AI Magazine Volume 10 Number 2 (1989) (© AAAI)

An Investigation of AI and Expert Systems Literature: 1980–1984

Tim Menzies

This article records the results of an experiment in which a survey of AI and expert systems (ES) literature was attempted using Science Citation Indexes The survey identified a sample of authors and institutions that have had a significant impact on the historical development of AI and ES However, it also identified several glaring problems with using Science Citation Indexes as a method of comprehensively studying a body of scientific research Accordingly, the reader is cautioned against using the results presented here to conclude that author A is a better or worse AI researcher than author B

0738-4602/89/$3 50 © 1989 AAAI

This article describes an attempt to systematically survey AI and expert systems (ES) literature. This survey was conducted for technical and historical reasons: First, I work in the commercial AI industry and was worried about missing significant intellectual contributions to my work. Second, this work was intended to test the thesis that there is a coherent body of study called cognitive science. If a new scientific discipline has emerged that is a fusion of psychology, computer science, linguistics, mathematics, philosophy, and neuroscience (as claimed by Gardner 1985), then there should be some evidence of this new discipline in the pattern of scientific publications and researcher biographies. For example, a paper about AI could cite a psychology paper, or a graduate in a mathematics department could migrate into the linguistics field.

I have been informally conducting this survey for the last three years. Through a variety of informal techniques—such as (1) browsing through the *Computer and Control Abstracts*, (2) looking up the papers of the authors that anecdotally are known to be good authors, and (3) inspecting the reference lists of interesting papers—I have collected several hundred papers. The bulk of these papers are not worth a second read. Further, such an informally collected group of papers could not serve as the basis for a historical or sociological analysis of the field. Obviously, a better method of finding worthwhile papers is required, hence this article.

The basic technique used in this survey was to defer reading any paper until some evidence had been collected that suggested the paper was worth reading. The results presented here show what information could be gathered from sources of information on published papers and publishing authors.

The following methodology was used: First, review articles were collected and their reference lists studied Second, *Science Citation Indexes* (SCI) were used to check the numbers of citations for various authors. (SCI records the citations received by the papers published in a large group of journals; hence, they allow for the forward searching of the literature. A paper can be linked to subsequent work when this subsequent work cites it) Third, the date and institution of origin for each author were recorded. Fourth, biographical data were sought on the studied authors.

It was hoped that this methodology would identify any patterns of author migration between institutions and disciplines and that by combining all this information, a list of significant institutions, dates, and authors in this field would emerge. However, this goal was only partially achieved: (1) Finding review articles proved to be problematic because the study area was inadequately covered by SCI (see Finding Review Articles). (2) Obtaining the number of online SCI entries required for a thorough analysis proved to be beyond the resources of this project (see Citation Analysis Is an Expensive Technique). (3) Without adequate SCI access, an alternative, less rigorous method had to be developed. The method, called *namedropping*, used review articles to identify commonly referenced authors (see A Heuristic Literature Search). (4) With the references of seven review articles as a start, a list of 585 AI and ES authors was compiled. Forty-four of these authors were found to be referred to in the majority (more than three out of seven) of the review articles (see Namedropping Results). (5) This group of 44 was termed the outstanding authors and was singled out for more detailed study. The 4762 SCI entries from the 1980–1984 cumulative index for the outstanding authors were then collected by hand. The SCI entries for 17 of these 44 authors (17/44 = 38%) proved to be incorrect, missing, or confusing. The usable data from SCI were then collected (see Namedropping Using SCI). (6) Using SCI entries, it was possible to confirm that the outstanding authors identified using namedropping are indeed significant authors (see Validating Namedropping). (7) The *American Men and Women of Science* and *Who's Who in Science* entries for the outstanding authors were collected. These entries proved to be inadequate for identifying migration patterns of authors between institutions and disciplines (see Searching for Biographical Data). (8) As a result of this work, the author analyzed the merits of citation analysis as a method of reviewing scientific work (see Advantages of Using SCI and Disadvantages of Using SCI).

Note that the data presented here were collected by hand from thousands of citations. It is more than possible that minor clerical errors resulted in a significant number of AI and ES authors not receiving appropriate credit from this study. In my defense, literature searches such as the one described here must be informally done every day by researchers attempting to come to terms with the mountain of paper being published on AI and ES. If injustices are committed in the following pages, these injustices must be occurring constantly. Researchers wishing to receive appropriate recognition for their work should look to altering our current methods for assessing this work.

Finding Review Articles

Initially, I intended to begin with review papers taken from the outstanding AI and ES journals. "Outstanding" was determined on the basis of the total number of citations for each journal.

The microfiche subject index at the University of New South Wales library lists seven journals under the heading of artificial intelligence. On consulting the source index of the 1980–1984 SCI, it was discovered that only one of these journals, *Artificial Intelligence,* belonged to the group of core journals used by SCI. The 1982 to 1986 issues of this journal were then searched for review articles. Of the two articles found (Bobrow and Hayes 1985, Newell 1982), only one had a list of references (the Bobrow and Hayes paper was more a discussion in which various AI and ES researchers wrote their replies to 10 questions posed by the editors of the article).

Review articles from any anecdotal source were then searched for. This search was surprisingly difficult; the AI and ES field appears to be marked by its lack of internal review. In the end, only nine review articles could be located, eight of which were used in this study.

All the subsequent work outlined in this article was performed on the two AI review articles from *Artificial Intelligence,* three ES review articles (Duda and Buchanan 1983; Ramsey, Reggia, and Nau 1986; Sandell and Bourne 1985), one introductory AI text (Rich 1983), and one introductory ES text (Hayes-Roth, Waterman, and Lenat 1983).

A concern with using anecdotal sources was that these articles might not represent the conventional AI and ES wisdom. To test this possibility, one well-known, introductory AI text was held in reserve (Winston 1984) and tested later (see Effect of Initial Selection of Reference Articles).

Citation Analysis Is an Expensive Technique

The names referred to by the seven review articles (either in the reference list or, in the case of the Bobrow and Hayes article, those AI and ES researchers who replied to the authors' 10 questions) were to have been used as pointers into SCI. A computer-based analysis of the citation links between papers written by these authors was planned.

This plan was not followed because the number of citations that would have been required was prohibitive. In one SCI-based study, Small and Sweeney (1985) used 3.9 million uniquely cited items and 7.6 million citing-cited pairs for their analysis. A humbler study by Beck, Pyle, and Lusted (1984) apparently used 3000 to 4000 citations. (I arrived at this figure from table 1 in Beck, Pyle, and Lusted, p. 454. Their study used the reference lists from 61 documents. Table 1 shows that the 20 most highly cited documents were cited a total of 2167 times. Based on the assumption that the remaining 41 articles were cited 30 to 45 times each, then the Beck, Pyle, and Lusted study used 3397 to 4012 citations.)

The funding available for this study was insufficient to download thousands of citations from online bibliographic databases. The alternative would be to manually key in the SCI entries. However, this alternative was impractical because even a small study (for example, Beck, Pyle, and Lusted) requires the entry of thousands of references.

A Heuristic Literature Search

A heuristic called namedropping was developed to minimize the work load. Tallies were kept of how often a surname was mentioned in the review articles. If an article mentioned a particular surname any number of times (greater than or equal to one), then this surname scored one more "namedrop." The amount of data that has to be recorded is minimal. A small table can be drawn up for each author with the columns labeled 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, and 10+. The SCI entries are read, and the number of citations received by each paper is added up. If a paper receives N citations, a tick is placed in the column labeled N. If a paper receives 10 or more citations, the actual number of citations is recorded in the 10+ column. Afterwards, the total number of citations in each column can be calculated by multiplying the number of ticks in each column by the number at the top of the column (seven ticks in the 2 column is 14 citations). The number of citations for each author is then the total of all the columns. Using this technique, I recorded 4762 ticks for 44 authors in about 16 hours.

Namedropping Results

Namedropping was performed on the initial group of seven articles. Forty-four outstanding authors were identified; that is, 44 names were mentioned in at least four of the seven articles.

The seven source references were found to drop 585 names (see table 1).

The value of heuristics is that they serve to restrict a search space Through namedropping, we could restrict the study from 585 names to however many we have time to study The cutoff level was chosen as N = 4 namedrops because this level provided the study with a group large enough to provide something of an overview but small enough to be manageable. The 44 names that were mentioned by four or more of the review articles are shown in table 2.

Citation counts on the outstanding authors were collected from the SCI 1980–1984 cumulative index. Before these counts are presented, I should digress and review the advantages and disadvantages of using the SCI indexes

N	Number of Surnames With N Name Drop	Cumulative Total
7	1	1
6	5	6
5	16	22
4	22	44
3	38	82
2	90	172
1	413	585

Table 1: Number of Review Articles that Refer to a Surname

7	6	5	4
Minsky	Hayes-Roth (B. or F.)	Aikins	Bobrow
	Nilsson	Amarel	Cohen
	Simon	Brachman	Doyle
	Shortliffe	Buchanan	Erman
		Davis	Fagan
		Duda	Gaschnig
		Feigenbaum	Hart
		Kulikowski	King
		Lenat	Kunz
		McDermott	McCarthy
		Newell	Michie
		Pauker	Miller
		Scott	Mitchell
		VanMelle	Nii
		Waterman	Pople
		Weiss	Riesbeck
			Stefik
			Swartout
			Szolovits
			Winograd
			Winston
			Zadeh

Table 2: Number of Times Surnames were Mentioned in the Review Articles. (Only those mentioned four or more times are represented.)

Advantages of Using SCI

It would be useful to be able to automate literature reviews. Such a method would save considerable time when studying a research field. Computer analysis of citation indexes has been proposed as a method of implementing such an automatic review.

Eugene Garfield, director of the Institute for Scientific Information, the group that publishes *Science Citation Indexes,* is most impressed by the promise of computer-assisted literature analysis and dreams of the day when . . .

> a historian or sociometrist will be able to sit before a computer console and specify some starting point—a person, a word, a citation, or a place. He will then ask the computer to display a list of pertinent papers. The computer will respond by drawing a historical road map which will show not merely a list of papers and books, but also a graphical approximation of the history of that subject (Garfield 1970, p. 670).

These indexes have been used by various authorities to generate seemingly impressive maps of the "structure of science " Maps of the interconnections between papers can be produced using large-scale citation analysis. Lines showing citation links can be used to identify significant journals A paper, or a group of papers, that sprouts many lines is obviously an important paper in the literature Greenlee and Small (1986) use the term "cocitation frequencies" to define a distance between papers Each possible combination of two papers from a reference list is a cocitation. During a cocitation analysis, all the cocitations from a large group of papers are generated and sorted alphabetically The number of repeated cocitations is recorded. Documents are then grouped based on their cocitation strengths; the greater the cocitation strength, the closer the grouping When the paper's strength is greater than some threshold, then the papers are grouped into a cluster of papers. Cocitation strengths can then be calculated between each cluster. It is claimed that this technique "identifies coherent areas of scientific research by tapping scientist's own foci of attention within the scientific

Name	Name-drop Group	Total No Cit	Total. No. Papers	Average Citations per paper	Institution
Pauker S.G.	5	228	26	8 7692	Tufts University, Boston
Shortliffe E H	6	210	31	6 7742	Stanford University
Zadeh L A.	4	929	158	5.8797	University of California, Berkeley
Szolovits P.	4	46	10	4 6	Tufts University, Boston
Duda R O	5	458	104	4.4038	Fairchild, Palo Alto
Winograd T	4	138	39	3 5385	Stanford University
Newell A	5	300	90	3.3333	Carnegie Mellon
Hayes-Roth B	6	36	11	3 2727	?
Simon H.A.	6	539	170	3 1706	Carnegie Mellon
Bobrow D G	4	139	45	3 0889	Xerox Corp , Palo Alto
Winston P & P H.	4	182	61	2 9836	MIT Cambridge, Mass
Minsky M	7	230	78	2 9487	?
Erman L.D	4	43	16	2 6875	University of Southern California
Nilsson N	6	68	26	2 6154	?
Waterman D & D A	5	77	30	2.5667	Syracuse University, New York
Pople H & H E	4	64	25	2.56	University of Pittsburgh
Mitchell T M	4	30	12	2.5	?
Nii H & H P	4	37	16	2.3125	Stanford University
Cohen P R	4	16	7	2.2857	?
McCarthy J	4	254	123	2.065	multiple affiliations
Stefik M & M.J	4	53	26	2.0385	Stanford University
Buchanan B G.	5	52	26	2.0	Stanford University
Riesbeck C & C K.	4	17	9	1 8889	?
Fagan L M.	4	13	7	1 8571	Stanford University
Aikins J	5	29	16	1 8125	Stanford University
Hayes-Roth F	6	68	38	1.7895	Teknowledge, Palo Alto
Feigenbaum E & E A	5	87	49	1 7755	Stanford University
Kulikowski C.A	5	23	13	1 7692	Rutgers University, New Jersey
Hart P.E	4	61	35	1 7429	?
Michie D	4	92	53	1 7358	University of Edinburgh
VanMelle W	5	25	15	1 6667	?
Lenat D & D A & D R	5	52	33	1.5758	Stanford University
Gaschnig J & J.G	4	32	21	1 5238	Fairchild, Palo Alto
Brachman R.	5	37	25	1 48	Fairchild, Palo Alto
Swartout W.	4	17	13	1.3077	University of Southern Calif.
McDermott J	5	39	30	1.3	?
Amarel S.	5	17	14	1.2143	?
McDermott D V.	5	24	23	1 0435	Yale University

Note: This table contains data that pertain to the period 1980–1984 Many of the authors' affiliations have changed since then The following list is an update as of Spring 1989: Aikins, *Aion Corp* ; Amarel, *Rutgers Univ* ; Brachman, *Bell Labs*; Cohen, *SRI*; Duda, *San Jose State Univ* ; Erman, *Cimflex Teknowledge*; Hart, *Syntelligence*; Hayes-Roth, B , *Stanford Univ* ; Lenat, *MCC*; McCarthy, *Stanford Univ* ; McDermott, *DEC*; Michie, *Turing Institute*; Minsky, *MIT*; Mitchell, *Carnegie Mellon Univ* ; Nilsson, *Stanford Univ* ; Riesbeck, *Yale Univ* ; Stefik, *Xerox PARC*

Table 3: Citation Analysis of the Outstanding Authors (ranked by average number of citations per paper)

literature (Greenlee and Small 1986, p. 95).

One complication with the mapping process described here is that a two-dimensional map of these clusters can be non-Euclidean. N number of papers spaced at arbitrary distances from one another can form an N-dimensional map. To visualize the map, the N-dimensional space is mapped repeatedly onto an (N-1)-dimensional space until N equals 2, and the map can be printed on paper. An analogous procedure would be to shine a light on a three-dimensional object made of wire and study its two-dimensional shadow on the wall. For more details on this process, see Beck, Pyle, and Lusted (1984); Greenlee and Small (1986); and Garfield (1979).

The mappings produced by citation analysis represent a snapshot of a body of literature for a certain period of time. By using reference lists from papers sampled at different periods of time, it is claimed that the evolution (or devolution) of a discipline can be studied. For example, Trofimenko (1987) uses citation analysis to draw graphs of a body of literature with a view to identifying the growth and decay of author groups (which are analogous to Greenlee and Small's clusters) His analysis identified a set of parameters derived from citation counts that he claims defines the growth-decay patterns of an author group. Some of these parameters appear to be phenomenally rigorous. For example, in table 2 (Trofimenko 1987, p. 239), Trofimenko quotes eight parameters with as much as five-figure accuracy (e.g., 0.0008)

Disadvantages of Using SCI

The validity of using SCI to review scientific work is questionable for the following reasons: (1) The basis of a citation analysis is a papers reference list, and these references might not accurately reflect the influences of a paper. (2) The use of citation analysis as an adequate method of assessing the worth of scientific research is problematic. (3) I have my own doubts about the meaning of the cocitation maps. (4) The SCI data can be significantly confusing, misleading, and incomplete.

Missing Influences. MacRoberts and MacRoberts (1986) challenge Garfield's assumption that references reflect the intellectual influences which acted on a paper. After a general discussion on this area (in which they identify 16 factors that can lead to an incomplete reference list), they describe their own study of the genetics literature After studying a random sample of 15 papers, they estimated that 719 references would be required to capture the influences of these papers. The papers contained only 216 references; that is, only 30 percent of the influences were cited As a postscript to their own paper, they mention that their own reference list only captures 20 percent of the influences on their paper.

If the MacRobertses are correct, and the reference list of a paper can contain as few as 20 to 30 percent of a paper's influences, then the Small and Greenlee-type maps can contain significant gaps. It is unclear what the MacRobertses study implies for Trofimenko's work. I can only speculate that Trofimenko's five-figure accuracy for his parameters might be spurious.

Inadequate as an Assessment Method. Small and Greenlee's maps are constructed solely from citation counts, and according to Healey, Rothman, and Hoch (1986), this basis is a poor one for assigning a value to scientific work. They caution against the use of citation analysis as the sole basis for assessing science because . . .

> when used in isolation, citation counts contain the seeds of a dangerous quantitative fetish-ism, especially in our current cultural climate which seems obsessed with "lists", "records" and "winners" (Healey, Rothman, and Hoch 1986, p. 235).

A United Kingdom Advisory Board research councilor surveyed in the Healey, Rothman, and Hoch (1986) study remarked that citation analysis was better than a mere anecdotal approach but should be used in association with other data such as peer-group judgments. Indeed, Irvine and Martin (1983) advocate using a whole host of "partially converging indicators" to assess scientific performance, such as the cost of the resources required to do the work reported in a paper, peer assessment, and others, then comparing these indicators between institutions to assess the significance of scientific establishments.

What Do the Cocitation Maps Mean? Not only can the Greenlee and Small cocitation maps contain significant gaps, but vital information can be lost when their N-dimensional cluster maps are mapped onto a two-dimensional object. Although I can not visualize the N-dimensional clusters, I have intuitive faith that the positions of these clusters reflect the cocitation counts. However, although I can visualize the two-dimensional clusters, I have far less faith in their geometry because I do not really understand what happens when an N-dimensional object is squashed into two dimensions.

Corrupt Data. When compiling the results of this study (see Namedropping Using SCI), I found that the raw SCI data for 17 of the 44 authors studied were confusing, misleading, and incomplete (17/44 = 38%). These data required significant amounts of post-processing to make them usable. Like the MacRobertses, I discovered that if the raw SCI data are read by a human rather than a computer, it becomes clear the SCI data are often unreliable.

As far as I have seen, this problem with corrupt data has not been discussed in the citation analysis literature, which suggests many of the SCI studies were based on the erroneous tacit assumption that the SCI data were reliable. Perhaps, in these days of computer-aided processing, it is considered pointless to do something so tedious as counting thousands of citation entries by hand. However, without such first-hand experience with the data, it would be all too easy to write programs that process large amounts of SCI data to produce impressive output which is grossly misleading.

Namedropping Using SCI

Bearing in mind my digression, we can now view the results in their appropriate perspective (that is, with a healthy degree of skepticism).

Citation counts, average citation counts per paper, and institutional affiliations of each author are shown in table 3. Institutional affiliations were found in the source index of the SCI 1980–1984 cumulative index. The number of outstanding authors from each institution is given in table 4.

N	*Institution with N Outstanding Authors*
10	?
9	Stanford University, Stanford
3	Fairchild, Palo Alto
2	Carnegie Mellon, Pittsburgh
2	Tufts University, Boston
2	University of Southern California, Marina Del Rey
1	MIT, Cambridge, Mass.
1	Rutgers State University, New Brunswick
1	Syracuse University, Syracuse
1	Teknowledge, Palo Alto
1	University of California, Berkeley
1	University of Edinburgh, Edinburgh
1	University of Pittsburgh, Pittsburgh
1	Xerox Corp., Palo Alto
1	Yale University, New Haven
1	multiple affiliations

Table 4: The Institutional Affiliations of the Outstanding Authors (as stated by the SCI)

Surname	N
Zadeh	13
Simon	6
Pauker	5
Winston	3
Shortliffe	3
Minsky	3
Duda	3
Newell	2
Bobrow	2
Buchanan	1
Winograd	1
Weiss	1
Waterman	1
Szolovits	1
Pople	1
Michie	1
Erman	1
Davis	1

The year of publication for each of the Top 49 papers is given in figure 1 References to the Top 49 papers are given in table 6

Table 5. Number of Times a Surname Appears in the Top-49 List

(An update to table 3 is given at the end of this article.) According to SCI, several authors with multiple papers seem to have multiple or missing affiliations. When in doubt and where possible, the most recent affiliation was used. Not all the names listed in table 2 are listed in table 3. Certain problems with the citation indexes made the data collection for some of the names difficult. In all, 38 percent of the entries (17/44) were problematic.

Citation index entries for three authors—King, Kunz, and Scott —could not be located. In the case of Scott, the reason seems to be that he was the second author of the Emycin manual (VanMelle was the first), and the citation indexes do not use the second, third, fourth, and so on, authors of a paper.

The index entries sometimes spread the same author over more than one entry; for example, Brachman R. and R. J.; Feigenbaum E. and E. A.; Gaschnig J. and J. G.; Lenat D., D. A., and D. R.; Riesbeck C. and C. K.; Stefik M. and M. J.; Waterman D. and D. A.; Winston P. and P. H.; and Zadeh L. and L. A. An inspection of the entries for Pople H. and Pople H. E. indicated that both these people had published articles with the same title in the same year and were being quoted by prominent AI and ES authors. In such cases, it seemed plausible to assume that these two entries referred to the same person, so their citation count could be combined.

The indexes also seem to confuse different authors with the same last name and initials. The entries for Davis R., Doyle J., Miller R. A., and Weiss S. M. all spanned too long a time period for them to have been produced by a single person. For example, the entry for Davis R. covered 1-1/2 pages and listed articles published from 1928 to 1984. I had the privilege of meeting Randall Davis (and his moustache) in May 1987. My subjective impression was that Davis was in his late thirties and, hence, would have been incapable of publishing anything in 1928.

The motivation for this study was to *identify papers worth reading*. A reasonable heuristic seemed to be "read the outstanding papers of the outstanding authors." By setting a sufficiently large number of citations as a threshold, the number of papers that require reading can be reduced to a manageable size. In the case of this study, a citation threshold of 15 citations per article excludes all but 49 of them. Table 5 shows the number of times various surnames appeared in the Top 49 group.

Note that half the papers in the Top 49 list were published from 1975 to 1979. Because this study used citation counts collected from 1980 to 1984, it appears that it takes five to seven years for a paper in the AI and ES field to have its maximum effect on the citation indexes. Apparently, the use of review articles to preselect surnames for further study resulted in a retrospective bias in the results. This result is not surprising. According to Bindon (1987), to appear in a review article, an author either needs to have been widely recognized as being of major significance; be the inventor of some technique that has proven to be of some operational significance; or be doing work which, in the opinion of the review article's author, will at some point be of major significance.

In the first instance, some time would have to elapse before the author's innovations spread into the mainstream and are recognized as important. In all three instances, a further delay would result from the lapsed time between when a review article is written and when it is published.

Validating Namedropping

Two possible problems with the namedropping technique were identified and analyzed.

Effect of Initial Selection of Reference Articles

To test the sensitivity of the method to the choice of initial review articles, namedropping was repeated using the seven articles plus another well-known introductory AI text (Winston 1984). It was found this additional information shrank the outstanding group from 44 to 42. Hence, it appears that no utility existed in collecting more than seven anecdotally selected review articles for use in namedropping. The initial seven articles seem to have collected all the names (and more) that a larger sample size would have collected.

Are the Outstanding Authors Outstanding?

Note that no connection seems to exist between the ranking assigned to authors using namedropping and that assigned using the average number of citations per paper. This result could suggest that namedropping can not be used to select outstanding authors. This notion was tested as follows:

De Solla Price (1965) conjectures that one-eighth of all papers are cited 20 or more times. The 44 outstanding authors have written 268 papers. If they were randomly selected groups of papers, then we would expect that 21 of these (268/12.5) would be quoted 20 times or more. However, 31 of the selected papers received 20 or more citations, suggesting that namedropping selects better-than-average authors.

A sample of the N = 3 namedrops group was examined to see if outstanding authors had been missed by namedropping. Ten names from this group were studied to see if they were outstanding authors For the purposes of this test, an author was considered outstanding if the author had received a total of more than 30 citations or had published one paper with more than 12 citations. Of the 10 names, 6 were outstanding.

Combining these results, I can say that namedropping can select a sample of truly outstanding authors. However, the sample is not exhaustive. This result is hardly surprising given the method of data collection. Because namedropping was performed by hand, and the amount of data processed was rather large, then it is possible that minor clerical errors might have occurred in compiling these results. Typically, these clerical errors would have been errors of omission; that is, a citation would not be credited to a particular surname. Consequently, certain outstanding authors might be missed However, as a heuristic technique to focus on a group of significant authors, these errors of omission do not threaten the validity of the technique. If an author survives through the possible omission with namedropping, then the author must have been mentioned frequently.

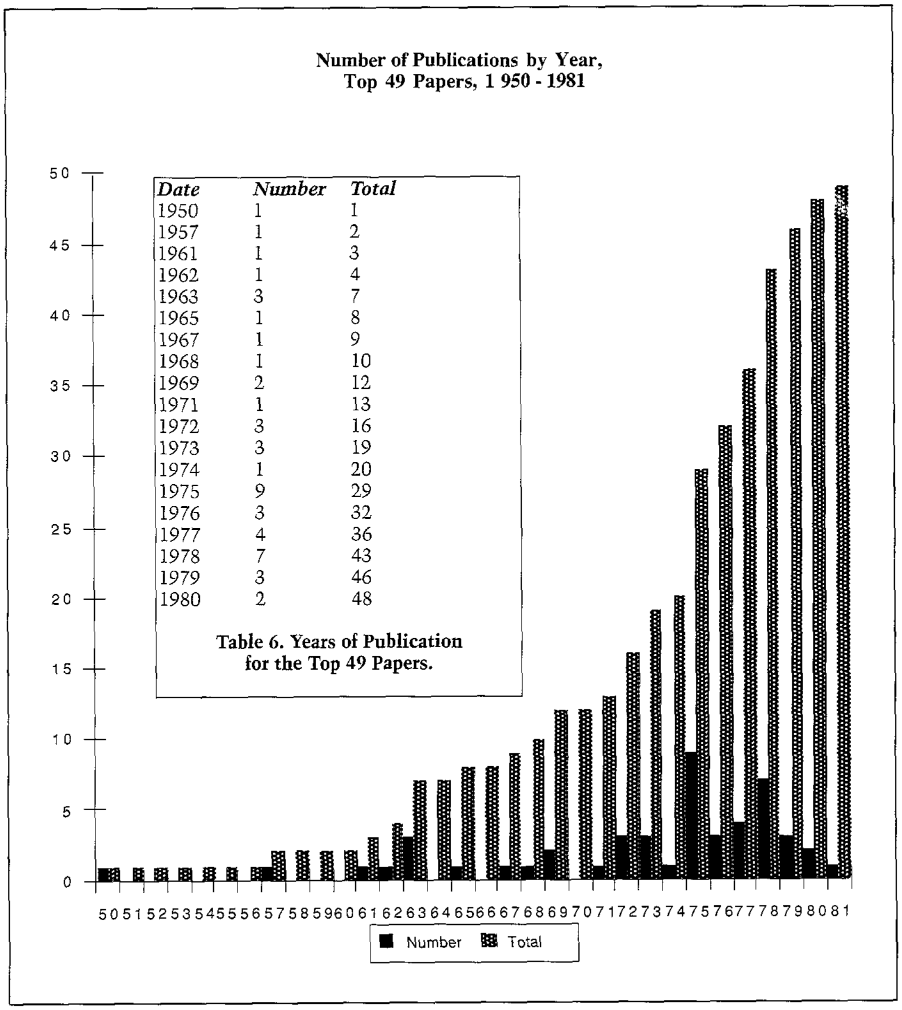

Date	Number	Total
1950	1	1
1957	1	2
1961	1	3
1962	1	4
1963	3	7
1965	1	8
1967	1	9
1968	1	10
1969	2	12
1971	1	13
1972	3	16
1973	3	19
1974	1	20
1975	9	29
1976	3	32
1977	4	36
1978	7	43
1979	3	46
1980	2	48

Table 6. Years of Publication for the Top 49 Papers.

Figure 1. Number of Publications by Year Top 49 Papers, 1950-1981
Note the historical bias in the sample; half the papers were published in the period 1975-1979

Searching for Biographical Data

The search for biographical data was not successful Two sources were used to collect these data. American Men and Women of Science and Who's Who in Science. Of a sample of 35 of the outstanding authors, 11 could not be found in the 1979, 1982, or 1986 editions of American Men and Women. The entries for the rest were brief. Sometimes, only their current institutional affiliations were stated. In other cases, uninformative initials were used to record a person's degree(s). For example, many of the entries for the authors recorded their degree(s) as BA or BSc. These data are uninformative because depending on the university, a BA or a BSc can be in any number of areas. Who's Who is a much smaller volume than American Men and Women and, consequently, even less informative.

A biographical search was considered using the author's publications. Sometimes, articles are published with an institutional afliliation or a brief biography. However, such a search would have been time consuming and, hence, was abandoned.

Conclusions

The goal of this work was to develop a list of significant authors, institutions, and dates within the AI and ES field. The resources available for this project were those available to the majority of researchers, that is, about five days of full-time work and limited funding. Within these restrictions, automated SCI-based analysis was not feasible This result proved to be a handicap but not an impossible one (indeed, given all the problems with SCI, this "handicap" might even have been an advantage).

A heuristic analysis method called namedropping was developed. The technique combined SCI with a peer review. Review articles were used to rapidly access a sample of the authors recognized within a field as being part of the field's elite. The analysis method could be used by the science historian as a basis for restricting the study to a limited number of institutions or authors. The historical overview provided by namedropping could be useful for those wanting to explore the foundations and moderately recent trends of their discipline. At the least, namedropping is useful for those who wish to reduce the amount of reading required when studying a new area.

A definite historical bias was seen

Author of Paper	*Number of Citations*	*Year of Publication*	*Name of Journal*
Duda	232	73	Pattern Classification
Zadeh	212	65	Information Control 8,338
Newell	145	72	Human Problem Solving
Zadeh	85	63	Linear Systems Theory
Shortliffe	82	76	Computer Based Med. C.
Pauker	58	80	New Eng J Med 302,1109
Zadeh	53	73	IEEE T Syst Man Cybe 328
Winston	53	77	Artificial Intelligence
Simon	51	69	Sci Artificial
Zadeh	49	78	Fuzzy Set Systems 1,3
Pauker	46	75	New Eng J Med 293,229
Winograd	44	72	Understanding Natura
Zadeh	42	75	Information Sci 1 8 199
Shortliffe	39	79	P.IEEE 67, 1027
Pauker	39	76	Am J Med 60,981
Zadeh	34	75	Fuzzy Sets Their App
Zadeh	33	75	Information Sci 8,301
Duda	32	72	Communications ACM 15,11
Zadeh	31	75	Inf Sci 9,43
Pauker	29	76	Ann Inter Med 85,8
Bobrow	28	77	Cognitive Sci 1,3
Zadeh	27	68	J Math Anal Appl 23,421
Waterman	27	78	Pattern Directed Inf
Szolovits	27	78	Artificial Intellig 11, 115
Winston	26	75	Psychol Computer Vis
Minsky	24	69	Perceptrons
Simon	23	62	P Am Phil Soc. 106, 467
Minsky	23	75	Psychol Computer Vis
Simon	22	74	Science 183,482
Simon	20	61	Econometrica 29,111 20
Davis	20	77	Artificial Intellig 8,15 20
Erman	19	80	Computing Surveys 12,213
Zadeh	18	71	Information Sciences 3,117
Simon	18	63	Psychol Rev 70, 534
Weiss	17	78	Comput Biol Med 8,25
Simon	17	57	Models Man
Minsky	17	67	Computation Finite I
Duda	17	79	IEEE T Pattern Anal 1,259
Buchanan	17	78	Artificial Intellige 11,5
Zadeh	16	50	P I Radio Engrs 38,391
Shortliffe	16	75	Mathemat. Bioscie 23, 351
Newell	16	73	Visual Information P
Michie	16	79	Expert Systems Micro
Bobrow	16	75	Representation Under
Zadeh	15	78	Int J Man Machine St 10,395
Zadeh	15	63	IEEE T Autom Control 8,59
Winston	15	81	Lisp
Pople	15	77	5th P Int Joint CAR
Pauker	15	78	Semin Nucl Med 8,324

Table 9: The Top 49 Papers AI & ES Papers (as found through namedropping)

in the results: The highly cited papers found using namedropping were five to seven years old. Another possible bias was a predominance of medical-based systems discussed in the review articles. It is possible that this situation is an artifact of the initial selection of reviewed articles. However, it is also possible it is a true reflection of the emphasis of the pioneering AI work. Further analysis is required to resolve this point.

Attempts to link namedropping with biographical data were less than successful. The readily accessible biographical data were all too brief. This study was therefore unable to find evidence for a coherent body of study called cognitive science based on the 19801984 SCI.

In my introduction, I commented that if something like cognitive science existed, then we should be able to find traces of it in the pattern of scientific publications and biographies of research. This study was unable to find any such traces. What can we infer from this? I can see four possibilities for why the evidence was not found, and further work would be required to endorse any of these: (1) I had insufficient resources, (2) The SCI is a poor resource for studying the structure of a new science like AI and ES, (3) The evidence doesn't exist. Perhaps the real achievement of the last three decades of work on synthesizing human cognition has been the development of a grab-bag of useful tricks for developing complicated programs rather than a comprehensive theory of human cognition. Recall the difficulties this study had with finding review articles on human cognition. As I pored through AI journals looking for articles that summarized the field, I became nervous about the lack of such articles. I began to wonder about the discipline (or lack of it) in AI. While in such a dubious mood, I found John McCarthy's 1987 Turing Award lecture, Generality in Artificial Intelligence (McCarthy 1987). It contains phrases such as the problem of generality in artificial intelligence is almost as unsolved as ever and concludes with all this [the ideas discussed in the article] is unpleasantly vague. Such articles do not serve to promote the notion of a grand theory of AI, and (4) The evidence isn't there yet. Recall that the historical bias in the sample half of the cited articles were published from 1975 to 1979. Perhaps cognitive science has formed since then. To check this possibility, this analysis should be repeated for the period 1984–1988.

As a postscript, I would like to repeat this study using the 1984–1988 SCI cumulative index. The starting point for this analysis would be review articles and texts published from 1983 to 1988 (which were not used in this analysis). I'm currently collecting a list of such articles and would appreciate any contributions from the AI Magazine readership. Please send any such references to me at Sixth Floor, Mayfair House, 73 Castlereagh Street, Sydney, New South Wales, Australia 2001.

Acknowledgments

Damien Laker (Honors, Science and Technology Studies, University of New South Wales) helped with the boring leg work associated with this project as well as acted as a testing ground for my ideas. Rosemary Robbins (Honors, Science and Technology Studies, University of New South Wales) led me to the work of Healey, Irvine, and the MacRobertses The enthusiasm of George Bindon (Lecturer, Science and Technology Studies, University of New South Wales) for the project was always encouraging

References

Beck J R , Pyle, K I , and Lusted, L B 1984. A Citation Analysis of the Field of Medical Decision Making. *Medical Decision Making* 4(4): 449–468

Bindon, G. 1987 Personal communication Bindon is a lecturer in the School of Science and Technology Studies, University of New South Wales, P.O. Box 1, Kensington, New South Wales, Australia 2033

Bobrow, D G , and Hayes, P. J. 1985. Artificial Intelligence: Where Are We? *Artificial Intelligence* 25: 375–415.

de Solla Price, D 1965 Networks of Scientific Papers *Science* 149: 510–515

Duda, R O , and Buchanan, B. G 1983. Principles of Rule-Based Expert Systems *Advances in Computing* 22: 163–216.

Gardner, H 1985 *The Mind's New Science* New York: Basic Books.

Garfield, E. 1979 *Citation Indexing—Its Theory and Application in Science, Technology, and Humanities* New York: Wiley.

Garfield, E 1970. Citation Indexing for Studying Science. *Nature* 227: 669–671.

Greenlee, E., and Small, E. 1986 Collagen Research in the 1970s *Scientometrics* 10(1–2): 95–117.

Hayes-Roth, F., Waterman, D. A., and Lenat, D B., eds. 1983 *Building Expert Systems* Reading, Mass.: Addison-Wesley.

Healey, P., Rothman, H., and Hoch, P. K 1986 An Experiment in Science Mapping for Research Planning *Research Policy* 15(5): 233–252

Irvine, J., and Martin, B. R 1983. Assessing Basic Research: The Case of the Issac Newton Telescope. *Social Studies of Science* 13: 49–86

McCarthy, J. 1987 Generality in Artificial Intelligence *Communications of the ACM* 30(12): 1030–1035

MacRoberts, M. H , and MacRoberts, B R 1986 Quantitative Measures of Communication in Science: A Study of the Formal Level. *Social Studies of Science* 16: 151–172

Newell A 1982. The Knowledge Level *Artificial Intelligence* 18: 98–127.

Ramsey, C L , Reggia, J A , and Nau, D S 1986 A Comparative Analysis of Methods for Expert Systems *International Journal of Man-Machine Studies* 24: 475–499.

Rich, E 1983 *Artificial Intelligence* Singapore: McGraw-Hill

Sandell, H S H., and Bourne, J R 1985. Expert Systems in Medicine: A Biomedical Engineering Perspective *Critical Review of Biomedical Engineering* 12(2): 95–129

Small, H , and Sweeney, E 1985 Clustering the Science Citation Index Using Co-Citations, 1) A Comparison of Methods *Scientometrics* 7(3–6): 391–409.

Trofimenko, A 1987 Scientometric Analysis of the Development of Nuclear Physics during the Last 50 Years *Scientometrics* 11(3–4): 231–250.

Winston, P H. 1984 *Artificial Intelligence.* Reading, Mass.: Addison-Wesley.

Tim Menzies is a Ph D student at the University of Technology, Sydney, Australia, and holds a Master's of Cognitive Science from the University of New South Wales, Sydney, Australia For the last three years, he has been active in expert systems implementation and has built Prolog-based systems for process control and farm management Menzies is currently working as a senior knowledge engineer for the Australian Mutual Provident Society. His research interests include combining programming paradigms, logic programming, object-oriented programming, knowledge-acquisition environments, and the maintenance of procedural knowledge