup> and wish Hendrik all the best in carrying out this important duty.

Stepping down as Editor-in-Chief brings benefits as well as some regrets. I look forward to having more time for research and interacting with PhD students, postdocs and colleagues in Bristol and further afield. But I already know I will miss having the finger on the pulse of contemporary machine learning research. And I'll also miss my regular interactions with all the people who have made the journal into what it is today. Here I particularly highlight the contributions of Dragos Margineantu and Tom Fawcett, who volunteered as editors for special issues and for surveys, respectively, and did a great job over many years. And finally I want to express a few words of gratitude to Melissa Fearon, who has been with the journal for over twenty years as senior editor at Kluwer, Springer, and now Springer-Nature. Working with Melissa has been an absolute pleasure: she cares about the journal and the community, and puts into practice the concept of reciprocity even if she might not have put it into those terms herself. To them, and to all the others who have contributed to the journal as action editors, guest editors, editorial board members, and reviewers, I express my heartfelt thanks.

Editor-in-Chief (July 2010–June 2020)

*Machine Learning*

---

<sup>4</sup> Avid data scientists will have notice the upward trend in time served by successive editors.