Some Super-Classics of Artificial Intelligence?

Jon Doyle

Laboratory for Computer Science
Massachusetts Institute of Technology
545 Technology Square
Cambridge, MA 02139
doyle@lcs.mit.edu

June 3, 1993
Revision of June 14, 1993

Abstract: *Artificial Intelligence* serves the field well with its recent retrospective on its most-cited papers, but the method used for identifying these "super-classics" exhibits an unfortunate sensitivity to demographic changes. In consequence, the retrospective slights influential papers of the sparsely-populated 1970's in favor of influential papers from the crowded 1980's.

The editors of *Artificial Intelligence* deserve everyone's applause for their inspiration and effort in soliciting retrospectives from authors of highly-cited *AIJ* papers from the first 50 volumes, recently published as numbers 1–2 of volume 59. These extremely interesting retrospectives provide a valuable resource to students and scholars alike. The volume recognizes great achievements, and constitutes a great testimonial to the productivity of our field.

Alas, few human efforts lack all blemishes. While one can safely and justly ignore most blemishes of great achievements like this volume, the modern tendency to glorify the new and ignore the historical perspective makes me concerned that this volume may encourage a seriously mistaken conclusion about the set of papers it identifies. But with this blemish noted, we can go on to celebrate this retrospective for a long time.

The selection of the papers for retrospective comment stems from a citation analysis of the papers from the first 50 volumes of *AIJ*. The editors examined two different citation indicies to produce two lists of the 50 *AIJ* papers receiving the most citations in the first five years after publication, and then combined these lists to yield a list of 69 most-cited papers. Citation analyses always pose major difficulties due to the great inaccuracies common in citations; see Schatz [5] for a discussion of these problems in the context of identifying some "super-classics" of twentieth-century mathematics, and Menzies [4] for a similar examination of AI papers from the early 1980's. But even assuming complete accuracy in the counting of references to *AIJ* papers, the list of 69 most-cited papers identified in this volume suffers a severe problem simply due to the changing demographics of the AI field (a consideration previously noted by Menzies).

One naturally interprets a collection of most-cited papers as representing the most influential papers from among the population sampled. One expects some anomalies (in the 1970's, at least, a sizable fraction of papers cited LISP manuals), but the general validity of this interpretation provides the main reason for conducting such exercises. In the case of the current retrospective, however, the citation statistics vividly reflect another factor beyond the influence of the papers, namely the growth in the size of the field. Cutting-edge papers published in the 1980's, when the field was expanding extremely rapidly were read by a much larger audience than papers published in the 1970's. One thus expects and observes a strong bias in the papers winning citation counts toward the end of the sampling period, when the field's population reached its maximum. Put another way, even 1980's

papers deriviative from and less influential than their sources in papers of the 1970's would garner larger citation counts than their sources simply because so many more people were writing papers in the later years, even if a much smaller fraction of the field read the later papers.

Table 1 lists the number of super-classics (if I may call them that) from the *AIJ* list according to year of publication. The first column lists the year; the second the number of papers from the *AIJ* list published in that year; and the third the two-year total of papers published that year and the previous year (as an approximation to the integral of the distribution). As one can

Year Y	#(Y)	#(Y)+#(Y-1)
1970	0	0
1971	1	1
1972	0	1
1973	0	0
1974	1	1
1975	1	2
1976	0	1
1977	5	5
1978	2	7
1979	1	3
1980	6	7
1981	10	16
1982	6	16
1983	7	13
1984	11	18
1985	9	20
1986	9	18

Table 1: Most-cited *AIJ* articles by year and by biennium

plainly see, while the late 1970's produced some years with several winners, their numbers hardly compare with those of papers published in the mid-1980's, which have the largest number of person-years in the citation period (assuming this ended in 1991). All told, 11 papers from the 1970's and 39 from the 1980's make the list.

Specific examples illustrate the difficulty even more vividly. While other early papers would have also made the list had the size of the field been constant, the absence of the seminal and (even at the time) widely influential papers of Fikes, Hart, and Nilsson [2], Knuth and Moore [3], and Bledsoe [1] seems astonishing until one recognizes the demographic changes influencing the citation counts. The absence of such papers astonishes even more when one recalls that the excitement engendered by such papers constituted a significant part of the reason for the later demographic growth.

One can also suspect other temporal factors as influences on the shape of the *AIJ* list, such as the tendency in the 1970's to cite some technical report versions of papers as readily as their appearances in journals (which probably diminished the citation counts of some *AIJ* papers). But the five-year window of citation counting poses a deeper problem. One can reasonably argue that the true classics of a field continue to receive citations beyond the first few years after their appearance, while most papers receive citations mainly in the first years following publication, if at all. If one takes this criterion for identifying the most influential papers, the statistical method used to construct the *AIJ* list counts citations on the wrong side of the citation-date cutoff.

The *AIJ* list identifies papers well worth our celebration, in spite of any methodological blemishes. But the compilation does some injustice to other papers which, in their day, exerted at least as great an influence on the field as those recognized in the retrospective. We should celebrate especially the earlier papers which, in spite of the demographic barriers posed by the statistical methods employed to find them, still made the *AIJ* list (and perhaps should marvel at the papers of the 1977 *annus mirabilis*).

References

[1] W. W. Bledsoe. Non-resolution theorem proving. *Artificial Intelligence*, 9:1–35, 1977.

[2] R. E. Fikes, P. E. Hart, and N. J. Nilsson. Learning and executing generalized robot plans. *Artificial Intelligence*, 3:251–288, 1972.

[3] D. E. Knuth and R. N. Moore. An analysis of alpha-beta pruning. *Artificial Intelligence*, 6:293–326, 1975.

[4] T. Menzies. An investigation of AI and expert systems literature: 1980-1984. *AI Magazine*, 10(2):53–61, 1989.

[5] J. A. Schatz. Some super-classics of mathematics. *Notices of the American Mathematical Society*, 18:723–725, 1971.