Indicator-Assisted Evaluation and Funding of Research: Visualizing the Influence of Grants on the Number and Citation Counts of Research Papers

Kevin W. Boyack
Sandia National Laboratories*, P.O. Box 5800, Albuquerque, NM 87185, E-mail: kboyack@sandia.gov

Katy Börner
School of Library and Information Science, Indiana University, Bloomington, IN 47405. E-mail: katy@indiana.edu

Abstract
This paper reports research on analyzing and visualizing the impact of governmental funding on the amount and citation counts of research publications. For the first time, grant and publication data appear interlinked in one map. We start with an overview of related work and a discussion of available techniques. A concrete example – grant and publication data from Behavioral and Social Science Research, one of four extramural research programs at the National Institute on Aging (NIA) – is analyzed and visualized using the VxInsight® visualization tool. The analysis also illustrates current existing problems related to the quality and existence of data, data analysis, and processing. The paper concludes with a list of recommendations on how to improve the quality of grant-publication maps and a discussion of research challenges for indicator-assisted evaluation and funding of research.

1. Introduction
In his book *The Structure of Scientific Revolutions*, Kuhn (1962, p. 4) points out that a scientific community cannot practice its trade without some set of established beliefs. In the case of research evaluation, the need to use human experts in a review and assessment process is an established belief, as evidenced by the prevalence of funding processes employing human peer review today. To their credit, many experts have established their own semi-quantitative methods for evaluation, which however, rely on desired attributes from institutional statements of need, personal biases, and perception of past performance, rather than actual quantitative measures. Despite this, expert input is often subjective and is frequently acknowledged as being "toothless," making it hard to objectively identify (interdisciplinary) research with high socio-economic benefit. Yet, limited resources require the setting of strategic priorities.

To solve this crisis, we propose a (paradigm) shift from experts working with their bare hands and intellects to experts utilizing advanced data analysis and visualization techniques. Very much like a calculator improves a human's computing capabilities, these techniques can be used as a tool to sift through and analyze very large amounts of data rapidly, and to explore findings interactively and understandably. For example, these techniques can help to objectively identify major research areas, experts, institutions, grants, publications, or journals in a research area of interest. In addition, they can assist in the identification of interconnections, the import and export of research between fields, the dynamics (speed of growth,

*** Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under Contract DE-AC04-94AL85000.**

diversification) of scientific fields, scientific and social networks, and the impact of strategic and applied research funding programs. This knowledge is not only interesting for funding agencies but also for companies, researchers, and society.

According to Kuhn (1962, p. 5), normal science "is predicated on the assumption that the scientific community knows what the world is like", i.e., that peer review by experts is the best way to evaluate research proposals. It is expected that scientists and government workers will take great pains to defend that assumption. However, we don't suggest replacing experts by automated techniques. Instead we propose to augment and accelerate the expert's intellect by the utilization of efficient tools. Efforts in this regard have already begun in the Netherlands, where bibliometric indicators are being used alongside the results of traditional peer review (Rinia, van Leeuwen, van Vuren, & van Raan, 1998).

While there is an increasing number of companies and national research laboratories that utilize commercially available systems[1] for science and technology management, to our knowledge neither the National Science Foundation, the National Institutes of Health, nor the National Academy of Sciences use data mining or visualization on a regular basis to aid in their decision making or to make resulting findings available to their researchers.

The proposed shift is timely as it is facilitated by the explosion of information available digitally (e.g., publications, grants, patents) in digital libraries, repositories, and the WWW; the decreasing cost of storage and computing power; fast graphics processors; and scalable data analysis and visualization algorithms.

The next section introduces diverse types of techniques available to map scientific progress. In general these techniques "measure" research vitality and productivity by determining and counting major research areas, experts, institutions, grants, publications, citations, and journals. They also seek to measure trends, interconnections, the import of research from other fields as measured via citations; the export factor of areas via references from other areas; and the relative speed of areas by means of time series, which helps to identify the most dynamic or static areas as well as new areas. Resulting visualizations typically show authors, papers or journals and their interconnections.

What makes the work presented here unique is the generation of interactive visualizations that show grant and publication data in one map allowing one to relate the dollar amount spent to the number and impact of the results. In addition, we present a set of recommendations and research challenges on how to improve grant-publication maps as a means to augment the management of science and technology (S&T).

2. Related Work

Starting with an explanation of the steps involved in the analysis and visualization of scientific areas, this section reviews qualitative and quantitative work aiming to determine the vitality of research areas, and to detect trends. We also discuss research on so called input-output studies that aim to relate research resources and quality of output.

The first step in a domain analysis or assessment of research vitality is the selection of a database (or databases) appropriate to the field in terms of subject specificity and breadth of coverage. Many different literature, patent, project, grant, and research opportunity databases are pertinent to assessment of science and technology areas. Examples of these include: INSPEC (http://www.iee.org/Publish/INSPEC/), Medline (http://www.ncbi.nlm.nih.gov/pubmed/), NEC's ResearchIndex (http://citeseer.nj.nec.com/cs), ISI Citation Indexes (http://www.isinet.com/ISI), EI Compendex (http://www.ei.org/ev2/home), Cambridge Scientific Abstracts (http://www.csa.com/), Chemical Abstracts (http://www.cas.org/SCIFINDER/SCHOLAR), NIH grants (http://commons.cit.nih.gov/crisp3/), NSF funding programs (http://www.nsf.gov), research opportunities (http://www.cos.com), US Patent and Trademark Office (http://www.uspto.gov/), Derwent World Patents Index (http://www.derwent.com/), and arXiv (http://arxiv.org/). The databases come in diverse formats and coverage. The ease and cost of raw data access differs widely. Recent standardization efforts such as the Open Archives Initiative (http://www.openarchives.org/) that develop and promote interoperability standards for e-print data will help facilitate the efficient access and utilization of digital material via value-added services.

Bibliometric Measures and Indicators

Derek J. deSolla Price was the first to examine the major transformation in the structure of science in his book entitled *Little Science, Big Science* (1963) and he laid out the foundations of the quantitative analysis of science and scientific development, called scientometrics or bibliometrics (see also the review by White and McCain (1989)).

Martin and Irvine (1983) conducted the first evaluation of 'big science' facilities using 'converging partial indicators', i.e., assessing the number of publications and citation counts for their degree of convergence. Both also pioneered the notion of 'foresight' as a tool for looking into the longer-term future of science and technology with the aim of identifying areas of research and technology likely to yield the greatest benefits (Irvine & Martin, 1984).

van Raan and co-workers at the Centre for Science and Technology Studies (CWTS), University of Leiden, conduct research performance assessment using advanced bibliometric methods. They point out that when mapping the socio-economical state of our society it is necessary to monitor both current S&T developments and those that may be of vital importance in the near future (van Raan, 1996).

---

[1] Example systems include VxInsight (http://www.sandia.gov/VxInsight or http://www.viswave.com/), SemioMap (http://www.semio.com/), VantagePoint (http://www.thevantagepoint.com/), and Internet Cartographer (http://www.inventix.com/).

Narin, Olivastro, & Stevens (1994) categorize bibliometric methods into activity measures, impact measures, and linkage measures that are explained and exemplified subsequently. King (1987) also reviews many of these indicators and their role in research evaluation.

Activity measures refer to counts of publications or patents, by topical area or institution over time. The number of publications produced by a researcher or group over time is the simplest indicator available. Although it does not provide an indication of quality, it does correlate reasonably well to other measures such as funding and peer ranking (King, 1987), and is thus commonly used.

Impact measures, such as citation counts, allow one to find out where and how often an article is cited. This provides an estimation of the importance of an article. Citation statistics are widely used for the allocation of funds, promotion and tenure decisions, and determining research influence. The number of references to a scientific paper or book generally peaks between two to five years after publication. Consequently, journal impact factors providing average citation rates for all papers published in a particular journal are used for younger papers. While the quality and impact of papers published in one journal may vary, the journal impact factors are simpler and less labor intensive to use and avoid the 2-5 years delay needed to produce meaningful citation counts, thus enabling timely results. The citation half-life (the length of time from publication to account for 50% of the citations received) can also be used to show the length of impact of seminal publications.

Linkage measures provide evidence of intellectual associations and are typically based on co-occurring words or citation links. These first two types of linkage measures are commonly and frequently used to determine similarity among documents, authors, terms, or journals, and have been described in detail elsewhere (Börner, Chen, & Boyack, 2003; White & McCain, 1997).

Interestingly, the set of retrieved documents based on followed citation links has very little overlap with the document set retrieved based on keywords (Pao & Worthen, 1989). In a similar study, McCain (1989) studied the overall performance of descriptor and citation retrieval as part of a Medline indexing evaluation project. The result was that there was little overlap between the two sets of relevant documents, one retrieved by descriptors and one retrieved by citations. Consequently, ISI defines new 'Research Frontiers' based on a mix of co-citation and co-word analysis, where the scope of the mix is adjustable by increasing or decreasing the threshold strength that refers to "the degree of association between co-cited pairs in terms of the proportion of their total citations that are co-citations."[2]

More interesting are new types of linkage measures. One, recently introduced by Kleinberg (1999), defines 'hubs' and 'authorities' to characterize the way in which a large 'community' of thematically related information sources links and refers to its most central, prominent members. Two types of nodes are distinguished: 'authorities' have a large number of incoming links from hub nodes, and 'hubs' link to many authorities. A recursive eigenvector-based algorithm is used to identify these hubs and authorities, of which multiple groups can exist in a given set of documents. Hubs act as high-quality guides directing users to recommended authorities. Authorities resemble high quality web pages or review articles. Authority and hub ratings can be used as linkage or impact measures.

A final type of linkage measure establishes relationships among different units, e.g., publications and grants. Studies using multiple units are rare, the author co-citation analysis by White and Griffith (1981) being the only one of which we are aware. Unfortunately, relationships between different types of units are rarely available in a complete and consistent form. If determined semi-automatically, then the lag time between grant duration and years of publication has to be determined and compensated for. Lewison, Dawson, and Anderson (1995) report on the behavior of authors in acknowledging their funding resources to be used for evaluation and policy-making purposes and conclude that acknowledgement depends heavily on the level of support given by the funding body.

Partial indicators of scientific performance (relying on publications, patents, R&D expenditures, equipment, and software as well as on case studies) have been used for R&D evaluation, research vitality assessment, technology opportunity analysis, and to set research priorities. To be successful, partial indicators need to be ranked and interpreted together with peer-ratings.

An automated approach (Zhu & Porter, 2002) to generating many indicators for a particular science or technology area is being perfected at the Technology Policy and Assessment Center (TPAC) at the Georgia Institute of Technology. Sample analyses are available at their website (http://tpac.iac.gatech.edu/hottech/).

Research Vitality Studies

Governmental institutions, companies, researchers, and society are interested in funding the most vital research areas, i.e., areas that promise the highest socio-economical benefits. While most companies need to focus on short term benefits and payoffs, grant agencies and tenured faculty have the luxury of supporting evolving research areas, which can aid in the merging of two areas that appear mutually beneficial. They can also support basic research with long term impact on more applied research. Consequently, companies typically fund highly vital research and development areas that promise high profit within a few month/years. Governmental agencies aim to steer the development of a larger research area, and in many cases can fund research areas that are not yet vital.

Keeping this in mind, we seek to define vital research areas that show some, but not necessarily all, of the following features:

[2] http://www.isinet.com/isi/hot/essays/citationanalysis/11.html

- A stable/increasing number of publications in prominent journals with high impact factors
- High export factors indicating that research is acknowledged and utilized in other domains
- A tightly knit co-authorship network leading to efficient diffusion of knowledge
- Funding resulting in larger numbers of high impact publications
- New emerging research fields

Input-Output Studies

To date, few studies have attempted to correlate research outputs with inputs. McAllister and Wagner (1981) studied the relationship between research and development (R&D) expenditures and publication output for US colleges and universities. Halperin and Chakrabarti (1987) examined the relationship between the quality of scientists and key financial characteristics of the corporations in which they work and the volume of scientific and technical publications. Results indicate a strong correlation between patenting and publications; firms with high annual sales produce proportionally fewer papers than small firms; and the number of elite scientists is more highly correlated with publications than with patents.

More recently, several studies have been done to investigate the influence of government funding on research output, giving a variety of results. Lewison and colleagues (Lewison, 1998; Lewison & Dawson, 1998; Lewison & Devey, 1999) report a number of studies on the impact of funding resources on a national level on research output in the fields of gastroenterology and arthritis research. Jain, Garg, Sharma and Kumar (1998) compared the output of SERC's funded project investigators to the Indian chemical sciences community as a whole, and found their output and impact to be higher as a result of the funding. Cronin and Shaw (1999) examined the impact of funding in four information science journals, and determined that citedness was not correlated with funding, but rather with journal of publication and the nationality of the researcher. Bourke and Butler (1999) report on the efficacy of different modes of funding research in biological sciences in Australia, concluding that research from full time researchers receives considerably higher citation counts. Their work related funding to the sector level. Butler (2001) followed this work up with a study of funding acknowledgement, finding that while acknowledgement data on the whole accurately reflected the total research output of a funding body, there was no ability to track research back to the grant level.

This inability to track research back to an individual grant precludes analyses of research vitality at the finest levels. Indeed, we are unaware of any published study which has been able to do so. In addition, none of the input-output studies of government funding have used visualization to try to show trends. In this study we start that process, and identify problems associated with data and the process that will need to be fixed before conclusive studies of vitality and impact at the grant level can be performed.

Science and Technology Maps

In addition to the types of measures and indicators described above, there have been recent efforts to produce interactive maps of science and technology areas. These maps use as their source the same types of data used for indicator studies. After data extraction, a unit of analysis needs to be defined (e.g., author, document, journal, grant, term), appropriate measures have to be selected (e.g. frequency counts of terms, citations, co-authorship, thresholds), the similarity/distance between units need to be calculated, and coordinates have to be assigned to each unit for special layout (called ordination). Interaction techniques need to be designed to facilitate an intuitive overview and navigation, rapid filtering of relevant data, and the display of details on demand (Shneiderman, 1996). These different steps as well as currently available techniques are reviewed in detail in Börner, Chen & Boyack (2003). The process concludes with the use of the resulting visualization for analysis and interpretation.

3. Mapping Behavioral and Social Science Research

This section presents results of a recent demonstration study conducted for the Behavioral and Social Research (BSR) Program[3], one of four extramural research programs at the National Institute on Aging (NIA). BSR supports training and basic social and behavioral research on the processes of aging at both the individual and societal level. The current structure and funding patterns of BSR reflect the current scientific paradigm and the issues and key research questions that BSR officials feel are pertinent today. For BSR, these issues and questions all have to do with the demographic, economic, social, psychological, and cognitive aspects of human aging, rather than with the specific diseases and biology of aging that are addressed by the other extramural research programs within NIA.

Another recent study has sought to map the entire field of human aging. Noyons and van Raan (2001) at CWTS have produced an interactive web map of human aging literature comprised of aging related papers from the Current Contents database from 1995-2000. Their data extraction was based on an extensive journal and keyword list. Clusters of documents were generated using keyword co-occurrence. The web interface to their maps is a wonderful development, allowing interested parties to see what authors, institutions, journals, and terms are associated with various subdomains in a scientific field. In contrast to the current work, the CWTS map does not include any grant data and thus does not show any

[3] http://www.nia.nih.gov/research/extramural/behavior/

relationships to funding. Perusal of the CWTS aging map shows that the areas of interest to BSR are an admittedly small portion of the overall field of human aging.

Data Acquisition & Selection

Data to create a map of human aging from an NIA/BSR perspective were received from two main sources: NIA grant data and BSR accomplishment reports. Researchers at ISI were involved in a similar demonstration study, and made their citation data available to us for this study.

Grant Data

The complete data set of grants funded by NIA covering the years 1975-2001 were supplied to us by NIA. Data for each record included grant number, sub-grant number (when applicable), principal investigator (PI) name, institution, funding year, award amount, title, abstract, and NIA supplied MeSH (Medical Subject Heading) terms. Including sub-grants, and noting that each grant and any corresponding sub-grants were listed separately for each year of their existence, the data contained a total of 33,448 records. Figure 1 shows that the total amount of research funded through NIA has increased significantly over the past 20 years to nearly $600M in FY2000. Over the same period of time, the average grant amount (per year) has grown much more slowly as the number of projects receiving funding each year has increased. Grant data for the BSR program alone could not be split out from the NIA grant data due to changing organization structure and changing program codes over the years.

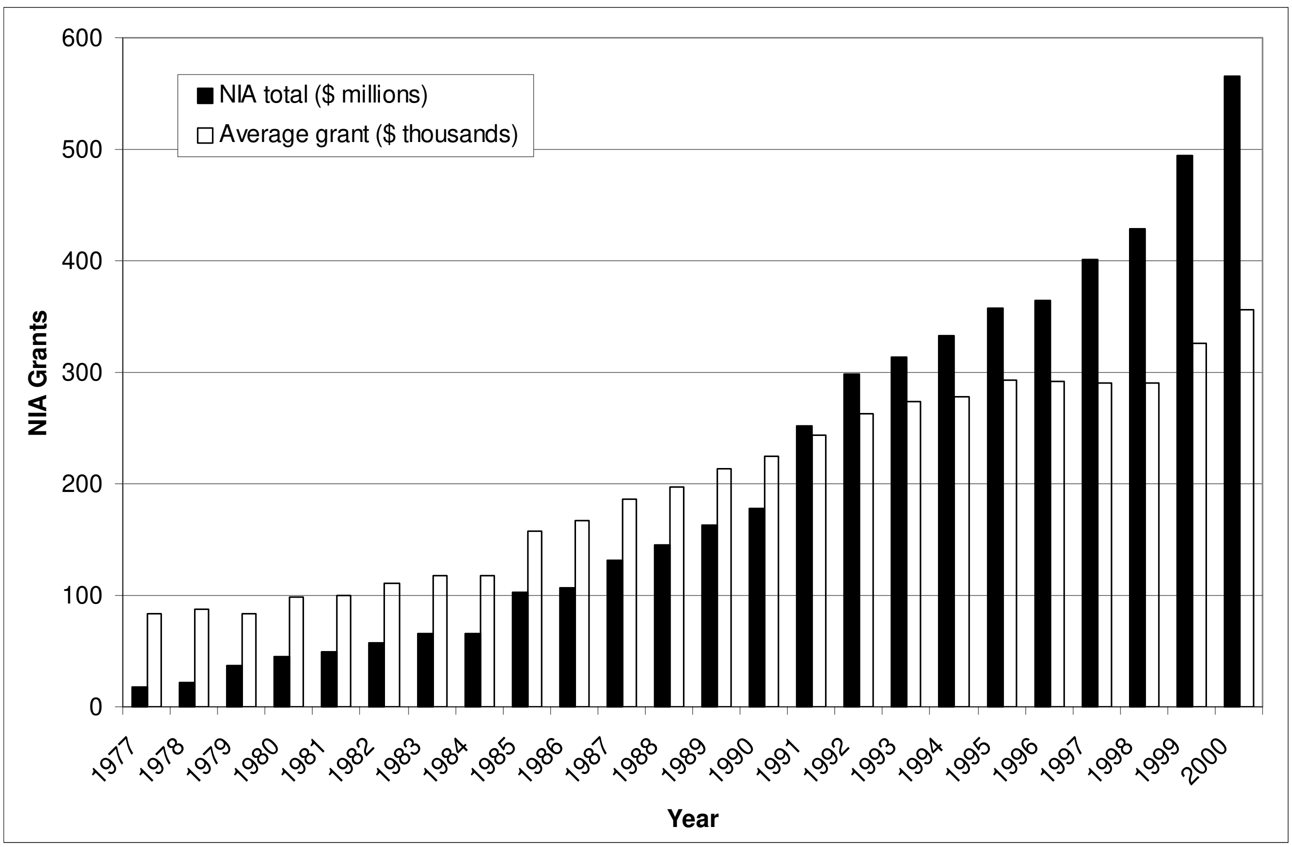

Figure 1. NIA total grant awards and average grant amounts by year.

Grant data current back to 1972 are publicly available online via CRISP[4], a searchable database of federally funded biomedical research projects conducted at universities, hospitals, and other research institutions.

Publication Data

BSR Data. Information on documents that BSR considered to have resulted from work funded through their extramural programs were made available to us in five accomplishment reports, each one corresponding to a different focus within BSR

[4] http://commons.cit.nih.gov/crisp3/Crisp_Query.Generate_Screen

as shown in Table 1. Note that although the headings for each of the five areas may not correspond to the current structure of BSR, as a whole they do cover the structure.

Table 1. NIA Behavioral and Social Research (BSR) accomplishment reports and related information.

	Years	# records	Grant #'s?	PI's?
Cognitive Functioning and Aging (**CFA**)	1995-1999	902	yes	yes
Demography and Population Epidemiology (**DPE**)	1992-2000	1462	no	no
Health Care Organizations and Social Institutions (**HCO**)	1995-1999	878	no	yes
Behavioral Medicine and Public Health (**PHBM**)	1995-1999	626	no	yes
Personality and Social Psychological Aging (**PSP**)	1995-1999	681	yes	yes

The accomplishment data from BSR were supplied in a bibliography-like form (see Figure 2) and included peer reviewed journal articles, conference papers, book chapters, books, encyclopedia articles, and other types of documents. Documents were in some cases listed by author, in other cases by the type of publication in which they appeared. The data from each of these five files were parsed into a common format suitable for combining with grant data, such that both could be placed in a single database. Parsing could not be totally automated for any of the five files since the bibliographic entries appeared in many different formats (see our recommendations in a later section). Indeed, formats changed within accomplishment reports by author, and seemed to correspond to whichever format and journal abbreviation a particular author was most fond of. Data cleaning and merging was accomplished using functions available in Excel or Word as well as simple parsing programs.

Carolyn M. Aldwin **University of California -- Davis**
AG13006 **Mental and Physical Health Trajectories in Adulthood**

Bibliography

Spiro, A. III, Aldwin, C. M., Ward, K. D., & Mroczek, D. K. (1995). Personality and the incidence of hypertension among older men: Longitudinal findings from the Normative Aging Study. *Health Psychology, 14,* 563-569.

Aldwin, C. M., Ober, B. A., & Levenson, M. R. (1996). The psychology of aging. In D. B. Reuben, T. T. Yoshikawa, R. W. Besdine, & B. R. Popecki (Eds.). *Geriatrics Review Syllabus (GRS-III)* (pp. 24-29). New York: American Geriatrics Society.

NIA Behavioral and Social Research Program
DEMOGRAPHY AND POPULATION EPIDEMIOLOGY UNIT
Partial Grantee Bibliography

Publications, 1992-present, in top journals in science (Science, Nature, PNAS) or medicine (NEJM, Lancet, JAMA, BMJ, Cell).

Bruce M. Alberts, Francisco J. Ayala, David Botstein, Ellen Frank, Edward W. Holmes, Ronald D. Lee, Eduardo, R. Macagno, Philippa Marrack, Suzanne Oparil, Stuart H. Orkin, Arthur H. Rubenstein, Carolyn W. Slayman, P., Frederick Sparling, Larry R. Squire, Peter H. von Hippel, and Keith R. Yamamoto. "Grant Review Process. Proposed Changes for NIH?s Center for Scientific Review" Science 1999 July 30; 285: 666-667. (in Policy Forum)

Bosma H, Marmot MG, Hemingway, H et al. 1997. Low Job Control and Risk of Coronary Heart Disease in Whitehall II (Prospective Cohort) Study. BMJ 314(7080): 558-65.

Figure 2. Example extracts (PSP and DPE, respectively) from BSR accomplishment reports.

A total of 4549 records were contained in the five accomplishment reports. An effort was made to remove duplicate records using common titles and sources. Duplicate records were attributed to several causes: the same publication being listed by multiple authors, the same publication belonging to multiple BSR areas, and the same publication being listed twice by the same author at different times (e.g. "In press" at one time and with the actual publication year at another). 546 of the records appeared to be duplicates and were removed, leaving 4003 unique records. Of these 2903 appeared to be journal articles or conference papers, while the remaining 1100 came from other sources.

It is instructive to see what journals are targeted for publication by researchers funded by BSR. A list of the top 30 journals represented in the BSR accomplishment data is shown in Table 2. Nearly half of the 2903 journal articles and conference papers are represented by the journals in the table, with 2521 in journals covered by the ISI Science Citation Index or Social Science Citation Index, and 1884 in journals covered by Medline. Given that a high percentage (~87%) of the BSR journal articles, including all of those in the top 30 journals, are available in ISI's indexes, a direct measurement of citation counts and impact could certainly be made by matching each BSR publication with its corresponding ISI record. However, we did not go to this effort since it would have required a great deal of manual data extraction and matching.

Although Medline has records for a smaller percentage (~65%) of the BSR publications, and does not have citation counts, it has the advantage of being a free source of data, and could thus be used to correlate abstracts and descriptors such as MeSH terms to BSR areas. This could be done using ISI data as well. A look at the full list of journals represented in the BSR publication data shows that journals related to the family and to economics are notably missing from Medline. Other key aging journals such as *Research on Aging* and *Aging Neuropsychology and Cognition* are also not available through Medline.

Table 2. Top journals represented in the BSR publication data (* 3-year impact factor)

	Number	Medline	ISI	IF9599
Gerontologist	165	yes	yes	1.73
Journals of Gerontology Series B - Psychological Sciences and Social Sciences	160	yes	yes	1.33
Psychology and Aging	102	yes	yes	2.03
Journal of the American Geriatrics Society	81	yes	yes	2.66
Journals of Gerontology Series A - Biological Sciences and Medical Sciences	75	yes	yes	1.02
Journal of Marriage and the Family	48	no	yes	1.43
American Journal of Public Health	46	yes	yes	3.23
Research on Aging	40	no	yes	0.61
American Journal of Epidemiology	40	yes	yes	3.86
Journal of Aging and Health	37	yes	yes	0.85*
Demography	33	yes	yes	1.70
American Economic Review	30	no	yes	1.77
Health Psychology	29	yes	yes	2.67
JAMA - Journal of the American Medical Association	28	yes	yes	9.44
International Journal of Aging & Human Development	23	yes	yes	0.48
Experimental Aging Research	21	yes	yes	0.60
Social Science & Medicine	21	yes	yes	1.30
Medical Care	21	yes	yes	2.24
Journal of Family Issues	20	no	yes	0.86
Journal of Clinical Epidemiology	20	yes	yes	1.80
Aging Neuropsychology and Cognition	19	no	yes	0.86*
Journal of Human Resources	19	no	yes	1.17
Annals of Behavioral Medicine	19	yes	yes	1.47*
Journal of Health And Social Behavior	19	yes	yes	2.53
Neurology	18	yes	yes	4.80
Journal of Applied Gerontology	17	no	yes	0.38
Psychosomatic Medicine	17	yes	yes	2.94
Archives of Internal Medicine	17	yes	yes	5.14
Science	16	yes	yes	23.84
Journal of General Internal Medicine	16	yes	yes	1.80

JCR impact factors averaged over the five years from 1995-1999, abbreviated IF9599, have been included in Table 2 for reference. The average impact factor for all articles represented by the journals in Table 2 using the IF9599 value as a surrogate for the correct years is 2.28. If papers in *JAMA* and *Science* are excluded, the average impact factor is only 1.82.

King (1988) shows that relative ranking of groups using impact factors (or expected citations) is the same as that obtained using actual citation counts. However, it is interesting to note that the majority of journals with an impact factor greater than 2 are medical journals rather than aging journals. Thus, aging work that can be published in medical journals may have a higher impact than that published in the aging and gerontology journals.

ISI Data. The Institute for Scientific Information extracted data for their demonstration project based on criteria supplied by BSR, and made those data available to us for this study. Data from ISI should provide a good view of research in BSR given that a large percentage of significant scientific results are published in a relatively small number of journals, most of which are available in ISI's indexes.

ISI's data set was based on the works of 32 merit awardees whose names were supplied by NIA. ISI extracted 2296 papers authored by these scholars between 1981 and July 2001. The data were hand checked by ISI to ensure that they were authored by the correct scholars and not by others with the same name. All were found to be relevant to human aging research. ISI then extracted 26,880 indexed papers that cite the merit awardee authored papers. Some of the citing papers were also in the group of core papers, thus a total of 27,851 papers were included in this data set. ISI also extracted papers cited by the group of core papers, but they have not been included in this analysis. Keywords and abstracts were not extracted by ISI for the work done here, and thus could not be used in the analysis. However, citation counts by paper were included with the data, and were used in the analysis, as will be shown later.

A data set based on the work of a number of highly cited scholars makes for an interesting analysis, but it is not certain that such a strategy provides proper coverage of the BSR or NIA aging field as a whole. Given that merit awardees must have a long track record with NIA to receive such a status, it is likely that this approach weights the data to older work, and excludes younger researchers whose work would bring them a merit award status in the future. On the other hand it could be argued that much of the current work builds upon the work of the merit awardees, and thus provides a good representation of the field. The current analysis will not settle this question.

Linking Grant and Publication Data

Not all of the grant data were directly related to the publication data supplied by BSR. Potentially relevant grants were extracted from the full list by finding all of the grants and subgrants having PI's in common with those from the BSR publication data. This was done for the CFA, HCO, PHBM, and PSP areas, since PI's were listed with the publication data (see Table 1). However, for the DPE area, no PI's were listed in the accomplishment report. Thus, relevant grants for DPE were extracted another way. Here, we found all grants whose PI matched any author in the list of DPE publications. For the five BSR areas, a total of 5284 records corresponding to 818 individual grant numbers were thus extracted. A single record for each grant number was thus used in a combined publication/grant data set. Grant duration in years, and average annual and total grant amounts were calculated and included with each grant record.

Use of a single record for each grant created some problems. These include the facts that many were multi-year grants with several subgrants or subcontracts, each having its own title. In addition, the title of the main grant often changed from year to year, thus giving us a choice of which title to use. The most recent title for the main grant for each grant number was used in all cases. The PI on a grant also changed from time to time, although this occurred less frequently than title changes.

The BSR publication data and grant data were combined into one data set for visualization using the field structure shown in Table 3. In addition to these fields, a separate table was made available with the following information for the grants: abstract, MeSH terms, initial review group (IRG), StartYear, EndYear, duration, total grant amount, and average annual grant amount.

The merged BSR grant and publication data set had a total of 4821 records, 818 of which were grants, and 4003 of which were publications.

Table 3. Field structure for combined BSR publication and grants data

Field	Pubs	Grants
Record_ID	yes	yes
Pub_Type	yes	yes
BSR_Area	yes	
Grant_No		yes
Year	yes	yes
PI_Name	yes	yes
Institution	yes	yes
Authors	yes	
Source	yes	
Title	yes	yes

Grant and publication data were linked based on the accomplishment reports provided by BSR. By a "link" we mean a citation-like connection from a publication (citing document) to a grant (cited document). We used two types of links in this study: author-supplied and inferred. Author-supplied links are those where a specific grant number (or numbers) was specified in the BSR accomplishment data as contributing to the work of a particular PI. An example of the information providing this type of link is shown in the first example in Figure 2, where the grant number AG13006 is associated with all of the papers authored by Carolyn Aldwin. These author-supplied links were generated from the accomplishment data for the CFA and PSP areas. These were the only two accomplishment reports containing grants numbers (see Table 1).

In many cases more than one grant number was listed for a particular PI without any distinction as to which publications in the list corresponded to which grant. In those cases we assumed that each grant contributed equally to each publication, although this was likely not the case in many instances. The data contained a total of 1803 "author-supplied" links.

Inferred links were generated for the three other BSR areas (DPE, HCO, and PHBM), where no grant numbers were listed. In these cases, we assumed that a link existed between each of the papers listed for a PI and each of the grants for the same PI. Thus these were "inferred" links, of which there were a total of 10859. We realize that there are many problems associated with inferring any connection in the place of real data. Here we give two examples. In the case of a PI with many grants and many papers, it is certain that each grant did not fund work published in all papers. Yet, inferred links assume that this is the case. Inferred links can also place the cart before the horse, or rather the publication before the grant, which is clearly impossible. The accuracy and usefulness of a map generated from those inferred linkages (in the same way that a citation or co-citation map could be generated) would be suspect given the unknown level of accuracy of the links. Thus, we have not generated such a map. However, we have made those links available to be shown on other maps (see e.g., Figure 7), in the hopes that they would provide some useful information in the context of a map based on more defensible relationships (such as a co-term analysis).

There are other issues in working with this type of publication data as well. For instance, funded publications by non-PI co-workers may not be represented in the accomplishment reports. Also, there are a substantial number of books, book chapters, encyclopedia entries, and so forth that are listed in these reports. We assume that these publications are of high quality, yet there are very little if any data available to do the same types of impact or vitality studies on these publications that can be performed on journal articles.

Characterization of the BSR Domain

Map of BSR Grants and Publications

The BSR domain comprised of 818 grants and 4003 publications was analyzed using latent semantic analysis (LSA) (Deerwester, Dumais, Landauer, Furnas, & Harshman, 1990; Landauer, Foltz, & Laham, 1998) and co-word (or co-term) analysis. Each technique was used to generate similarity values for each pair of documents needed to calculate document positions on the map.

A description of the LSA method used here can be found in Börner, Chen, and Boyack (2003). In short, a document by term matrix is generated in which each cell entry denoted the frequency with which a term occurs in a document title. The resulting matrix constitutes the input to an advanced statistical technique, namely singular value decomposition (SVD), which constructs an n dimensional abstract semantic space. LSA models the term-document relationships using a reduced approximation for the column and row space computed by the SVD. Only the most important dimensions, here 68, were kept and used to generate a document-by-document similarity matrix. Then, all similarity values greater than or equal to 0.7 were used as input to the VxOrd force-directed placement clustering algorithm (Davidson, Wylie, & Boyack, 2001), which generated x,y positions for each document for visual display.

As an alternative method, a co-word analysis was used. Recall that MeSH terms assigned by NIA were included with the grant data. A list of unique words found in the MeSH terms for the 818 grants was generated. All of the occurrences of these MeSH words in the titles of the 818 grants and 4003 publications were found and placed in an index. A traditional cosine co-word similarity was then calculated from the [document, MeSH word] index for each pair of documents, and used as input to the VxOrd clustering routine.

Each of the two BSR domain maps, one from the LSA analysis, and one from the co-word analysis, were viewed with the VxInsight database visualization tool (Boyack, Wylie, & Davidson, 2002; Davidson, Hendrickson, Johnson, Meyers, & Wylie, 1998). VxInsight uses a landscape metaphor and portrays the structure of a literature space as mountain ridges of document clusters. The size of a cluster (or peak) and its relative position in the layout provide valuable clues to the role of the cluster in the overall structure. Labels on dominant peaks are based on the two most common words (or, alternately, MeSH terms) in the titles that comprise that peak, thus revealing the content of the various peaks. Users can navigate the map terrain by zooming in and out, querying metadata fields (e.g., titles, MeSH terms, authors, PI's), or by restricting the data displayed to a certain time span and sliding through sequences of years with a slider. Relationships among the data may be displayed and understood at many levels of detail. Detail about any data record is also available upon demand.

The MeSH co-word map provided a more topic-based clustering than the LSA map, perhaps since it was based on a controlled vocabulary and did not have any contribution from 'junk' words that can appear in titles.[5] Thus, we have chosen to concentrate on the MeSH-word based BSR map, which is shown in Figure 3.

[5] Grant titles are often descriptive for the proposed research. This is not necessarily the case for titles of publications, which can lead to a distortion of combined grant-publication maps based on titles. It is expected that LSA applied on abstracts would result in a much higher quality similarity measure between documents, and a correspondingly higher quality visualization and analysis. Unfortunately, only abstracts of grants, and not of publications, were available to us at time of this study.

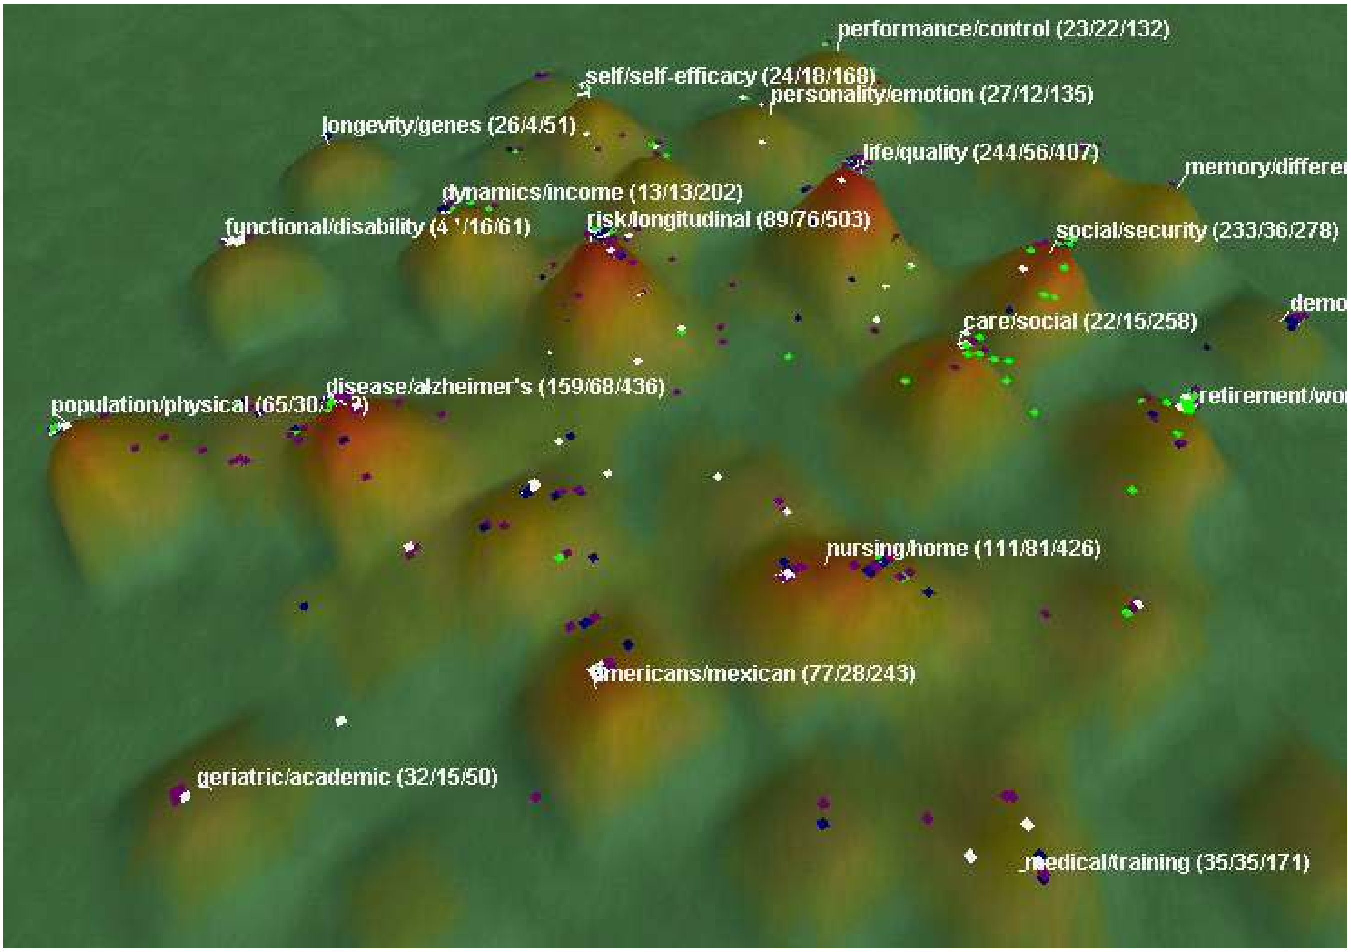

Figure 3. Map of BSR grants and publications using MeSH words as a basis for clustering. The height of each peak is proportional to the number of documents in the peak. Labels show the two most common words in the titles of the documents in each cluster (peak). Query markers for papers with certain words in either the title or MeSH words are shown as colored dots: disability – white, retirement – light gray, economics – dark gray, demography(ics) - black.

This map shows distinct areas of research in topics such as Alzheimer's disease, nursing homes, retirement, functional disability, memory or cognition, personality, population, longitudinal risk, and quality of life, all of which correspond well to the BSR areas of interest. A comparison of the query results (colored markers on the terrain) with the labels shows that papers dealing with economics are associated with many other topics including diseases, quality of life, risk, and nursing homes. Likewise, retirement-related documents are not confined to the "retirement" peak alone, but are associated with social issues, populations and risk. The fact that documents related to a certain topic can be found in many parts of the terrain correlates well with the overlap between the five main BSR areas seen we identified and removed duplicate documents from the data.

The impact of funding is somewhat more difficult to show from the map, but can be done. Figure 4 shows the BSR map in two different time periods, 1995-1996 and 1999-2000. Grants are shown on the map as colored markers above the terrain. Light colored markers correspond to grants of less than $300k per year, while dark markers correspond to grants of greater than $300k per year.

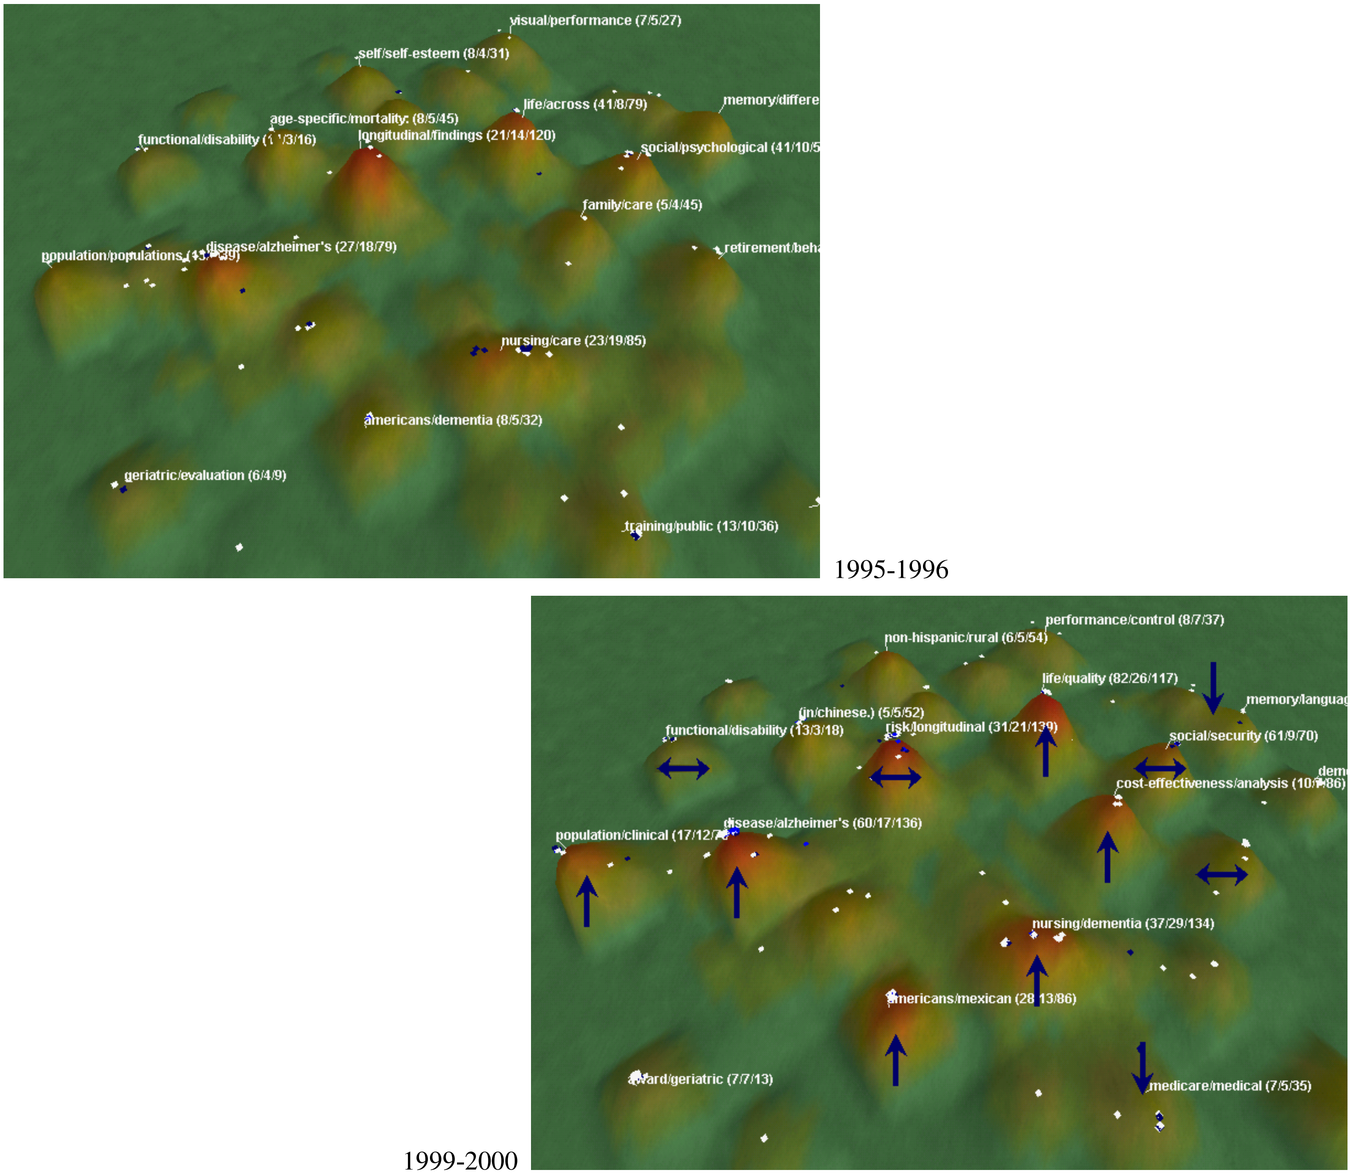

Figure 4. BSR map in two different time periods showing the impact of funding on the number of articles (size of peaks). Small grants are shown by light colored markers on the terrain. Large grants are shown by dark colored markers. Significant changes in the number of publications over time by peak are shown by dark arrows.

In general, peaks of documents containing large grants show significant growth in the number of publications from the 1995-1996 time period to the 1999-2000 time period. This is true for the population, Alzheimer's disease, quality of life, nursing care/dementia, and studies focused on ethnic groups (labeled "americans/mexican") peaks. However, some other correlations can be seen as well. The family care/cost effectiveness peak has few small grants, but shows a significant increase in publication. The longitudinal risk peak seems stable in terms of publication, while receiving many large and small grants. The retirement peak at the right of the map also shows stability in publication, while receiving a few small grants. Two peaks showing a relative downward trend are the memory peak and the training/medicare peak. The memory peak contains only a few small grants, so a downward trend can be understood. However, the medicare peak contains many medium to large grants, and thus might be expected to at least maintain its publication rate. These relative upward and downward shifts in the numbers of grants and publications in different clusters of activity may indicate a shift in the perceived importance by funders and researchers, respectively, of different areas in the BSR-related aging fields of study, and thus may denote shifts in the scientific paradigm.

The BSR map also allows the grants awarded to and publication output of an institution to be shown in the context of the BSR domain. For instance, all documents (grants and publications) for the University of Michigan (blue markers) and Duke University (white markers) during the years 1999-2000 are shown in Figure 5. These two institutions received the highest amounts of NIA funding between 1993-1997. Figure 5 shows the areas of focus for each of the two institutions, as well as how tightly focused their work is in certain areas. For example, Duke University places most of their focus in the population, disease, risk, cost-effectiveness, and memory topics with very little scatter. However, the University of Michigan has a much more scattered portfolio, working on topics such as quality of life, social security, nursing care/dementia, risk, Social Security, and performance. Similar profiles could be projected for any institution for any time period.

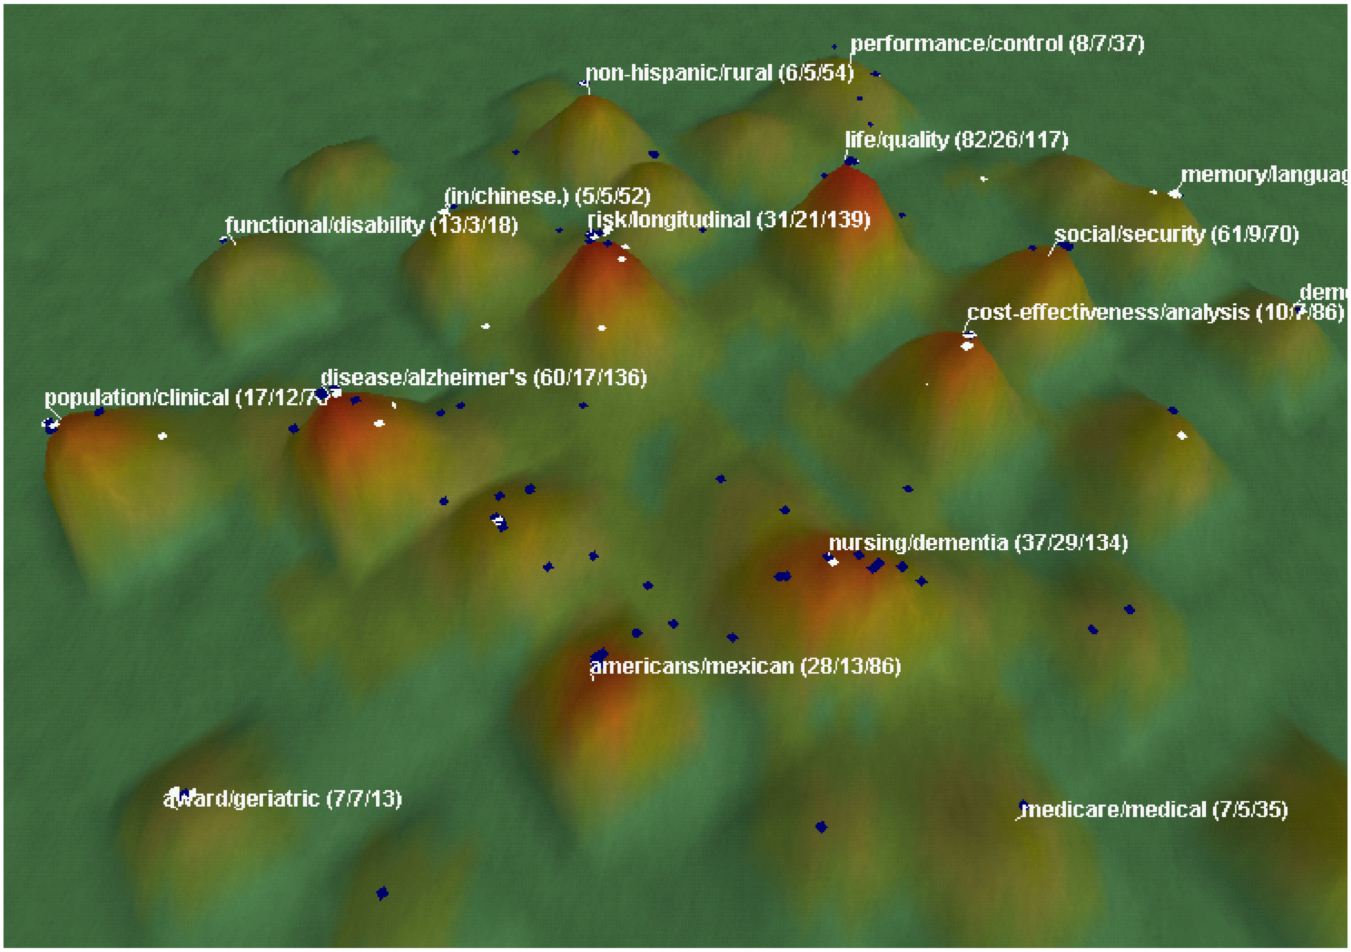

Figure 5. NIA grants awarded to and BSR-related publication output of the University of Michigan (dark markers) and Duke University (white markers) in the BSR domain during the 1999-2000 time period.

Impact of Funding Based on Publications

In addition to generating the BSR domain map shown above, we made an effort to correlate the NIA grant data with citation counts from the ISI data that was supplied to us. The total amount of money funded to various institutions by NIA was calculated for the five-year time period 1993-1997. The top 30 institutions ranked by the amount of NIA money received are listed in Table 4. Average citation counts for two different groups of papers authored from 1995-1999 from the ISI data set are also included in the table. We focus on this five year time period since all five of the BSR accomplishment reports included this time period (see Table 1). We introduce a two-year time lag between grants and publications here to account for the time necessary to perform and publish research. McAllister and Narin (1983) found a very high correlation between the amount of NIH money received and the number of biomedical publications produced two years later by 120 U.S medical schools.

Table 4. Grant and citation data for 30 institutions funded by NIA.

	Grants 1993-1997		Awardee papers 1995-1999		Citing papers 1995-1999	
Institution	# Grant Records	$M	# Papers	Cites / paper	# Papers	Cites / paper
University of Michigan	373	89.9	81	11.73	254	1.76
Duke University	314	69.1	71	13.04	198	1.79
Johns Hopkins University	252	63.3	2	21.00	189	1.40
University of California at San Diego	225	55.1	4	5.25	82	1.23
University of Washington	267	53.7	8	7.63	161	1.47
University of California at San Francisco	218	49.2	29	10.76	184	1.58
University of Southern California	230	48.8	19	10.05	113	1.88
University of Texas (system)	302	43.9	8	12.63	207	1.45
Case Western Reserve University	266	40.8	22	4.77	90	1.67
Washington University	251	36.7	3	8.00	83	1.67
University of California at Los Angeles	165	31.7	10	9.70	228	1.61
Massachusetts General Hospital	169	31.4			46	1.17
University of Pennsylvania	161	30.2	23	9.83	122	1.44
University of Pittsburgh	152	29.3	6	11.50	135	1.54
Columbia University	138	29.0	5	1.60	121	1.57
Rush Presbyterian – St. Luke's Medical Center	119	27.8			33	1.97
University of Kentucky	174	27.6			37	1.41
Boston University	146	27.2	23	5.17	123	1.59
New York University	165	27.0			43	1.47
Harvard University	154	25.5	19	13.16	321	1.60
Indiana University	105	25.1	30	12.43	80	1.61
Mayo Clinics & Mayo Foundation	101	24.8			30	1.27
University of Maryland	158	24.6	2	2.00	90	1.70
Mt. Sinai School of Medicine	150	23.6			71	1.34
University of Wisconsin	108	21.9	3	19.67	133	1.65
University of Colorado	148	21.7	3	26.33	67	1.76
Pennsylvania State University	99	21.2	12	4.67	115	1.54
Stanford University	134	21.0	39	14.36	180	1.63
University of Alabama	121	19.0	12	6.33	95	1.44
Brigham & Women's Hospital	97	17.8			53	1.64
TOTALS / *AVERAGES*	**5462**	*35.3*	**434**	*10.85*	**3684**	*1.58*

Many interesting things can be seen from the information in Table 4. First, we see the institutions where merit awardees were actively publishing large numbers of papers in the late 1990s. Institutions with only a few papers by merit awardees would likely indicate a collaboration with a merit awardee from another institution. We also see many institutions that are getting large grant sums from NIA who have published no papers by merit awardees. Given the peer review process that one must endure to receive funding from NIA (which is a part of NIH, and uses its proposal and funding process), we assume that all of this funding is going to high quality research. Thus, institutions without merit awardees that are high on the list are likely to have younger scholars who some day may receive a merit status.

There is a large amount of scatter in the average citation rates among the merit awardee paper sets as shown in both Table 4 and Figure 6. However, most of the scatter occurs where there are few papers. Institutions with larger numbers of papers tend toward the average citation rate of 10.85 with much less scatter. There are only three merit awardee papers in the data summarized in Table 4 that have been cited over 100 times. Thus, the data are not swayed by just a few highly cited papers.

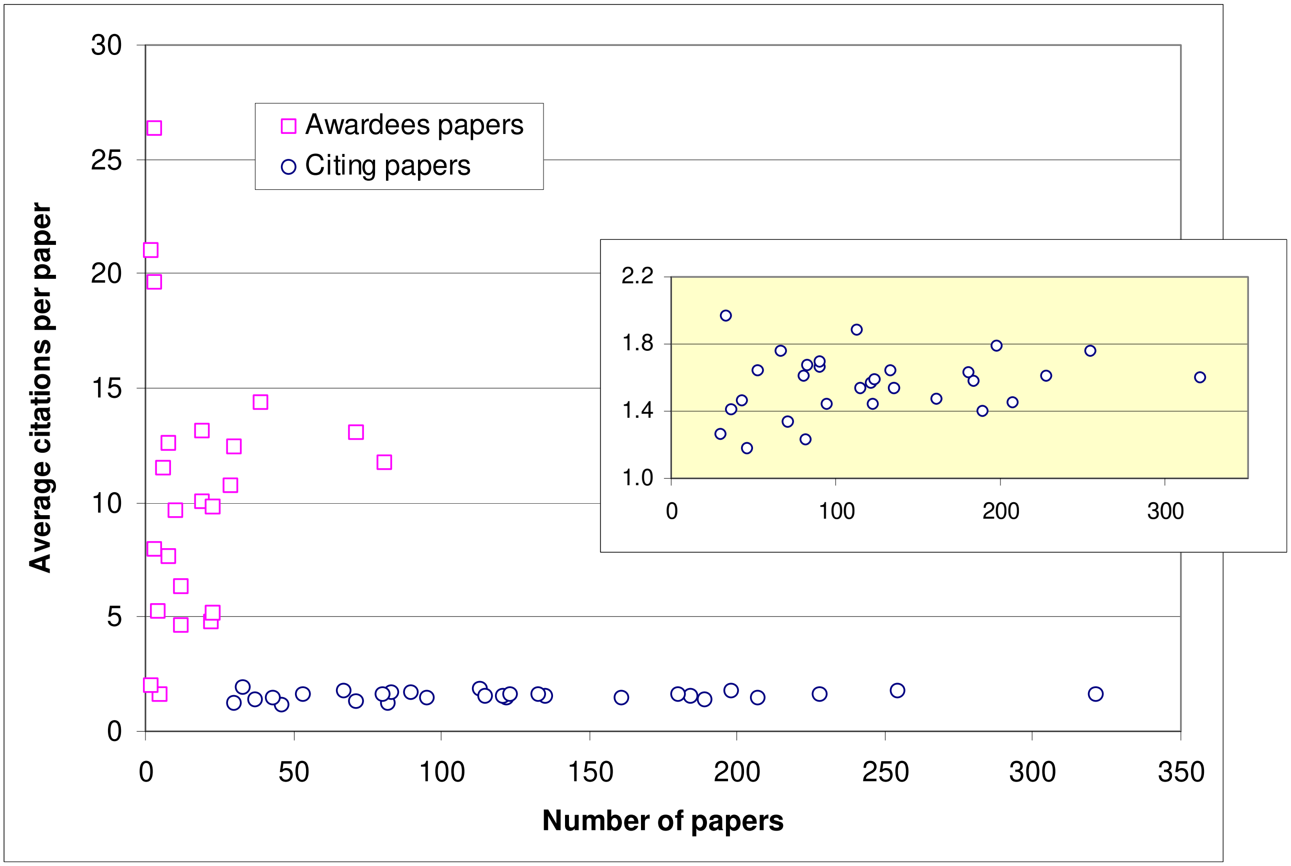

Figure 6. Correlation between the average citations per paper and number of papers for papers authored by NIA merit awardees and authors who cite the merit awardees for the 30 institutions receiving the most funding by NIA.

Perhaps the most notable observation from Table 4 is the difference in average citation rate between papers authored by merit awardees (10.85 cites/paper) and those authored by those who cite the merit awardees (1.58 cites/paper). Such a large difference (a factor of 7) is higher than what we would expect from other literature studies. For example, Bourke and Butler (1999) found differences in expected citation rates between Australian Research Council (ARC) fellows and all grantees corresponding to a factor of 1.5. This leads us to believe that a large part of this difference is an artifact of the data set and the way it was constructed. A test of this hypothesis would require an exact list of all papers published as a result of NIA funding. Such a list may not be possible to obtain due to the effort and cooperation that would be required of a great number of people and institutions.

The data in Table 4 also suggest that there is little correlation between citation rates and funding for this data set. The average citation rate for CITING papers is surprisingly constant across funding and numbers of papers published, and tends toward about 1.6 for institutions publishing large numbers of papers (see Figure 6 inset).

Discussion

The purpose for this study was to investigate research vitality and the influence of grants on publications for NIA/BSR. On balance, and with these data and the effort expended to date, the results are inconclusive. One the one hand, journal impact factors for BSR publications suggest an average impact of near two for journal articles, which supports the argument that either the field of human aging itself or the funding is responsible for a higher than average impact. However, the citation rate data for merit awardees' papers, and for papers citing the merit awardees' papers have such a difference in citation rate that one might think that author status is more responsible for impact. Of course, author status is a chicken and egg type of question – did the author receive status because of the impact of his or her work, or does the work have a seemingly higher impact due to the status of the author.

Qualitatively, the VxInsight map of the BSR domain shows a correlation between grants and increased publication rates in most cases, which qualitatively argues for a certain amount of vitality or momentum in the field.

A case can be made for doing another literature data extraction to try to build a more comprehensive data set to correlate with grant data. An alternate method could be to have BSR contact their PI's and have them explicitly state the correct links between grants and publications in the BSR accomplishment data. In that case, correct citation counts or impact factors could be found from ISI data for each journal article, and a much more accurate correlation between grant and publication data could be obtained.

We conclude our findings with one last picture of how a combined grant and publication map could be used to visualize impact. Figure 7 shows "author-supplied" linkages from grants to publications on the BSR domain map. The links from one particular grant (near the bottom right of the figure) to all of the papers listed by that author as being related to that particular grant are shown as dark arrows. At the macro view shown in the figure, it can be seen that the impact of the one grant spreads across the entire BSR domain. Dark arrows extend to all corners of the map. This particular grant was active from 1993 through 2000 (and perhaps beyond) with an average funding of $560K annually. It is heartening to visually see the impact of this grant upon the BSR domain in terms of links to over 100 publications touching many different topics.

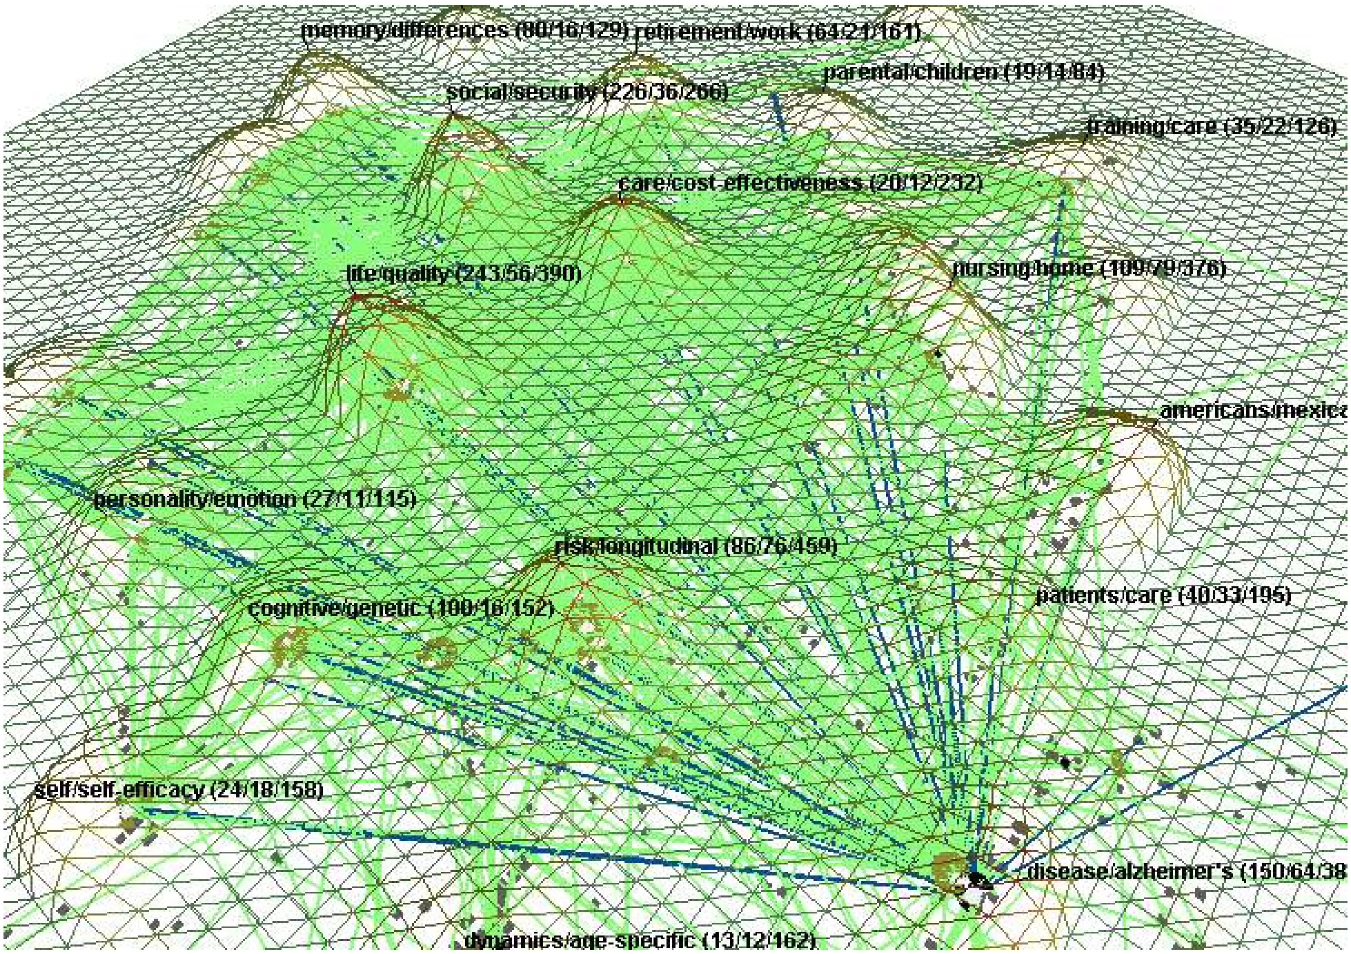

Figure 7. Author supplied linkage patterns (light gray lines) from grants to publications with links highlighted as dark lines for grant 01 P50 AG11715-01.

4. Research Challenges & Recommendations

Once again, the techniques presented in this paper can't replace human knowledge gathering and decision-making but they can support and complement it. Given complete and accurate data and an agreed upon set of partial indicators, domain visualizations can be objective and scalable using standardized processes. Resulting maps can be used by researchers, governmental workers, and industry researchers to accelerate their understanding of large data sets and to improve their decision-making.

Using grants and publication data from the BSR domain, we showed how major research areas, experts, institutions, grants, and publications in (aging) research and their interconnections can be determined, the evolution of areas can be visualized, as well as the influence of funding on the number and citation counts of publications. For the first time, grants and papers have been visualized together in one map.

The work presented here is certainly not complete. Many issues, especially regarding the completeness and accuracy of available data, exist. Among the major research challenges are:

- To utilize larger citation data sets (including citation links, citation counts, and abstracts as provided by ISI)

- To improve existing data mining techniques to incrementally process huge amounts of data
- To incorporate advanced text analysis (e.g., LSA, VantagePoint) to improve document clustering and labeling
- To make visualization easier to understand and use

Therefore, we conclude with a list of recommendations that aim to improve the quality and amount of data, analysis, and availability of resulting findings.

Recommendation 1 (to all funding agencies): Create a clean and maintainable database for grants and resulting publications as a basis for the application of bibliometric methods. Ask PI's to provide complete information on funding results – including co-authored papers, patents, and changes in public policies. Use a standard form entry to ensure a consistent data format. In the long run it might be advantageous to acquire, store, and utilize PI's resumes as a similarly consistent format as this would help, e.g., to disambiguate identical author names.

Recommendation 2 (to all funding agencies): In addition to measuring technical reports, lecture notes, grant proposals, publications in scholarly journals and conference proceedings, patents, R&D expenditures, equipment, and software, it would be advantageous to track and measure economic, environmental, social outcomes – contributing to the quality of life. Therefore, require PI's to provide a short "new result(s)/impact headline" that can be used to incorporate this data and improve the accuracy of mapping and labeling.

Recommendation 3 (Digital library and information visualization researchers, and grant agencies): Make data and results of science mapping analysis publicly available in a 'Scholarly Database' to

- Enable cross-disciplinary information access & transfer between different research areas
- Determine 'export factor' and 'import factor' for different research fields
- Reveal unproductive duplication, unrealized complementarity, gaps & opportunities, and overlapping topics
- Help to facilitate (cross-disciplinary) collaborations and to establish research priorities
- Support universal high information density and facilitate the creation of strong connections among major experts
- Provide opportunities for each scientist to think systematically about the state of their field
- Assess socio-economic impact of research
- Ideally, help achieve consensus on what should be funded

Acknowledgements

Symlin Lin and Raghu Mukkamalla, Information Science graduate students at Indiana University have been involved in the enormous amount of data parsing and cleaning necessary for the data analysis and visualizations shown in this paper. The ISI data set utilized in this study was extracted and made available by David Pendlebury and Henry Small, ISI. The SVDPACK by Michael Berry (1993) was used for computing the singular value decomposition of large sparse matrices. The authors gratefully acknowledge the support of Sandia National Laboratories, U.S. Department of Energy, under contract DE-AC04-94AL85000, and an NIH/NIA demonstration fund for Mapping Aging Research in 2001.

References

Berry, M. (1993). *SVDPACKC (Version 1.0) User's Guide, University of Tennessee Tech.* (CS-93-194).

Börner, K., Chen, C., & Boyack, K. W. (2003). Visualizing knowledge domains. *Annual Review of Information Science & Technology, 37, to appear*, 56 pages.

Bourke, P., & Butler, L. (1999). The efficacy of different modes of funding research: Perspectives from Australian data on the biological sciences. *Research Policy, 28*(5), 489-499.

Boyack, K. W., Wylie, B. N., & Davidson, G. S. (2002). Domain visualization using VxInsight for science and technology management. *Journal of the American Society for Information Science and Technology, 53*(9), 764-774.

Butler, L. (2001). Revisiting bibliometric issues using new empirical data. *Research Evaluation, 10*(1), 59-65.

Cronin, B., & Shaw, D. (1999). Citation, funding acknowledgement and author nationality relationships in four information science journals. *Journal of Documentation, 55*(4), 402-408.

Davidson, G. S., Hendrickson, B., Johnson, D. K., Meyers, C. E., & Wylie, B. N. (1998). Knowledge mining with VxInsight: Discovery through interaction. *Journal of Intelligent Information Systems, 11*(3), 259-285.

Davidson, G. S., Wylie, B. N., & Boyack, K. W. (2001). Cluster stability and the use of noise in interpretation of clustering. *Proc. IEEE Information Visualization 2001*, 23-30.

Deerwester, S., Dumais, S. T., Landauer, T. K., Furnas, G. W., & Harshman, R. A. (1990). Indexing by latent semantic analysis. *Journal of the American Society for Information Science, 41*(6), 391-407.

Halperin, M., & Chakrabarti, A. K. (1987). Firm and industry characteristics influencing publications of scientists in large American companies. *R & D Management, 17*(3), 167-173.

Irvine, J., & Martin, B. R. (1984). *Foresight in Science: Picking the Winners*. London & Dover: Frances Pinter Pub Ltd.

Jain, A., Garg, K. C., Sharma, P., & Kumar, S. (1998). Impact of SERC's funding on research in chemical sciences. *Scientometrics, 41*(3), 357-370.

King, J. (1987). A review of bibliometric and other science indicators and their role in research evaluation. *Journal of Information Science, 13*, 261-276.

King, J. (1988). The use of bibliometric techniques for institutional research evaluation: A study of avian virology research. *Scientometrics, 14*(3-4), 295-313.

Kleinberg, J. (1999). Authoritative sources in a hyperlinked environment. *Journal of the ACM, 46*(5), 604-632.

Kuhn, T. S. (1962). *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.

Landauer, T. K., Foltz, P. W., & Laham, D. (1998). Introduction to Latent Semantic Analysis. *Discourse Processes, 25*, 259-284.

Lewison, G. (1998). Gastroenterology research in the United Kingdom: Funding sources and impact. *Gut, 43*(2), 288-293.

Lewison, G., & Dawson, G. (1998). The effect of funding on the outputs of biomedical research. *Scientometrics, 41*(1-2), 17-27.

Lewison, G., Dawson, G., & Anderson, J. (1995). *The behaviour of biomedical scientific authors in acknowledging their funding sources.* Paper presented at the Fifth International Conference of the International Society for Scientometrics and Informetrics, Medford, NJ.

Lewison, G., & Devey, M. E. (1999). Bibliometric methods for the evaluation of arthritis research. *Rheumatology, 38*(1), 13-20.

Martin, B. R., & Irvine, J. (1983). Assessing basic research: Some partial indicators of scientific progress in radio astronomy. *Research Policy, 12*, 61-90.

McAllister, P. R., & Narin, F. (1983). Characterization of the research papers of US medical schools. *Journal of the American Society for Information Science, 34*(2), 123-131.

McAllister, P. R., & Wagner, D. A. (1981). Relationship between R&D expenditures and publication output for US colleges and universities. *Research in Higher Education, 15*(1), 3-29.

McCain, K. W. (1989). Descriptor and citation retrieval in the medical behavioral sciences: Retrieval overlaps and novelty distribution. *Journal of the American Society for Information Science, 40*(2), 110-114.

Narin, F., Olivastro, D., & Stevens, K. A. (1994). Bibliometrics: Theory, practice and problem. *Evaluation Review, 18*(1), 65-76.

Noyons, E. C. M., & van Raan, A. F. J. (2001). *MRC - Mapping the field of human ageing research*. Retrieved June 5, 2002, from the World Wide Web: http://www.cwts.nl/ed/mrc-aging/home.html

Pao, M. L., & Worthen, D. B. (1989). Retrieval effectiveness by semantic and citation searching. *Journal of the American Society for Information Science, 40*(4), 226-235.

Price, D. J. D. (1963). *Little science, big science*. New York: Columbia University Press.

Rinia, E. J., van Leeuwen, T. N., van Vuren, H. G., & van Raan, A. F. J. (1998). Comparative analysis of a set of bibliometric indicators and central peer review criteria : Evaluation of condensed matter physics in the Netherlands. *Research Policy, 27*(1), 95-107.

Shneiderman, B. (1996). *The eyes have it: a task by data type taxonomy for information visualizations.* Paper presented at the Symposium on Visual Languages, Boulder, CO.

van Raan, A. F. J. (1996). Advanced bibliometric methods as quantitative core of peer review based evaluation and foresight exercises. *Scientometrics, 36*(3), 397-420.

White, H. D., & Griffith, B. C. (1981). Author co-citation: A literature measure of intellectual structure. *Journal of the American Society for Information Science, 32*, 163-172.

White, H. D., & McCain, K. W. (1989). Bibliometrics. In M. E. Williams (Ed.), *Annual review on information science and technology. Volume 24* (pp. 119-186). Amsterdam, Netherlands: Elsevier Science Publishers.

White, H. D., & McCain, K. W. (1997). Visualization of literatures. *Annual Review of Information Science and Technology, 32*, 99-168.

Zhu, D., & Porter, A. L. (2002). Automated extraction and visualization of information for technological intelligence and forecasting. *Technological Forecasting and Social Change, 69*(5), 495-506.