
Open Scholarship and Peer Review: a Time for Experimentation

David Soergel
Adam Saunders
Andrew McCallum

SOERGEL@CS.UMASS.EDU
SAUNDERS@CS.UMASS.EDU
MCCALLUM@CS.UMASS.EDU

University of Massachusetts Amherst, Amherst, MA 01003

Abstract

Across a wide range of scientific communities, there is growing interest in accelerating and improving the progress of scholarship by making the peer review process more open. Multiple new publication venues and services are arising, especially in the life sciences, but each represents a single point in the multi-dimensional landscape of paper and review access for authors, reviewers and readers.

In this paper, we introduce a vocabulary for describing the landscape of choices regarding open access, formal peer review, and public commentary. We argue that the opportunities and pitfalls of open peer review warrant experimentation in these dimensions, and discuss desiderata of a flexible system.

We close by describing OpenReview.net, our web-based system in which a small set of flexible primitives support a wide variety of peer review choices, and which provided the reviewing infrastructure for the 2013 International Conference on Learning Representations. We intend this software to enable trials of different policies, in order to help scientific communities explore open scholarship while addressing legitimate concerns regarding confidentiality, attribution, and bias.

Introduction and Background

Continuing evolution of scientific communication practices on the Internet and the recent flourishing of the Open Access movement have prompted an explosion of interest in reconsidering the nature of peer review of journal and conference articles. It is widely recog-

nized that the traditional peer review and publication system has serious deficiencies. Among many other concerns: it is slow; anonymous reviewers may exhibit biases or hold grudges; and reviewers are not credited for their work. These deficiencies have become more acute in recent years, as the Internet is making everything in our society faster and more transparent. The success of the Open Access movement, and the resulting upheaval (both economic and cultural) in the publishing industry, has now made it possible to consider more profound changes in the structure of the publication process. We are consequently in a period of widespread debate as to whether papers should be made public prior to peer review; whether reviewing ought to be single-blind, double-blind, or otherwise; and whether the reviews themselves can or should be made public (Kriegeskorte et al., 2012; Desjardins-Proulx et al., 2013). At the same time, questions have been raised regarding the apportionment of credit for scientific contributions—not only among the authors of a paper, but also among reviewers (who can only be credited if named), authors of insightful blog entries and forum comments, and so forth.

One view holds that scientific progress is impeded by the historical practice of single- or double-blind, secret reviews, because findings are withheld for months to years until that process is complete, and because the community is not privy to the potentially illuminating communications between authors and reviewers along the way. This view has produced calls for a policy of complete transparency (LeCun, 2009), and even calls for doing away with peer review entirely (Wasserman, 2012). Conversely, fully closed reviewing procedures are intended to encourage more candid reviews; to save readers' time by filtering out bad papers before they are made public; and to safeguard against biases that may occur due to gender (Nature, 2012; Budden et al., 2007), perceived institutional prestige, the Matthew effect (Merton, 1968), personal vendettas, and so forth. And yet there is widespread concern that these safeguards are frequently ineffective.

The traditional system is based on the notion that work should not be disseminated until after it has been evaluated. Historically this made sense due to the cost of physical printing and distribution. Now, this approach stifles progress, especially when the evaluation procedure is slow or unreliable. And, in any case, the idea that information can be withheld at all is an anachronism. Researchers who are proud of their work are naturally excited to share it, and are now able to do so with abandon in blogs, Twitter, open lab notebooks, and data- and slide-sharing sites such as Data Dryad (Vision, 2010), Figshare (Hahnel, 2011), and Slideshare (Sinha et al., 2006).

Incentives that prevent researchers from sharing their work (e.g., fear of lost attribution) are increasingly perverse and unnecessary. It is now essential for scientific communities and publishing organizations to develop modern approaches to scientific communication which increase the rate and efficiency of review, publication, and discussion while simultaneously alleviating concerns around confidentiality, attribution, and bias. This must be done while recognizing that different communities have very different cultures and requirements in these matters.

Many parties have embarked on the project of making article dissemination and peer review more open and efficient. The arXiv preprint server is now 20 years old, and has expanded its scope from Physics and Mathematics to include Biology, Finance, Statistics, and Computer Science (Ginsparg, 2011). In the life sciences, a number of efforts have recently launched to provide various flavors of high-speed, open reviewing, including PeerJ (Binfield et al., 2012), eLife (Schekman, 2013), F1000Research (Lawrence & Tracz, 2012), and GigaScience (Goodman et al., 2012). There is also increasing interest in pre-publication discussion of papers, as at Haldane’s Sieve (Coop et al., 2012), and in reviewing and endorsement of preprints via “overlay journals”, which serve to evaluate and filter papers already available from another source such as ArXiv (Demailly et al., 2013; Desjardins-Proulx et al., 2013). Unsurprisingly, for every service or product launched in this space, there are innumerable published papers, blog posts, comments, and tweets proposing and debating alternative approaches.

We believe that the most important innovation in this discussion is the separation of dissemination from evaluation—that is, the idea that scientific works should be made public as soon as possible, and that the peer review process and resulting endorsement by a trusted institution (such as a journal or conference) is a completely orthogonal issue. Furthermore we expect

that making the peer review process rapid, transparent, and participatory will provide great benefits.

That said, we do not know which set of reviewing policies will optimize both the rate of scientific progress and fair treatment of all scientists in the long run. We therefore advocate an infrastructure in which a variety of policies can be tested. The essential task is to allow conference organizers, journal editors, article authors, and other users to set policies which determine who is privy to what information (e.g. articles, reviews, author identities, etc.), and when. This requires a flexible scheme for encoding reviewing workflows, for instance to specify how responsibility for a paper is delegated among journal editors or conference organizers, how reviewers are chosen and what is expected of them, and who is authorized to make a decision to publish or reject. It also requires tracking the progress of papers through these workflows, and—for some policies—a system for ongoing public discussion of papers after (and perhaps prior to) publication.

1. The Dimensions of Open Scholarship

This explosion of ideas about peer review and publication processes, while highly valuable, is also diffuse and confusing; and many of the ideas being discussed may eventually die out. Researchers from a wide variety of scientific disciplines are weighing in, but may have quite divergent experiences and expectations of the peer review process. The discussion also lacks a common vocabulary—for instance, the term “open peer review” is frequently used to mean only that the identities of the reviewers of a paper are revealed to the authors (van Rooyen et al., 1999; Walsh et al., 2000; Smith, 1999). It may also be used to mean that the reviews and reviewer identities are published along with the paper (Koonin et al., 2013); or that anonymous community input is solicited prior to publication (Walker & Holt, 2006); or that anonymous official reviews are published and openly discussed prior to publication (LeCun, 2009).

Our first task in discussing levels of transparency in scholarly work is therefore to clarify exactly which aspects of the peer review and publishing processes might be made transparent under different proposals. Here we provide a general framework for describing and thinking about the many dimensions of openness that may apply to scholarly work, so as to classify the many journals, products, and companies in this space and to understand the relationships among them.

The word “open” denotes access to information. To characterize a system, then, we must state **who** has

access to **what** information, and **when**. (Additionally there may be special **conditions** on that access).

In the world of scholarly communication, the **who** typically includes:

- researchers / authors
- moderators (journal editors and conference program chairs)
- reviewers
- journal subscribers or conference attendees, and
- the general public.

What information those parties may (or may not) access includes:

- primary research materials, e.g. lab notebooks
- "completed" experimental protocols, source code, raw data, and analysis workflows
- manuscript drafts
- completed manuscripts (including supplementary materials)
- identities of manuscript authors
- official peer reviews
- identities of official peer reviewers
- author responses to reviewers, and consequent paper revisions
- unofficial peer reviews, annotations, and comments
- identities of unofficial peer reviewers
- press releases
- presentation slides
- presentation videos

When parties may access information can be described in terms of the typical lifecycle of a scientific work intended for "publication", such as a journal paper, conference paper and presentation, or conference poster. Significant time points include:

- During research
- During manuscript writing
- Upon manuscript "done"
- Upon manuscript submission
- During formal peer review & revision
- Upon journal or conference decision

- Upon journal publication or conference presentation
- N months post publication
- Never

Finally, **conditions** may include:

- The authors authorize an optional openness feature
- The authors pay a fee to enable an optional openness feature
- The reviewers allow their review to be read by some selected set of people, anonymously or non-anonymously.
- A reader pays a fee to gain access to some information

From this set of primitives we can construct sentences describing nearly any openness claim or publication policy. For example, the core idea of Open Access is: the **general public** has access to **final manuscripts**, effective **upon publication**.

2. Exploration of the Landscape of Open Scholarship

A complete policy for a venue consists of a set of such propositions, summarized for a small sample of different venues in Figure 1.

Some interesting highlights of alternative policies—as seen against a background of traditional reviewing—are:

- Reviewer identities are revealed to authors, but not to the public. (Biology Direct, BioMed Central journals, and others).
- Attributed reviews are made available to the public together with the paper upon publication. (F1000Research).
- Reviewers discuss a paper amongst themselves and write a consensus review (eLife).
- Papers are publicly posted immediately, and are not reviewed (arXiv).
- Papers are publicly posted immediately, and are publicly discussed but not formally reviewed (Hal-dane's Sieve, Philica).
- Reviewers may choose whether or not to reveal their identity, and authors may choose whether or not to publish peer reviews together with the paper (PeerJ).

- Authors pay for “portable” peer reviews managed by a third party (Rubriq (Mudunuri & Collier, 2013)).

2.1. Concerns about too much openness

It is not a foregone conclusion that maximal openness of everything is the end goal. We certainly do believe that making things generally more open will accelerate scientific progress, and may help to correct problems with the traditional peer review system. At the same time we are sensitive to social issues that may arise. These include:

- Fears that holding an unpopular opinion will have social or career consequences.
- Concerns that personal reputations and power dynamics may unduly impact evaluations, to the detriment of scientific quality.
- Fears of retribution for negative reviews, especially reviews by younger investigators.
- Fears that releasing results too early risks losing credit for downstream developments.
- Concerns about bias on the basis of gender, nationality, institution, etc.
- Fears that unmoderated commentary will be of low quality.
- Concerns that commenting threads will deteriorate into unproductive emotional arguments (flame wars).

Clearly, at least some these concerns apply to varying degrees not only to open procedures but also to closed ones.

It may be that judicious limitations on the transparency of an otherwise open peer review process can substantially alleviate these concerns. For instance, given that gender bias is widespread and leads to underrepresentation of women among accepted papers, we expect that blinding reviewers to author identities would restore the balance (Budden et al., 2007). Conversely, the argument has been made that complete openness (including non-anonymous reviews) will quickly reveal cases of gender bias and generally encourage more fair and thoughtful reviewing. This is disputed on the grounds that bias is not a matter of a few “bad apples” but rather is a pervasive community-level phenomenon, unconsciously propagated even by the female participants (Nature, 2012; Laba, 2011-2013).

We hope that creative solutions can be found to solve these kinds of issues while at the same time exploring

avenues for accelerating scientific progress. For example: a review process might be simultaneously open, in the sense that papers and reviews are immediately available for public discussion, and double-blind, so that neither the authors nor the reviewers know each others’ identities during the review period; these identities would then be revealed when a decision is made. Authors might be concerned about disseminating their work without attribution (temporarily), but the assurance that the work would become attributed within a reasonable timeframe might assuage this concern. In this case, a trusted party (most likely the venue) effectively acts as an information escrow service, revealing information such as author identities at a later time, while vouching that this information had been submitted previously.

3. Open Scholarship Platform Desiderata

Given the diversity of proposed peer review systems, and the above framework for describing them, we now propose that new infrastructure for peer review and publishing workflows should be flexible enough to accommodate a wide variety of openness policies. The idea of public-yet-doubleblind reviewing is just one example of a policy that might be worth testing; we are not necessarily promoting it, but are simply pointing out that reviewing entities ought to have the option to easily specify policy variants such as this. The fact of the matter is that we do not know which peer review and publication procedures will best address concerns around bias, attribution, and career advancement, nor how these considerations may vary across communities. We therefore advocate development of a platform that is very flexible, to allow designing creative solutions to social and career concerns, while simultaneously allowing communities to collect data about the consequences of different policies.

A unified system for multiple venues. We envision a system in which any person or group can establish a publication venue (or “reviewing entity”, (LeCun, 2009)). Such a system could support a wide variety of venue types, including traditional journals and conferences; overlay journals; private lab meetings and journal clubs; and even personal blogs on which papers are discussed and endorsed. Upon creating a venue, the editors or conference chairs would configure the policy of their choice (or copy a policy from a preexisting venue). Hosting multiple venues in one system provides a fluid path for a paper to mature from a private draft, to a public preprint, to a paper with an official endorsement resulting from a formal peer-review

Open Scholarship and Peer Review

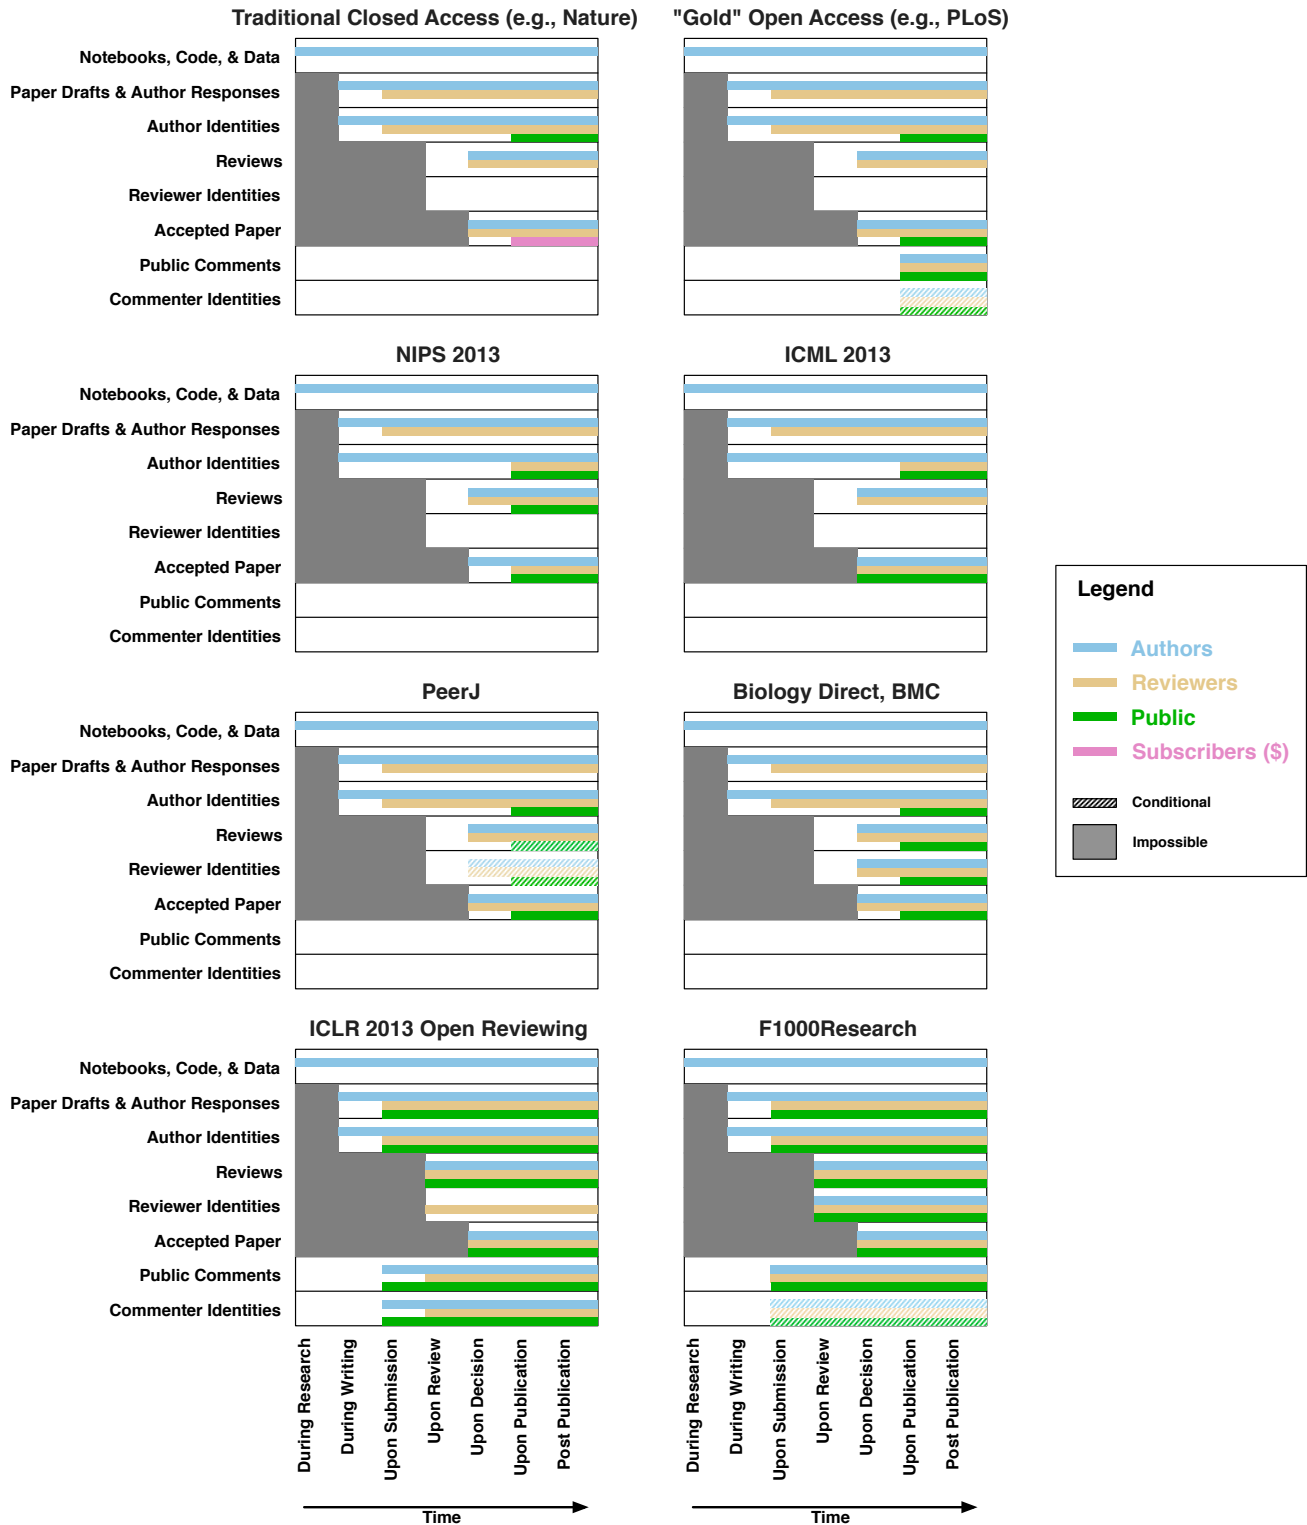


Figure 1. A sampling of openness policies. Different venues employ policies which reveal different pieces of information to different parties at different times. Colored bars indicate that a party has access to a given piece of information. In general, more “open” venues reveal more information, to more parties, earlier in the publication process. These examples are roughly ordered by openness from top to bottom and left to right; of course, given the many dimensions of openness, a strict ordering is often not possible. “Upon submission” here incorporates initial editorial filtering and reviewer assignment. Hatched bars indicate that the release of information is conditional on someone’s approval; in the PeerJ case, reviewers decide whether or not to be anonymous, and authors decide whether the reviews are published.

process. In this case the history of reviews, comments, and revisions is preserved, being attached to the paper and not necessarily to any particular venue, so the idea of “portable peer review” becomes automatic. Indeed, in such a system, multiple entities might endorse the same paper.

Benefits of user-configurable review policies.

A platform for hosting multiple venues with configurable reviewing policies would have numerous benefits. First, it would encourage scientific communities to explore the space of possible policies. In this sense it would provide a testbed for controlled experiments, allowing communities to gain experience with different policy choices while holding other aspects of the infrastructure constant. Computer science is particularly amenable to this kind of exploration, because there are many conferences, each of which may set a different policy (journals, in contrast, have greater longevity and tend to keep their policies stable).

Second, a flexible platform would provide a path for a community to change its conventions slowly, starting from a status quo policy which community members presently accept (for better or for worse), and making incremental changes, building user comfort with increased openness along the way.

Finally, even among communities that feel ready to make a radical change in policy, there will remain substantial disagreements about *which* radically different policy to use. This is a time of exploration and upheaval, so we feel it is far too soon to cast judgment—in the form of a concrete implementation—on one aspect of a policy or another. To do so would risk losing potential users who might feel strongly that we made the wrong choice. For example: in the course of managing the reviewing process for ICLR 2013 (described below), we and the program chairs had a spirited discussion about whether members of the public should be allowed to remain anonymous when commenting on a paper. For this venue, it was decided not to allow anonymous commenting; but a different program committee might easily have made a different decision. From the perspective of providing an infrastructure service intended for wide adoption, we feel it is important to support the widest possible range of user choices.

Furthermore, one long-term result of experimentation with peer review will surely be that the most effective policy will differ by community for cultural reasons: for instance, some fields might be more or less prone to gender bias; some might have historically developed a more congenial or more adversarial culture;

and so forth. Finally, different communities are starting from very different assumptions and expectations about sharing data, ideas, and results. For instance, physicists and mathematicians are already comfortable with releasing preprints on arXiv—many of which are never published in any other forum—while biologists have historically been much more protective, competing for publication slots in high-prestige venues and trying to avoid getting scooped.

4. The OpenReview.net Platform

We built a prototype system, available at OpenReview.net, to begin to test these ideas.

The core challenge was to design a data representation that would provide the desired flexibility. We reasoned that the process of submitting, reviewing, editing, and commenting on scholarly articles is well described in a messaging metaphor, in which each message reveals some piece of information to the recipient. Every step of the process consists essentially of one party sending a message to another: for instance, authors send a paper to conference chairs; conference chairs send the paper to reviewers; reviewers send reviews back to conference chairs; conference chairs send decisions back to authors; and finally conference chairs publish papers by “sending” them to the world.

We therefore implemented our system based on such a messaging metaphor, in which nearly every event is represented by the propagation of messages through a graph of nodes representing individuals or groups of various types, including authors, reviewers, and venues.

The key message types which provide the flexibility described above are the License, which grants read permission on a document to the recipient, and the Identity License, which allows the recipient to know who a (previously anonymous) person is. The “documents” in question may include papers, reviews, author responses, and other comments. By arranging that Licenses and Identity Licenses be sent to different recipients at different times, one can precisely specify **what** is visible to **whom**, and **when**.

Additional message types encode Requests and Responses, which facilitate tracking the state of a paper through a peer review workflow (e.g., to know that a review was requested but has not yet been received).

Such exchanges of messages could of course be handled by simpler systems such as email; however the unstructured nature of that medium makes it very difficult to control and track the status of each paper in a review

process, especially when a large number of papers are in play. Furthermore email is inherently private, and so cannot easily support open models. We therefore provide a structured environment that offers several advantages for managing the flow of messages.

- First it allows describing an expected pattern of messages (e.g, a review process as sketched above), so as to reason about the state of a conversation from the point of view of a given participant. For instance, if a reviewer has agreed to review a paper but has not yet done so, then this can be displayed on that user's "todo" list (or an "overdue" list). From the point of view of a program chair, the same review request appears on a "waiting for" list.
- Second, we provide a flexible permissions system, to allow restricting visibility of papers, reviews, author and reviewer names, and so forth, depending on the policies of the journal or conference, and of the authors themselves.
- Third, we provide public visibility of documents and messages (subject to the permissions model), allowing for an open and transparent reviewing model in which members of the public may discuss papers and read the "official" reviews in a unified environment.
- Fourth (and really as a consequence of the first three advantages), a structured system allows for flexible reporting. A program chair can view the list of reviews requested but outstanding; the public may view a list of papers accepted to a given conference; and so forth.

5. Pilot test at ICLR 2013

We deployed our prototype system to provide paper submission, reviewing, and public discussion for the International Conference on Learning Representations (ICLR) 2013. At this early stage, the system provided a basic feature set; some tasks were handled manually; and the conference policy was not user-configurable. However: the system employed the full messaging model described above. This design proved to be an effective means of tracking the state of a paper through the review process, and of handling permissions and anonymity.

This conference employed the following policy:

- Papers were first submitted to arXiv, where they were publicly visible and archived in perpetuity.
- Papers and author identities were visible to everyone upon submission.

- Official peer reviews were visible to everyone upon submission.
- Reviewer identities were not visible to the authors or to the public, but were visible to other reviewers of the same paper and to the program chairs.
- Reviews of and comments on a paper were hidden from official reviewers until they completed their review.
- Anyone could comment on the papers and on the reviews at any time, with visible author identities. Official reviewers could however maintain their anonymity throughout a conversation with paper authors and other commenters.
- Comments were not moderated.
- Comments could be made publicly, or privately to the authors, reviewers, conference organizers, or combinations of these.
- Discussions remained active before, during, and after the official peer-review phase.

Sixty-seven papers were submitted; these received a total of 178 reviews. Many of the papers produced vigorous discussion threads. Across the entire conference, paper authors gave 119 responses, and anonymous reviewers made 19 follow-up comments. 20 comments were made by conference participants on papers with which they were not directly associated (i.e., on which they were neither an author nor a reviewer). Notably, 9 papers received comments from users who had no affiliation with the conference at all (i.e., who were not authors or reviewers of any paper). In fact, one such commenter reported his discovery of an important bug in a proof, leading to a revision.

We received very few comments after the acceptance decisions were announced (Figure 2). It remains to be seen whether papers receive comments after a significant delay, and whether authors continue to revise papers after the conference has occurred. Because users are not likely to actively visit discussions regarding past papers, we feel that our feature of notifying interested users of new activity by email is essential to keep discussions alive.

While our implementation and the ICLR 2013 policy were oriented towards public discussion, we also provided for comments to be made privately to the conference chairs, or among the reviewers, and so forth. Reviewers were asked to give the program chairs private rankings, and did so; aside from this, the private message feature was rarely used. It may be that users simply found that feature unnecessary, but it may also be that the option was insufficiently clear in the user




Figure 2. Timecourse of ICLR 2013 paper submission and commenting.

interface. Also, the conference policy for ICLR 2013 did not emphasize the use of private messages to the extent that another policy might—for instance, in the eLife model, reviewers are asked to discuss a paper amongst themselves and to provide a consensus review.

6. Community response to open reviewing

Anecdotally, the reviews, author responses, and public comments were highly substantive and congenial throughout. The conference organizers and participants were on the whole very enthusiastic about the prospect of open reviewing and our implementation. The organizers of ICLR 2013 declared the experiment an “unmitigated success”, and have chosen to use OpenReview.net again for ICLR 2014.

We believe it is important to gather more than anecdotal feedback, however. Thus we administered an anonymous user survey to the ICLR 2013 participants to get more quantitative information about users’ likes, dislikes, and concerns. Of the 160 conference participants, over 100 submitted survey responses, indicating tremendous interest and engagement in the topic of open reviewing. The respondents were diverse in their career status, ranging from graduate students to tenured professors, and including industry employees.

Both authors and reviewers were well represented. Of course we recognize that the results of a single survey may not generalize: perhaps the ICLR 2013 community is unusual in some way; for example, the participants were largely white (not measured, but visible at the meeting) and male (88%). We look forward to comparing these results with future surveys of different communities using different reviewing policies.

The overall reaction revealed by the survey was overwhelmingly positive. In fact, *nearly every respondent* expressed a positive view of future conferences with similar open reviewing policies. Summary results of the survey (Figs. 3-8) include:

- The quality of the papers was increased by the open model, in part due to public commentary.
- Public visibility of reviews made reviewers more constructive and diplomatic, despite their anonymity.
- The decision process was more fair in the open model.
- The question of who should be anonymous at what point in the process was far more unresolved than the question of what content should be made visible.
- To a great extent, any fears authors had at the outset of the experiment were not realized.

Open Scholarship and Peer Review

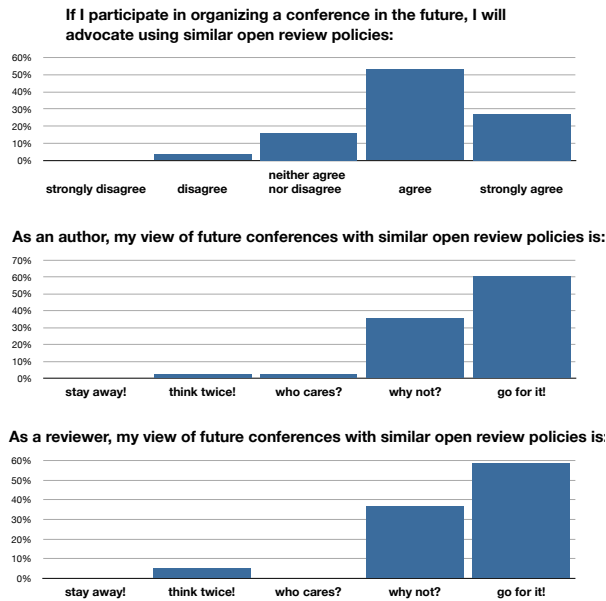



Figure 3. Respondents felt positively about open reviewing at future conferences.

Revisions made in the course of open reviewing increased the quality of individual papers, to a greater extent than I would expect in a closed model.



Insights from commenters other than the anonymous reviewers helped us to improve the paper. (n=24)




Figure 4. The quality of the papers was increased by the open model.




Figure 5. Reviewers felt more pressure to be clear and constructive, knowing that their reviews would be publicly visible.





Figure 6. The decision process was more fair in the open model.

It was important to me that my reviews be anonymous.



When reading papers of which I was neither an author nor a reviewer:




Figure 7. Reviewers strongly preferred to remain anonymous, and some readers withheld comments because they could not be made anonymously.



Figure 8. Author anxiety about open reviewing was reduced as a result of the experience.

The greatest differences of opinion revealed by the survey responses concerned who should be anonymous, and when. In addition to the quantitative results (Fig. 7), respondents made numerous textual comments regarding anonymity, such as:

- “I think the amount of junk papers gets fairly low when the author’s name is associated publicly with them.”
- “Reviewers can be very much biased by the knowledge of the identities of the authors.”
- “Reviewers should not be anonymous, so it would be clear in the future who is making wrong judgements or predictions.” (However, 84% of the reviewers considered anonymity important, largely because they were uncomfortable criticizing friends and because they feared potential social or career repercussions.)
- “Anonymity should be available to all commenters.”
- “I think anonymity (outside the trusted assigned reviewers) would likely lead to flame wars.”
- “I think anonymity should be optional. Sometimes I want to criticize a famous or irritable person without risking fallout. Most of the time I would be fine posting non-anonymously.”

Many other comments provided ideas for future experimentation, expressed appreciation for existing features, or raised interesting concerns.

- Limited ability to cite papers: “Formal publication of the accepted papers in conference proceedings is important and now lacking.”
- Lost credit for work that is rejected but nonetheless publicly available: “The authors disclosed

their ideas and may not be able to publish (in the conventional sense) before others.”

- Perceived dilution of paper quality when rejected papers remain visible: “This makes the review process less helpful as a filter.” “...has the potential to introduce more garbage...”
- Increased accountability of authors and reviewers: “Reviewers cannot ignore rebuttals in this format, and authors cannot ignore errors pointed out by reviewers and commenters.” Conversely, the same effect limited participants’ ability to ignore unproductive comments: “My students, desperately chasing acceptance, did experiments which clearly would show us nothing.”
- One respondent suggested that “a private discussion board for the reviewers would be nice so we can discuss papers before posting the actual review.” On the other hand, several other respondents were concerned about undermining the independence of different reviews: “I think it encourages uniformity in reviews, and causes opinions to be biased by whomever completes their review first.”
- The ability of reviewers to seek out authors of papers they reviewed at the conference was “a huge positive.”

Respondents were quite enthusiastic in their closing comments:

- The best part of the system was “All papers immediately visible so ideas spread faster.”
- The best part of the system was “discussion with people who are not reviewers.”
- “Congratulations for the amazing work and a successful novel approach to reviewing!”

- “The experiment is worth continuing, but the jury is still out.”
- “Keep up the good work. I think it was a good experiment and hopefully more conferences will start adopting similar system.”
- “I find this system so much better, I think it should be used more widely.”
- “Many thanks for the effort in putting up the conference and the site. I believe the model will change the community in the next years.”

Taken together, the survey responses express a great deal of excitement around open reviewing, together with substantial differences of opinion around specific policy choices—particularly regarding anonymity of authors, reviewers, and public commenters. This combination underlines the need for further experimentation in this area, and suggests that such experiments will enjoy continuing community support.

7. Future plans

In addition to the ongoing hosting of discussions for ICLR 2013, we have now hosted two ICML 2013 workshops (Infering 2013 and Peer Review and Publishing Models 2013), as well as the Third International Workshop on Automated Knowledge Base Construction (AKBC) 2013. For these conferences, the organizers chose policies similar to ICLR 2013 with only minor variations. One such variation is that we now host the PDFs directly rather than relying on arXiv for this service, because the 1-2 day delay between submission and posting there was found to be problematic.

We are developing the functionality of OpenReview.net on an ongoing basis and expect to mature beyond the prototype stage within the year. We hope to support reviewing for several conferences and journals throughout 2014. At the same time we hope to host commenting on retrospective papers from past conferences.

Once all of the essential infrastructure is in place, including policy configuration, reviewer assignment, and revision tracking, we look forward to developing more advanced features—such as a reputation system and automated reviewer suggestion—which will benefit from a machine learning approach.

Acknowledgements

We are grateful to the chairs of ICLR 2013—Yoshua Bengio, Yann LeCun, Aaron Courville, Rob Fergus, and Chris Manning—for their helpful feedback and

patience during the initial implementation of OpenReview.net, and to Pamela Burke for discussions of the sociological implications of open peer review. This work was supported in part by the Center for Intelligent Information Retrieval at the University of Massachusetts and in part by Google.

References

- Binfield, Peter, Hoyt, Jason, and O’Reilly, Tim. Peerj, 2012. <http://peerj.com>.
- Budden, Amber E, Tregenza, Tom, Aarssen, Lonnie W, Koricheva, Julia, Leimu, Roosa, and Lortie, Christopher J. Double-blind review favours increased representation of female authors. *Trends in Ecology and Evolution*, 23(1), 2007.
- Coop, Graham, Howie, Bryan, and Pickrell, Joe. Haldane’s sieve, 2012. <http://haldanessieve.org>.
- Demailly, Jean-Pierre, Kloeckner, Benoît, Rolland, Ariane, Berthaud, Christine, Capelli, Laurent, and Magron, Agnès. Episciences project, 2013. <http://episciences.org>.
- Desjardins-Proulx, Philippe, White, Ethan P., Adamson, Joel, Ram, Karthik, Poisot, Timothée, and Gravel, Dominique. The case for open preprints in biology. *PLoS Biology*, (to appear), 2013.
- Ginsparg, Paul. It was twenty years ago today ... 08 2011. <http://arxiv.org/abs/1108.2700>.
- Goodman, Laurie, Edmunds, Scott C, and Basford, Alexandra T. Large and linked in scientific publishing. *GigaScience*, 1(1):1, 2012. doi: 10.1186/2047-217X-1-1.
- Hahnel, Mark. Figshare, 2011. <http://figshare.com>.
- Koonin, Eugene V, Landweber, Laura F, and Lipman, David J. Biology direct: celebrating 7 years of open, published peer review. *Biology Direct*, 8(1):11, 2013. doi: 10.1186/1745-6150-8-11.
- Kriegeskorte, Nikolaus, Walther, Alexander, and Deca, Diana. An emerging consensus for open evaluation: 18 visions for the future of scientific publishing. *Frontiers in Computational Neuroscience*, 6, 2012. doi: 10.3389/fncom.2012.00094.
- Laba, Izabella. The accidental mathematician (blog), 2011-2013. <http://ilaba.wordpress.com/2013/02/09> (and other posts).

- Lawrence, Rebecca and Tracz, Vitek. F1000research, 2012. <http://f1000research.com>.
- LeCun, Yann. A new publishing model in computer science. 2009. <http://yann.lecun.com/ex/pamphlets>.
- Merton, R. K. The matthew effect in science: The reward and communication systems of science are considered. *Science*, 159(3810):56–63, Jan 1968. doi: 10.1126/science.159.3810.56.
- Mudunuri, Shashi and Collier, Keith. Rubriq, 2013. <http://www.rubriq.com>.
- Nature. Nature’s sexism. *Nature*, 491(7425):495–495, Nov 2012. doi: 10.1038/491495a.
- Schekman, Randy. elife, 2013. <http://www.elifesciences.org>.
- Sinha, R., Boutelle, J., and Ranjan, A. Slideshare, 2006. <http://www.slideshare.net>.
- Smith, Richard. Opening up bmj peer review: A beginning that should lead to complete transparency. *BMJ: British Medical Journal*, 318(7175):4, 1999.
- van Rooyen, S., Godlee, F., Evans, S., Black, N., and Smith, R. Effect of open peer review on quality of reviews and on reviewers recommendations: a randomised trial. *BMJ*, 318(7175):23–27, Jan 1999. doi: 10.1136/bmj.318.7175.23.
- Vision, Todd J. Open data and the social contract of scientific publishing. *BioScience*, 60(5):330–331, May 2010. doi: 10.1525/bio.2010.60.5.2.
- Walker, Ian and Holt, Nigel. Philica, 2006. <http://philica.com>.
- Walsh, E., Rooney, M., Appleby, L., and Wilkinson, G. Open peer review: a randomised controlled trial. *The British Journal of Psychiatry*, 176(1):47–51, Jan 2000. doi: 10.1192/bjp.176.1.47.
- Wasserman, Larry. A world without referees, 2012. <http://www.stat.cmu.edu/~larry/Peer-Review.pdf>.