



Moshe Y. Vardi

DOI:10.1145/2753507

Incentivizing Quality and Impact in Computing Research

Over the past few years, the computing-research community has been conducting a public conversation on its publication culture. Much of that conversation has taken place

in the pages of *Communications*. (See <http://cra.org/scholarlypub/>.) The underlying issue is that while computing research has been widely successful in developing fundamental results and insights, having a deep impact on life and society, and influencing almost all scholarly fields, its publication culture has developed certain anomalies that are not conducive to the future success of the field. A major anomaly is the reliance of the fields on conferences as the chief vehicle for scholarly publications.

While the discussion of the computing-research publication culture has led to general recognition that the “system is suboptimal,” developing consensus on how the system should be changed has proven to be exceedingly hard. A key reason for this difficulty is the fact the publication culture does not only establish norms for how research results should be published, it also creates expectations on how researchers should be evaluated. These publication norms and research-evaluation expectations are complementary and mutually enforcing. It is difficult to tell junior researchers to change their publication habits, if these habits have been optimized to improve their prospects of being hired and promoted.

The Computing Research Association (CRA) has now addressed this issue head-on in its new Best Practice Memo: “Incentivizing Quality and Impact: Evaluating Scholarship in Hiring Tenure, and Promotion,” by Batya Friedman and Fred B. Schneider (see <http://cra.org/resources/bp-memos/>). This memo may

be a game changer. By advising research organizations to focus on quality and impact, the memo aims at changing the incentive system and, consequently, at changing behavior.

The key observation underlying the memo is that we have slid down the slippery path of using quantity as a proxy for quality. When I completed my doctorate (a long time ago) I was able to list four publications on my CV. Today, it is not uncommon to see fresh Ph.D.’s with 20 and even 30 publications. In the 1980s, serving on a single program committee per year was a respectable sign of professional activity. Today, researchers feel that unless they serve on at least five, or even 10, program committees per year, they would be considered professionally inactive. The reality is that evaluating quality and impact is difficult, while “counting beans” is easy. But bean counting leads to inflation—if 10 papers are better than five, then surely 15 papers are better than 10!

But scholarly inflation has been quite detrimental to computing research. While paradigm-changing research is highly celebrated, normal scientific progress proceeds mainly via careful accumulation of facts, theories, techniques, and methods. The memo argues that the field benefits when researchers carefully build on each other’s work, via discussions of methods, comparison with related work, inclusion of supporting material, and the like. But the inflationary pressure to publish more and more encourages speed and brevity, rather than

careful scholarship. Indeed, academic folklore has invented the term LPU, for “least publishable unit,” suggesting that optimizing one’s bibliography for quantity rather than for quality has become common practice.

To cut the Gordian knot of mutually reinforcing norms and expectations, the memo advises hiring units to focus on quality and impact and pay little attention to numbers. For junior researchers, hiring decisions should be based not on their number of publications, but on the quality of their top one or two publications. For tenure candidates, decisions should be based on the quality and impact of their top three-to-five publications.

Focusing on quality rather than quantity should apply to other areas as well. We should not be impressed by large research grants, but ask what the actual yield of the funded projects has been. We should ignore the *h*-index, whose many flaws have been widely discussed, and use human judgment to evaluate quality and impact. And, of course, we should pay no heed to institutional rankings, which effectively let newspapers establish our value system.

Changing culture, including norms and expectations, is exceedingly difficult, but the CRA memo is a very promising first step. As a second step, I suggest a statement signed by leading computing-research organizations promising to adopt the memo as the basis for their own hiring and promotion practices. Such a statement would send a strong signal to the computing-research community that change is under way!

Follow me on Facebook, Google+, and Twitter.

Moshe Y. Vardi, EDITOR-IN-CHIEF

Copyright held by author.