Report to NIPS board on New Publication Models for the NIPS Community

This report has been drafted by the following committee named by the NIPS board at NIPS'2009: Yoshua Bengio, Tom Dietterich, Daphne Koller, John Lafferty, John Platt, Laurence Saul and Chris Williams. It benefited from the feedback of many others, including several other members of the NIPS board as well as colleagues from our community (in particular Yann LeCun and Zoubin Ghahramani), following the public debate held December 8th, 2009, Vancouver (during the NIPS conference).

Objectives: What we think could be improved

* **Improve the quality of the reviewing and selection process**. The quality of the reviewing process in conferences (especially large ones like NIPS where most of the community is involved) is substantially less than with journals. Efforts to introduce feedback/revision mechanisms have helped, but not enough, and it is difficult to improve the selection of reviewers (even with bidding mechanisms) simply because there is a '*reviewer crunch*' when the whole community is reviewing 1200 papers all at once. As pointed out by Mike Jordan, the NIPS reviewing pool "includes a significant number of people who aren't able to review very broadly, and they end up reviewing papers that they aren't really able to review. It's human nature to not admit weakness, so you'll get these overly negative reviews that don't actually make much connection with the ideas in the paper. Great consternation about the lack of comparison to 45 different methods and the lack of experiments on 45 different datasets because they can't think of anything else to say." See also this recent discussion and paper on the overly strong effect of poor reviewers on the quality of the paper publication/selection process.

* **Improve the rate of scientific progress**. Quoting from Yann LeCun's pamphlet on new publishing models for CS: "Our current publication system should be redesigned to maximize the rate of progress in our field. This means accelerating the speed at which new ideas and results are exchanged, disseminated, and evaluated. This also means minimizing the amount of time each of us spends evaluating other people's work through reviewing and sifting through the literature. A major issue is that our current system, with its emphasis on highly-selective conferences, is highly biased against innovative ideas and favors incremental tweaks on well-established methods. Ideas that turn out to be highly influential are sometimes held up for months (if not years) in reviewing purgatory, particularly if they require several years to come to maturity (there are a few famous examples, mentioned). The friction in our publication system is slowing the progress of our field. It makes progress incremental. And it makes our conferences somewhat boring."

* **Harmonize with other communities, where journals have greater value than conference proceedings**. See Viewpoint: Journals for Certification, Conferences for Rapid Dissemination,

by Joe Halpern and David Parkes, for a nice discussion of the needs for publicity and certification and the difference between the CS community and others. That difference hurts us both in grant/promotion committees and in collaborative work with researchers from other communities. Halpern and Parkes propose an alternative model for CS, in which a public archive serves as the initial publication medium, followed by submission to fast turn-around journals (which exist in other communities), followed by presentation in a conference. In this new model (already taking shape in some CS sub-communities), conference and journals share and re-use reviews and reviewers (on the same or revised paper) to improve reviewing efficiency.

Long-term Plan: where we propose to go

* **Allow archives to be the first publication mode, before conferences, to make results available early.** A primary mode for publication is a place like arXiv, where people who want to publish their results earlier (without certification) put them online and then wait for certification to take place.

* **Share reviewers and reviews between different certification authorities**. Halpern followed this policy as editor-in-chief of the Journal of the ACM, and it is common in other communities outside of CS; authors of rejected papers could request reviews to be passed on to the journal of their choice. Sharing reviews and reviewers (of the same paper) between journals and conferences would avoid unnecessary re-reviewing, especially if a higher-quality reviewing process has already been done for a journal, and then the paper is submitted to a conference. It is much more efficient to ask a reviewer who has already reviewed version 1 of some work to also review version 2, even if this is submitted in a different review. This should be balanced by the possible introduction of new reviewers, but would greatly reduce the overall reviewing load.

* **Formalize comments posted online and incorporate them in the editorial decisions, in a centralized location grouping them along with revisions and certifications**. We would like a flexible commenting mechanism, so interesting papers will get read and commented on, at which point journals would want to get credit for them and would be likely to review them. The papers that no-one reads are less likely to get reviewed, although an author can still submit to a journal, at which point the journal has the 2-week window to decide whether to review it or not (during which authors are not allowed to submit their paper to another reviewing entity). Once a paper has received a certification ("been accepted for publication") by a reviewing entity, the paper (possibly in a somewhat different version) could still receive another certification by another entity, where it makes sense (e.g., by a journal targeting a different community, or for public presentation in a conference). Ideally, the different versions of a paper (first archive, conference, journal) should all be centralized, along with the comments and certifications.

* **Improve reviewer motivation**. It might be easier to get reviewers for a paper that was posted online and raised interest, because potentially commentators could be asked to review the paper. To facilitate that, we could allow three levels of anonymity for comments or parts of

comments - fully anonymous, anonymous to the public but not to an editorial board, and fully visible. To motivate fully visible reviews, these should be citable, in particular in the revision of the paper that they comment on.

* **The reviewing process should be, as much as possible, evenly distributed across the year and the same (or revised) paper should not reviewed again and again by different conferences and journals.** Papers that were recently posted online can also be submitted to a conference. Reviews can be shared between conferences and journals (reviewing entity), at the discretion of the journal and conference boards (with bilateral or multilateral agreements). Papers can be posted online and submitted to a conference (e.g., NIPS) long before the conference date, to avoid the logjam of getting 1200 papers reviewed at the same time. Papers that have been reviewed by a journal (which is willing to share reviews with other reviewing entities) can be submitted to a conference, and the area chair can decide what to do regarding reviews (e.g., gather 0, 1 or 2 more reviews, or possibly ask a previous reviewer to re-review the new version) before taking a decision, taking the journal reviews and editorial decision into consideration. The reverse is also possible (a paper submitted at the conference can be later submitted, along with its reviews and editorial decision, to a journal), although the advantage may be lower. Papers that have been reviewed and rejected can be submitted to another reviewing entity; but when reviews are shared, the previous reviews and the extent to which the paper has been revised can play a role in the reviewing entity's decision on whether to review the (revised) paper.

* **Compromise between fairness (double-blind reviewing) and quick accessibility (archive posting)**. To address concerns about double-blind reviews, we could allow for semi-anonymous postings on arXiv while the paper is under review - the paper is first shown to a viewer without author names, and then the viewer has to actively click to see the author names. For reviewers searching for relevant literature, we'd have to trust them not to deliberately seek out author names by clicking. Other viewers who just want to cite the work can click and get the author names. Reviewers will have to login to look at the authors' names, and their looking at it will be recorded during the semi-anonymous period.

Proposed Short-term Actions

* NIPS makes review(er)-sharing deals with JMLR, Neural Computation, MLJ, and possibly other relevant journals; it encourages the boards of ICML, AISTATS, COLT, and UAI to do the same, possibly starting a growing network of similar deals. The deals allow other conference and journals to know who reviewed what, so that at least that can help re-use reviewers that have already spent time on some version of a paper, and possibly re-use reviews.

* Make it clear to authors that archive submissions prior to NIPS are OK, in spite of the double-blind reviewing. Make it clear to reviewers that they are expected to first read the paper and make up their mind without actively seeking the author identities.

* Discuss with the journals in our field to encourage them to create or improve a ‘fast track’, with quick turnaround, starting with a 2-week period during which the submissions are filtered using a single reviewer or editor. These fast tracks could have page limits, as suggested by Halpern & Parker, and implemented by Neural Computation and many top journals in other fields.

Proposed Long-term Actions

* Set-up an overlay on top of arXiv which will add the following features:
  - comments (with different degrees of anonymity), possibly citable (in the same arXiv)
  - certifications (pointers to the conference / journal versions)
  - semi-anonymized manuscripts during a review process
  - alerts and karma system? (see Yann LeCun’s suggestion).
* Study the possibility of creating a PLoS-1-like fast-evaluation venue that certifies correctness rather than impact, i.e., something in between arXiv and a regular journal or conference proceedings.

Discussions

* We should encourage authors to first publish in arXiv, then journals, then in conferences to improve the distribution of reviewing across the year, yield better quality reviews, and harmonize with other communities. The conference --> journal route should be possible but not encouraged by the system’s incentives. Workshops can be used as a dissemination medium right after the arXiv paper has been posted. It might be easier not to accept NIPS submissions during the year (to address the ‘reviewer crunch’), encouraging people to go the Journal route first. The disadvantage is that we don’t get as much of a relief (not as much spreading across the year), especially initially, before people take the habit of first submitting to journals and while journal reviewing delays remain high. This decision does not change the long-term target state, but it is not clear what would be the best way of getting there.

* We do not propose that NIPS becomes a journal. As pointed out by Mike Jordan, this might create a pull for a large number of low-quality papers. There are already enough journals. What we need is to reform the way conferences and journals work together to achieve the best progress in our community.

* Review-sharing deals with a set of journals should not prevent or discourage NIPS publication from authors involved in other communities, through direct submission to NIPS. The proposed system should not prevent people from submitting a long version of their NIPS paper to established journals (e.g. in statistics) where publication delays are long, but we should strive to make review-sharing deals with these journals as well.

* Certification could come with an explicit or implicit duty to review. Halpern & Parkes propose that “for every paper that is reviewed by a journal, some author of that paper must be available to review another paper (Crowcroft et al. [2] make a similar point). This approach is in fact

used in the B.E. Journal of Theoretical Economics, which also has financial penalties for late reviews." Zoubin Ghahramani proposed a kind of journal-conference-college hybrid in which reviews earn points. This proposal also shares some of the objectives discussed here.