The definition of cancer research: journals, titles, abstracts or keywords?

Grant Lewison

g.lewison@ucl.ac.uk and *grantlewison@aol.co.uk*
Department of Information Studies, University College London, Gower Street, London WC1E 6BT, UK

Abstract

Three versions of a "filter" used to identify papers on cancer research, as defined by Cancer Research UK and interpreted by four experts, were compared. The first was based only on specialist journals and had unacceptably low recall. The second was based also on title words, and had both precision and recall above 0.9. The third was based additionally on words in the abstract and/or keywords provided with the paper: it improved the recall to almost unity but the precision was severely degraded, with many false positives. The three filter versions were compared in terms of the outputs of 15 countries in the *Web of Science* in recent years, and in some instances gave differing indicators of their performance (numbers of papers and citations) which could give conflicting messages for science policy.

Introduction

The evaluation of research often involves making comparisons between the outputs from different actors (countries (May, 1997a; King, 2004; Prathap, 2010), regions (Lewison, 1991; Shapira *et al*., 2003; Holbrook & Clayman, 2006; Zhou *et al*., 2009; Levitt & Thelwall, 2010), universities (Adams, 1998; Prathap & Gupta, 2009; Aguillo *et al*., 2010; Docampo, 2011), institutes and departments (Yi & Kang, 2000; Hyvarinen, 2009; Torres-Salinas *et al*., 2009), even persons (Qiu *et al*., 2008; Alonso *et al*, 2010; Claro & Costa, 2011)) on a number of criteria. These typically include the numbers of papers and some of their parameters, usually including numbers of citations. As an alternative, the output of one or more actors may be compared with the world average, particularly when countries are being compared in terms of their citation performance. It is well attested that the norms of production and of citation vary greatly between fields (Schubert & Braun, 1996; Kostoff, 1997; Nederhof & Visser, 2004; Podlubny, 2005), and between subject areas within them (Zitt *et al*., 2005), so that any comparisons must respect these differences if they are to be valid.

Another common activity for bibliometricians is to examine a particular scientific field in order to determine its dynamics (how fast it is growing relative to all science, for example; *e.g*., Gupta & Dhawan, 2008), its structure (the relationships between sub-areas and how they are changing, often shown as maps: Cambrosio *et al*., 2006; Calero-Medina & Noyons, 2008; Jeong & Kim, 2010) and the principal actors (Karisiddappa *et al*., 2002; Gupta & Dhawan, 2008; Kostoff & Morse, 2011).

Both of these tasks require the field or subject area to be defined, for details of the relevant papers to be extracted from a database by means of a "filter", and for the filter to be calibrated in terms of its precision (or specificity) and recall (or sensitivity). Somewhat surprisingly, the first and last of these three jobs are often omitted. But they are fundamental to a rigorous analysis of a subject area that will command confidence among the study's readership. Very often, the "filter" simply consists of a set of journals allocated to pre-set subject areas by the database publisher (*e.g*., the *Web of Science* – WoS, Ugolini & Mela, 2003; Adams, 1998; Sooryamoorthy, 2011; *Scopus*, Gupta and Dhawan, 2009) or determined from cognitive relationships (Leydesdorff, 2008). However, now that several databases also contain searchable abstracts of many of the papers that they process, these have sometimes been used to generate additional papers for the analysis (Meyer *et al.,* 2010; Kostoff & Morse, 2011). Sets of keywords are also increasingly being added to the paper record – some given by the paper authors, some by the journal, or by the database provider (*e.g.,* MedLine). More complicated filters have also been devised, based on citations either from or to papers to or from a "core set" (Schwechheimer & Winterhager, 2001; Glänzel *et al*., 2009; Bolaños-Pizarro *et al*., 2010). If the subject area of interest is not too large, then it may be possible to improve the precision of the filter by inspection of the individual papers with a view to the rejection of ones deemed irrelevant.

The lack of attention to how well the filter performs is surprising, as a poorly-designed filter can give spurious and misleading information about a subject area – how big it is, how well cited it is, and its structure and the principal actors within it. Moreover, it is often difficult for others to check the stated results and see how sensitive they might be to small changes in the filter used to generate them.

This paper examines three filters that could be used to define the subject of "cancer research", based first on oncology journals, second on journals and title words, and third on these plus terms in the abstract and keywords. The three successive filters will yield increasing numbers of papers. Which is best in terms of precision and recall? And how much difference does the choice of filter make to the dynamics of the subject area, its citation norms and the relative ranking of some individual countries?

The first task in any work of this type is to provide a simple and clear definition of the subject, see Webster (2005 – Appendix) and as was done recently for nanotechnology by Maghrebi *et al*., (2011). Usually 50-100 words are enough, and they tell readers what is included and what is excluded, so that they know the definition used, even if they might have defined it differently. For cancer research, we use the definition provided by Cancer Research UK, a leading charity, which reads as follows and has just 53 words:

The study and treatment of cancer or tumours. This incorporates academic oncology and clinical oncology. Academic oncology is aimed at identifying the causative agents or underlying genetic defects producing cancer and at developing these discoveries into effective drugs and other therapies. Clinical oncology is oriented towards the treatment, management and cure of cancer.

Methodology

The process of filter development is, or should be, a progressive process and it needs to be tested at each stage to check that precision and recall are improving and approaching unity. The simplest way to start is to select some very obvious title words, or address words, that indicate the subject. In the present study, title words could be *cancer*, carcinoma*, leukemi*, oncol*, tumor** (where * denotes any character(s) or none) and address words or contractions could be *CANC, ONCOL, TUMOR*. These should be used to search the database (the Web of Science, Science Citation Index Expanded was used in this study, which was limited to articles, proceedings papers and reviews) for 2005 and 2009 publication years. The sources (*i.e*., journal, year, volume, issue, pages) should then be downloaded to file and the names of all the journals listed that have one or more papers. From this list, all those journals with appropriate strings in their titles, such as CANCER, ONCOL, ONKOL, LEUKAEM, LEUKEM, TUMOR would be marked, plus a few others clearly relevant such as CHEMOTHERAPY.

The first filter is then the list of all these journals, but for practical purposes it can be collapsed into a much shorter set of search strings by the use of asterisks, thus the six journals: *ADVANCES IN CANCER RESEARCH* or *AMERICAN JOURNAL OF CLINICAL ONCOLOGY-CANCER CLINICAL TRIALS* or *ANTI-CANCER AGENTS IN MEDICINAL CHEMISTRY* or *ANTI-CANCER DRUGS* or *ANTICANCER RESEARCH* or *ASIAN PACIFIC JOURNAL OF CANCER PREVENTION* can all be represented by the contracted statement: A*CANCER*. This procedure ensures that the journal list is up-to-date, but it can be repeated for an earlier year in order that the filter should capture papers in specialist journals that are no longer in existence, or no longer processed for the database, although the contracted statements will mostly do this automatically.

The second filter uses both specialist journals and title words. At this stage, it is necessary to engage the services of an expert in the subject area. The titles of all the papers in the specialist journals in the most recent year available are downloaded from the database to a file, and after some cleaning to remove punctuation marks, all the title words are listed in descending order of frequency of occurrence. Many of them will be common words not relevant to the subject, but the expert will be able to identify relevant ones and mark them. Some may need to be qualified by the presence (or absence) of another word in order to ensure that they are used in the correct sense or context. Thus "tumor" needs *not* to be accompanied by "necrosis factor" in order to be relevant to cancer, and "irradiation" *must* be accompanied by "fractionated". The title words are conveniently sorted alphabetically and formed into a set of search statements, which can be combined with the search statements based on specialist journal names.

The third filter is similar to the second, but the list of words is applied not only to the titles of the papers but also to the abstracts and keywords.

The list of title words, and possibly also the list of specialist journals, needs to be tested to check that it does not generate too many false positives or false negatives. There are several ways to perform this calibration (Lewison, 1996; Lewison, 1999) but the simplest is based on the assumption that research teams whose addresses contain one or more of the selected contractions (here, CANC, ONCOL, TUMOR) will publish similar papers to those without such address strings. Three sets of papers are then identified and the bibliographic details (title, source) of samples of them (perhaps a few hundred) are downloaded to file:

A papers captured by the filter AND with the contractions in their address(es)
B papers with the contractions in their address(es) but NOT captured by the filter
C papers captured by the filter but WITHOUT the contractions in their address(es)

The expert is then invited to mark these papers as relevant (1) or not relevant (0); she/he may shade the mark with a decimal fraction for papers where the title does not give enough information for a firm decision to be made. In the interests of fairness, it is advisable to mix up papers from the three sets so that the expert marks them without knowing to which set they belong – they should, of course, have hidden codes or other markings so that they can subsequently be identified for analysis purposes.

In order to calibrate the filter, we need to determine the number of missing papers, *i.e.*, ones not captured by the filter and not having the contractions in their addresses. If the number of papers in the database in a selected year or years in set A is a, and the precision of this set, based on the sample, is p(a), then the true number of papers = $a^* = a$ x p(a), and similarly for b^* and c^*. The assumption in the previous paragraph on the similarity of papers from groups with eponymous (cancer) address strings to those without them yields $d^* = c^*$ x b^* / a^*, and the true total of papers is $a^* + b^* + c^* + d^*$. The true number retrieved is $a^* + c^*$ and the actual number is $a + c$, so precision $p = (a^* + c^*) / (a + c)$ and recall $r = (a^* + c^*) / (a^* + b^* + c^* + d^*)$.

Filter development proceeds in steps, and at each step it is necessary to check that p and r are increasing until a point is reached when gains in one are offset by losses in the other. The calibration factor, $CF = p / r$, and may be either greater than or less than unity.

Results

Precision and recall of the filters

The latest version of the cancer research filter, labelled ONCOL, is actually the sixth, earlier versions having been supplemented with the names of new drugs and newly-discovered genes that code for an increased cancer risk. It is now quite complex, with 55 journal search strings, 8 title/abstract words with Boolean conditions, and 293 single title/abstract words or pairs (*e.g.*, *xeroderma pigmentosum*). When applied to the WoS for 2009, the numbers of papers identified were as in Table 1.

Table 1. Numbers of papers retrieved from the WoS by the three versions of the ONCOL filter, publication year = 2009.

Filter	*based on:*	*Set A*	*Set B*	*Set C*	*Retrieved*
1	journals only	15147	28629	10365	25512
2	journals & titles	28374	15402	40081	68455
3	journals, titles, abstracts, keywords	35143	8633	89283	124426

For the first version of the filter, based only on specialist journals, the calculation of the numbers of cancer research papers and *p* and *r* went as follows:

Table 2. Calculation of the precision, *p*, and recall, *r*, of the first version of the filter.

Set	*n*	*Sample*	*OK*	*p*	*n**
A	15147	213	210.8	0.990	14991
B	28629	524	69.7	0.133	3808
C	10365	80	73.3	0.916	9494
D					2412
Total	25512				30705

For this filter, p = (14991 + 9494) / 25512 = 0.96, and r = 25512 / 30705 = 0.83. This might be deemed fairly satisfactory, but we need to investigate the other two versions before we can rely on the assumption that researchers in eponymous (cancer) departments publish in a similar range of journals to those in non-eponymous departments.
The second version of the filter, which uses title words as well as specialist journals, gives a much larger estimate of the size of the cancer research output:

Table 3. Calculation of the precision, *p*, and recall, *r*, of the second version of the filter.

Set	*n*	*Sample*	*OK*	*p*	*n**
A	28374	509	492.4	0.967	27438
B	15402	524	69.7	0.133	2048
C	40081	470	422.4	0.899	36033
D					2690
Total	68455				68209

For this version, *p* = 0.93 and *r* = 0.93, but the estimated true total is more than twice as large as with the first version. Evidently, there are many cancer research papers not in specialist journals – in fact, the majority.
Finally, we apply the terms in the filter also to abstracts and keywords with the following result:

Table 4. Calculation of the precision, *p*, and recall, *r*, of the third version of the filter.

Set	*n*	*Sample*	*OK*	*p*	*n**
A	35143	1008	830.5	0.824	28947
B	8633	1000	6.5	0.0065	56

C	89283	967	456.4	0.472	42134
D					81
Total	124426				71218

The estimated true total is now somewhat larger, but only by 4.4%. The addition of papers retrieved because of words in the abstract or keywords has apparently given almost complete retrieval (r = 0.998) but the precision is now severely degraded to p = 0.571. Many of the papers marked as false positives were on topics unrelated directly to cancer but described clinical manifestations or treatments that could also have occurred when cancer was present. As a result, a "cancerous" word appeared in the abstract or keywords plus, but seldom in the author-selected keywords. It is reasonable to conclude that the second version of the filter is the most nearly correct one, as version 1 has a low recall (based on the results of versions 2 and 3) and version 3 a low precision.

Country outputs compared

The three versions of the filter were applied to the Web of Science for publication years 2005 and 2009, and the numbers of papers world-wide and from 15 leading countries were determined. The results are shown in Table 5, and are given as percentages of the world totals using integer (whole) counting.

Between 2005 and 2009, world cancer research output increased by 20% according to version 1 of the filter, but by 29% according to both version 2 and version 3. Since these two versions gave very similar totals (about 70,000 papers per year), the latter growth rate can be accepted rather than the former. This means, incidentally, that an increasing percentage of cancer papers are *not* published in specialist cancer journals but in general journals.

The results for the individual countries are somewhat varied, as would be expected. Some countries have similar percentage presences in the world on all three versions of the filter, such as France (in 2005), Germany and Spain. A few have a higher presence according to the fuller versions of the filter (F3 > F2 > F1), such as China and South Korea; but most show the reverse, with a higher presence in the specialist journals and a lower one in the titles, abstracts and keywords, notably Italy, the Netherlands, Sweden and Belgium.

Between 2005 and 2009, despite the steady increase in international co-authorship, the four leading countries in Table 5 (the USA, Japan, Germany and the UK), and also the Netherlands and Sweden, all decreased their percentage presence in cancer research according to all three versions of the filter. Six countries (Canada, China, Spain, South Korea, Australia and Switzerland) all increased their presence, again based on all three filter versions. But for the other three (France, Italy and Belgium) the message was mixed, and the change could have been reported as either a gain or a loss of presence.

We turn now to the rating of countries based on the mean citation scores of their papers. In Figure 1, these have all been compared with the world mean values in a five-year window, *i.e.,* the numbers of citations in 2005 thru 2009 for the 2005 publications. These were respectively 18.35, 16.1 and 16.04 cites for filter versions 1, 2 and 3. It appears that papers in the specialist cancer journals received more citations than ones in the general journals that were retrieved because of their titles or abstracts/keywords.

Table 5. Numbers of cancer research papers (articles, proceedings papers and reviews) for 15 leading countries in the WoS, publication years 2005 and 2009, according to the three versions of the filter: percentages of world total.

ISO	*Country*	*2005F1*	*2005F2*	*2005F3*		*2009F1*	*2009F2*	*2009F3*
	World (papers)	*21236*	*53072*	*96106*		*25512*	*68455*	*124426*

US	United States	41.95	38.21	38.51		38.77	34.59	34.70
JP	Japan	10.50	11.12	10.64		9.67	9.19	8.70
DE	Germany	9.10	9.20	8.99		8.32	8.12	8.20
UK	United Kingdom	8.11	7.73	7.78		7.47	6.97	7.21
IT	Italy	7.60	6.55	5.85		7.22	6.63	6.07
FR	France	5.53	5.58	5.53		6.12	5.54	5.42
CA	Canada	4.82	4.24	4.34		4.93	4.28	4.44
CN	China (P. R.)	3.18	3.71	3.82		7.43	8.11	8.18
NL	Netherlands	4.14	3.37	3.02		4.09	3.20	2.89
ES	Spain	2.68	2.71	2.64		3.06	3.09	3.10
KR	South Korea	1.99	2.56	2.59		3.51	4.29	4.06
SE	Sweden	2.71	2.29	2.16		2.49	2.00	1.86
AU	Australia	2.35	2.24	2.31		2.91	2.57	2.58
CH	Switzerland	1.76	1.71	1.82		2.19	1.85	1.89
BE	Belgium	1.78	1.55	1.48		1.88	1.45	1.41

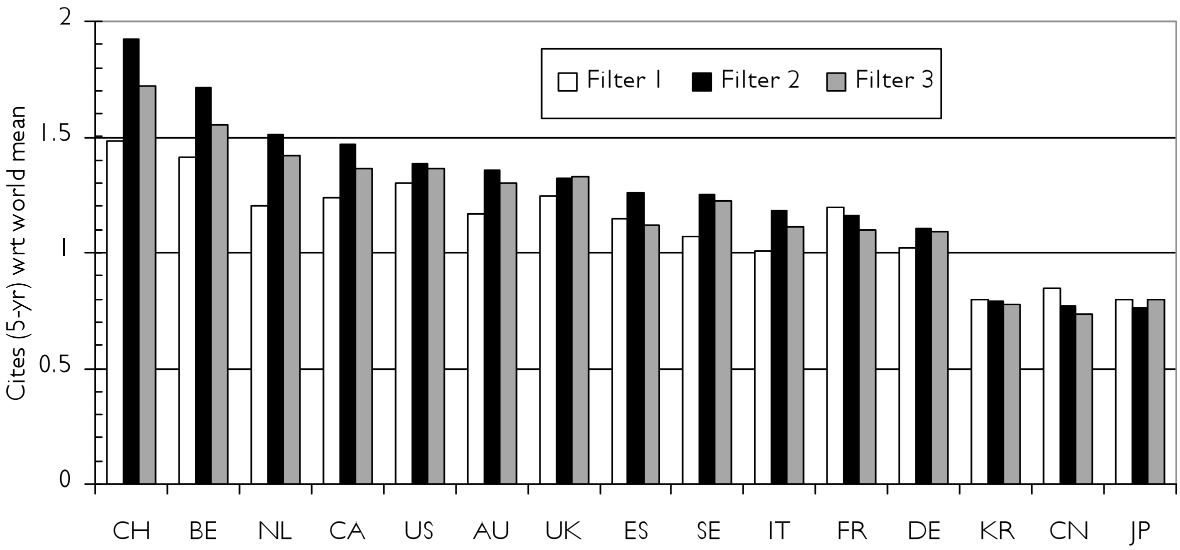

Figure 1. Five-year citation scores relative to the world mean values for cancer research papers from 15 leading countries (for codes, see Table 5) published in 2005 and cited 2005 thru 2009, based on three versions of the cancer research filter.

The countries have been ordered in Figure 1 on the basis of their citation performance on the second version of the filter, and this ranking puts three small European countries ahead of Canada and the USA. Their performance, and that of Canada and several other European countries, is much better than that shown by the specialist journals, where the USA shows to advantage, but is still behind Switzerland and Belgium. The three East Asian nations all score relatively low on all three filter versions, as has been found elsewhere (López-Illescas *et al.*, 2008).

Conclusions

This paper has examined one particular field, namely cancer research, in some detail and has shown that the world output in the Web of Science was of the order of 70,000 research papers per year in 2009. The best filter in terms of both precision and recall was one based both on specialist journals and title words. The omission of title words meant that fewer than half the relevant papers were identified, and the addition of words in abstracts and/or keyword lists was not helpful as nearly all the additional papers identified were false positives.

The effects of using version 1 or version 3 of the filter instead of version 2 were rather variable, and some countries benefited in terms of their percentage presence or relative citation score, and some were disadvantaged. Few of the differences were large, but countries are often looking (Grant & Lewison, 1997; May, 1997b) for evidence of small improvements to their relative position in order to claim that their science is being well managed and providing good value for money, as with the European agri-food research programmes (Borsi & Schubert, 2011), where different search strategies sometimes produced very different outcomes. It seems important, therefore, that any such claims should be based on the best approximation to the true set of papers in the selected field or subject area, even though it is really a "fuzzy set" rather than one that can be precisely defined without argument.

Acknowledgment

I am grateful to the four cancer specialists who kindly marked the sets of papers from the three filters: Dr Lynne Davies of Cancer Research UK, Dr Katie Hyde of the National Cancer Research Institute and to Professor Arnie Purushotham and Professor Richard Sullivan of King's College London. I should also like to thank Dr Philip Roe who wrote the macros that enabled me to download and analyse the papers on the Web of Science and their citations.

References

Adams, J. (1998) Benchmarking international research. *Nature* 396, 615-618.

Aguillo, I.F., Bar-Ilan, J., Levene, M. & Ortega, J.L. (2010) Comparing university rankings. *Scientometrics,* 85, 243-256.

Alonso, S., Cabrerizo, F.J., Herrera-Viedma, E. & Herrera, F. (2010) hg-index: a new index to characterize the scientific output of researchers based on the h- and g-indices. *Scientometrics,* 82, 391-400.

Bolaños-Pizarro, M., Thijs, B. & Glänzel, W. (2010) Cardiovascular research in Spain. A comparative scientometric study. *Scientometrics*, 85, 509-526.

Borsi, B. & Schubert, A. (2011) Agrifood research in Europe: a global perspective. *Scientometrics,* 86, 133-154.

Calero-Medina, C. & Noyons, E.C.M. (2008) Combining mapping and citation network analysis for a better understanding of the scientific development: the case of the absorptive capacity field. *Journal of Informetrics*, 2, 272-279.

Cambrosio, A., Keating, P., Mercier, S., Lewison, G. & Mogoutov, A. (2006) Mapping the emergence and development of translational cancer research. *European Journal of Cancer*, 42, 3140-3148.

Claro, J. & Costa, C.A.V. (2011) A made-to measure indicator for cross-disciplinary bibliometric ranking of researchers' performance. *Scientometrics,* 86, 113-123

Docampo, D. (2011) On using the Shanghai ranking to assess the research performance of university systems. *Scientometrics,* 86, 77-92.

Glänzel, W., Janssens, F. & Thijs, B. (2009) A comparative analysis of publication activity and citation impact based on the core literature in bioinformatics. *Scientometrics*, 79, 109-129.

Grant, J. & Lewison, G. (1997) Government funding of research and development. *Science*, 278, 878-879.

Gupta, B.M. & Dhawan, S.M. (2008) Condensed matter physics: an analysis of India's research output, 1993-2001. *Scientometrics*, 75, 123-144.

Gupta, B.M. & Dhawan, S.M. (2009) Status of India in science and technology as reflected in its publication output in the Scopus international database, 1996-2006. *Scientometrics,* 80, 473-490.
Holbrook, J.A. & Clayman, B.P. (2006) Research funding by city: an indicator of regional technological competitiveness? *Research Evaluation*, 15, 221-231.
Hyvarinen, J. (2009) Evaluation of TEKES funding for research institutes and universities – the role of talent. *Research Evaluation*, 18, 365-373.
Jeong, S. & Kim, H-G. (2010) Intellectual structure of biomedical informatics reflected in scholarly events. *Scientometrics*, 85, 541-551.
Karisiddappa, C.R., Gupta, B.M. & Kumar, S. (2002) Scientific productivity of authors in theoretical population genetics. *Scientometrics,* 53, 73-93.
King, D.A. (2004) The scientific impact of nations. *Nature,* 430, 311-316.
Kostoff, R.N. (1997) Citation analysis cross-field normalization: a new paradigm. *Scientometrics,* 39, 225-230.
Kostoff, R.N. & Morse, S.A. (2011) Structure and infrastructure of infectious agent research literature: SARS. *Scientometrics,* 86, 195-209.
Levitt, J.M. & Thelwall, M. (2010) Does the higher cittion of collaborative research differ from region to region? A case study of Economics. *Scientometrics,* 85, 171-183.
Lewison, G. (1991) The scientific output of the EC's less favoured regions. *Scientometrics,* 21, 383-402.
Lewison, G. (1996) The definition of biomedical research subfields with title keywords and application to the analysis of research outputs. *Research Evaluation*, 6, 25-36.
Lewison, G. (1999) The definition and calibration of biomedical subfields. *Scientometrics*, 46, 529-537.
Leydesdorff, L. (2008) The delineation of nanoscience and nanotechnology in terms of journals and patents: a most recent update. *Scientometrics*, 76, 159-167.
López-Illescas, C., de Moya-Anegón, F. & Moed, H.F. (2008) The actual citation impact of European oncological research. *European Journal of Cancer*, 44, 228-236.
Maghrebi, M., Abbasi, A., Amiri, S., Monsefi, R. & Harati, A. (2011) A collective and abridged lexical query for delineation of nanotechnology publications. *Scientometrics,* 86, 15-25.
May, R.M. (1997a) The scientific wealth of nations. *Science,* 275, 793-796.
May, R.M. (1997b) Government funding of research and development – response. *Science,* 278, 879-880.
Meyer, M., Debackere, K. & Glänzel, W. (2010) Can applied science be "good science"? Exploring the relationship between patent citations and citation impact in nanoscience. *Scientometrics,* 85, 527-539.
Nederhof, A.J. & Visser, M.S. (2004) Quantitative deconstruction of citation impact indicators – waxing field impact but waning journal impact. *Journal of Documentation*, 60, 658-672.
Podlubny, I. (2005) Comparison of scientific impact expressed by the number of citations in different fields of science. *Scientometrics,* 64, 95-99.
Prathap, G. & Gupta, B.M. (2009) Ranking of Indian universities for their research output and quality using a new performance index. *Current Science*, 97, 751-752.
Prathap, G. (2010) An iCE map approach to evaluate performance and efficiency of scientific production of countries. *Scientometrics*, 85, 185-191.
Qiu, J.P., Ma, R.M. & Cheng, N. (2008) New exploratory work of evaluating a researcher's output. *Scientometrics*, 77, 335-344.
Schubert, A. & Braun, T. (1996) Cross-field normalization of scientometric indicators. *Scientometrics,* 36, 311-324.
Schwechheimer, H. & Winterhager, M. (2001) Mapping interdisciplinary research fronts in neuroscience: a bibliometric view to retrograde amnesia. *Scientometrics*, 51, 311-318.
Shapira, P., Youtie, J. & Mohapatra, S. (2003) Linking research production and development outcomes at the regional level. *Research Evaluation*, 12, 105-116.
Sooryamoorthy, R. (2011) Scientific publications of engineers in South Africa, 1975-2005. *Scientometrics,* 86, 211-226.

Tores-Salinas, D., Lopez-Cozar, E.D. & Jimenes-Contreras, E. (2009) Ranking of departments and researchers within a university using two different databases: Web of Science versus Scopus. *Scientometrics*, 80, 761-774.

Ugolini, D. & Mela, G.S. (2003) Oncological research overview in the European Union. A 5-year survey. *European Journal of Cancer*, 39, 1888-1894.

Webster, B.M. (2005) International presence and impact of the UK biomedical research, 1989-2000. *Aslib Proceedings*, 57, 22-47.

Yi, C.G. & Kang, K.B. (2000) Developments of the evaluation system of government-supported research institutes in Korean science and technology. *Research Evaluation*, 9, 158-179.

Zhou, P., Thijs, B. & Glänzel, W. (2009) Regional analysis on Chinese scientific output. *Scientometrics,* 81, 839-857.

Zitt, M., Ramanana-Rahary, S. & Bassecoulard, E. (2005) Relativity of citation performance and excellence measures: from cross-field to cross-scale effects of field-normalisation. *Scientometrics*, 63, 373-401.