

Journal of Machine Learning Research, MIT Press

<http://mitpress.mit.edu/JMLR>

The *Journal of Machine Learning Research* was founded in 2001, in its words “as an international forum for the electronic and paper publication of high-quality scholarly articles in all areas of machine learning.” Machine learning is a fascinating and important area of computer science: basically, can computers learn? It is both a theoretical area and a very practical area, where computer programs are developed and run on data to see what they can learn. There are enormous and widespread practical applications for machine learning, from labour negotiation, medical treatment, agriculture, to generally making computers more ‘intelligent.’ The particular slant of the journal will make it interesting for theoreticians, biologists and psychologists, who of course are interested in how animals and humans learn and what the theoretical limits to learning are. There are even applications in machine learning for countering terrorism. And as the world wide web fills up with vast amounts of unstructured information, we need all the help we can get to learn how to use it effectively. It says something about the huge relevance of machine learning that one managing editor of this journal works at Google and the other has published papers on financial markets.

The bulk of the journal’s papers are devoted to discussing and evaluating learning methods. I was interested to see how many of the ideas talked about in the journal actually worked, because that’s really the whole point. So, as the journal is available on-line, I looked at every paper and then emailed the authors to ask them about their ideas.

After a few weeks, I had 99 replies (plus 10 broken replies from authors who had moved and not left a forwarding email address).

I had asked whether the system described in the paper was available. Of course, some papers were theoretical; I had 5 replies saying my question was irrelevant. Of the remaining 94 replies, 76 talked about systems, and of these only 25 said the systems were available if I wanted them. A further 29 said their systems were unavailable. Of these, 27 admitted their systems were broken, private or incomplete in some way, and 8 said their systems were commercial confidential —of these, 4 were also broken in some way. And three people want jobs.

Here are some quotes from different replies I got: “The system is a research prototype developed in my group, and is not appropriate for public dissemination.” Or “We never really built a system to do this.” Or “I’ve a set of perl scripts which [...] are a bit of a mess and really should be cleaned-up sometime.” Or “There is no user friendly software package. I would not mind sharing the source code of either version, provided you promise to give credit, not to distribute it, and so on.” Or “Unfortunately, I do not have the system in a state where I can give it away right now.”

Further quotes are quite telling: “The implementations we had were very much ‘research code’, and [are] not suitable for public consumption” and “we are indeed developing code along these lines but it is research code at present and not at the polish for distribution. So we will not be able to share that with you at this time.” Or “The

"system" consists of different parts which rely on training data. I do not have tied the different parts together in way that makes it easy to release it to others. The fact that the training data is not mine [...] makes that hard as well." With this, "You can download the source code [...] We have just completed a re-write [...]" the authors have at least made their system available, but it's no longer the one their paper was based on.

Many replied saying they were continuing to write more papers: "Sorry that we don't have the system available to public yet. We did a simple extension [...] a paper should come out soon in *Neural Computation* [another journal]. You can also get a short version of the paper from *NIPS2002* [a conference]." So here we have a system, which is not available to scholarly scrutiny, that nevertheless is the basis of at least three papers. Some replied that it was someone else's work, as in "This is the core topic of the PhD research of my student, [...], whom I cc'ed on this email. I suggest you correspond with him directly regarding system availability." Well, it's nice that the supervisor credits his PhD student, but in what sense has this student contributed to effective world knowledge if the developments are inaccessible, as they apparently are, even to his supervisor? I mean, machine learning is *supposed* to work!

As I wrote elsewhere (see <http://www.ucl.ac.uk/harold/warp>), there is a very, fine line between publishing stuff that does not work or that cannot be got to work because it is described vaguely, and publishing stuff that *never* worked, or never worked quite so well as it has been written up, let alone publishing stuff without any description of how it works at all. This is the computer science equivalent of deleting experimental data. Indeed one reply to me said, "We don't have the data ready to be published." In the established scientific disciplines this sort of behaviour would lead to outrage

Generally, there is little regard for conventional scientific qualities: reproducibility, testability, openness, and availability of data, methods (or programs). The low level of scholarly background for the papers in the journal is ironic since it has an excellent web site that could easily make resources available to the wider community.

This unaccountability is a sad but standard problem in computer science — so widely acceptable, in fact, that I feel I'm being hard on the journal, its editorial board and its authors. It isn't the journal's fault! Indeed, promisingly, the *Journal of Machine Learning Research* tries to encourage authors to add electronic appendices with source code, data, even Quicktime demonstrations: anything, as the journal puts it, that will make life easier (or more interesting) for readers and researchers who follow in the authors' footsteps. Machine learning will change our uses of computers dramatically, so let's hope the journal achieves its goals with more and more success.

Harold Thimbleby is Director of UCLIC, the UCL Interaction Centre, and Gresham Professor of Geometry. www.ucl.ac.uk/harold