Identifying Distinctive Competencies in Science

SciTech Strategies, Inc. Working Paper 2008-01

Richard Klavans (rklavans@mapofscience.com)
Kevin W. Boyack (kboyack@mapofscience.com)

Abstract

Science mapping is an alternative method for identifying areas of scientific leadership for a university. This method reveals the fine structure of science using bibliographic analysis and visualization algorithms. We compare the results from science mapping to those from traditional ranking approaches. Using the University of California, San Diego (UCSD) as a test case, we find that science mapping more accurately identifies the university's distinctive competencies in science.

Introduction

An accurate assessment of the scientific strengths of a university is critical for identifying areas of scientific leadership. Students want to know the best university in their specific area of interest when they are applying for graduate school. Faculty need to provide evidence of their scientific leadership when being considered for promotion and tenure. Administrators know that they can attract people (students and faculty) and funding (from alumni, foundations and government agencies) if there is well documented evidence that they are (or know exactly how to become) the leader for a specific area of science. Overall, accurately identifying areas where a university is, or can become, the scientific leader benefits a wide set of stakeholders.

University rankings are an increasingly popular method for comparing universities. Common sources of such rankings include U.S. News & World Report's annual entry on America's Best Colleges[1], the THES QS World University Rankings[2], and the Academic Ranking of World Universities from Shanghai's Jiao Tong University[3]. These three ranking systems are similar in that each takes a variety of inputs and generates a single measure of attractiveness for a university which forms the basis for the ultimate ranking. Values for each of the inputs are typically reported so that each customer of these rankings can modify their opinions based on the inputs they value most highly.

None of these university ranking systems measure how well the university performs in specific areas of science. To accomplish this, we must turn to methods for ranking nations typically used by government bodies, such as the National Science Foundation (NSF) in the United States, or the Organization for Economic Co-operation and Development (OECD) in Europe. These methods adopt a journal-based classification system (journals are assigned to categories) and then count articles and/or citations. The number of highly cited articles indicates historical

[1] http://colleges.usnews.rankingsandreviews.com/usnews/edu/college/rankings/rankindex_brief.php
[2] http://www.topuniversities.com/worlduniversityrankings/
[3] http://www.arwu.org/rank/2007/ranking2007.htm

leadership, while the number of current publications indicates current leadership. This method is widely accepted as fair and unbiased although it does have its critics and known methodological deficiencies.

In this report we investigate an alternative methodology for identifying scientific leadership. This method adopts a reference-based classification, in which 2.1 million highly cited references are clustered into 40,400 categories; a total of 5.6 million current articles are then assigned to this classification system. Each category is called a paradigm. Paradigms in which a university has a leadership position are identified. Distinctive competencies are defined as clusters of related paradigms where the university has a leadership position. This clustering is done on a university-by-university basis, so that the subsequent aggregations reflect the unique character of each university.

The balance of the report proceeds as follows. First, we use the traditional ranking approach to identify where the University of California, San Diego (UCSD) is the scientific leader. We then use an alternative science mapping methodology to identify where UCSD is the scientific leader. A comparison of these methods shows that science mapping generates a far more comprehensive list of areas in which UCSD is the scientific leader. We conclude with a discussion about the strengths and weaknesses of the new methodology.

Section 1: Traditional Methods for Measuring Scientific Strength

Traditional methods for measuring scientific strengths of a nation rely on counting and weighting publications according to a journal-based classification system. First, we present two journal-based classification systems: a highly aggregated system used by the NSF, and a more disaggregated system developed by the authors with support from UCSD. We then review some of the primary ways that publications are weighted and counted to determine strength and rank. We conclude with the identification of areas of science where UCSD is ranked #1 using these two journal-based classification systems.

Journal-Based Classification Systems: The most common type of classification system used by government bodies to rank nations is based on the assignment of journals to categories. For example, NSF, in its (NSB, 2008) Science and Engineering Indicators report[4], uses a classification that focuses on 4,906 journals. These journals *"give good coverage of a core set of internationally recognized peer-reviewed scientific journals, albeit with some English-speaking bias… Journals of regional or local importance may not be covered"*. This classification system has 13 fields and 127 subfields; the size distributions of categories at both levels are shown in Figure 1. The NSF subfields span two orders of magnitude in size, with the largest, *Biochemistry and molecular biology*, comprising 36,230 papers, and the smallest, *General engineering*, covering only 373 papers, using publication levels from 2005.

We initiated a research project in 2006 to generate an alternative journal classification system. This alternative was intended to overcome two shortcomings in the NSF journal classification system. The first shortcoming was coverage – the number and scope of journals in the classification system. The 4,906 journals selected by the NSF were from an initial set of 8,000

[4] http://nsf.gov/statistics/seind08/

journals that were being tracked by Thomson Scientific (TS, formerly the Institute for Scientific Information or ISI). Elsevier had recently introduced a competitive database that covered over 15,000 source titles. We were concerned that a classification system with only 4,906 journals and an English-speaking bias would not pick up the thousands of journals that were nationally recognized peer-reviewed journals or international journals with a more applied nature.

The second shortcoming was the level of aggregation. NSF used a highly aggregated classification system – only 127 subfields to characterize all of science. We believed that a more detailed classification system was required by students, faculty and university administrators. Those familiar with the structure of a typical university know that there are different levels of disaggregation (university, colleges, centers/institutes, schools and laboratories). The NSF classification system might be more appropriate if one were comparing universities or colleges. We were more interested in rankings that would apply to a more detailed level of analysis – that of the school or the large focused laboratory.

The journal-disciplinary classification system we developed (SciTech Strategies, STS) was designed to overcome these two shortcomings. We increased coverage to over 16,000 source titles (journals, proceedings, and book series) by combining the journal-journal relatedness statistics from the two largest citation database providers, Thomson Scientific and Scopus/Elsevier. We increased the level of detail (from 127 to 554 categories) by using the journal mapping techniques described in Boyack (2009). Size distributions for these two journal-classification systems are shown in Figure 1.

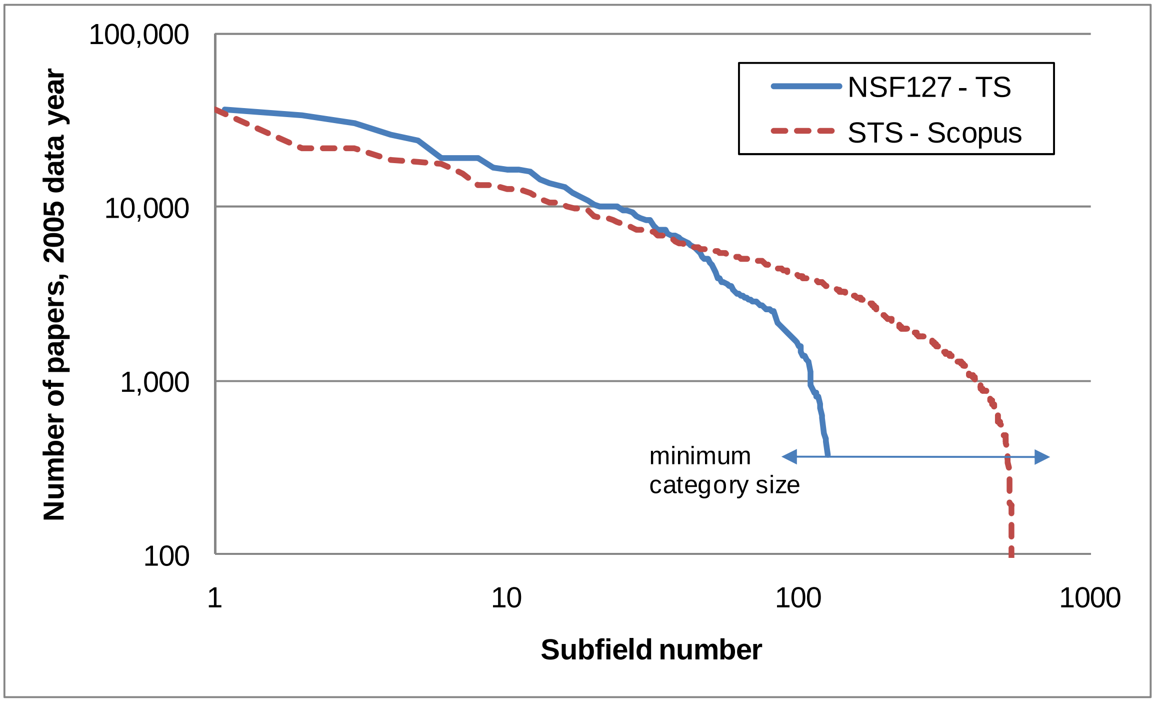

Figure 1: Sizes of subfields in two different journal-based disciplinary classification systems – the NSF classification system used in the 2008 Science and Engineering Indicators report, and the Scitech Strategies' (STS) classification system.

The largest STS category is nearly as large as the largest NSF category, with 36,218 papers (numbers based on a single year, 2005). The slopes of the two distributions (see Figure 1) are very similar for the first 30 categories, after which the slope of the NSF distribution falls off rapidly. The STS slope stays relatively flat for another 300 categories. Interestingly, 520 of the 554 STS categories are larger than the smallest of the NSF categories. The remaining 34 STS categories were dropped from this analysis; only 34 of the 187 journals in these 34 STS categories appear in the NSF journal classification system.

The structures of the two journal classification systems are quite different. The largest NSF category in terms of number of article is *Biochemistry and molecular biology*. All but two of the 155 journals in this NSF category are in the STS classification system, but they are dispersed into over 40 different STS categories. The largest concentrations of these journals are found in the *Protein science*, *Clinical cancer research*, and *Genomics* categories, which are the 11th-, 4th-, and 85th-ranked STS categories, respectively.

The largest STS category is *Data mining*. Only 6 out of the 99 source titles in this category are covered by NSF. Many of the source titles are refereed conference proceedings rather than journals, are not tracked by NSF, but are critical to the communication of research findings in computer science. Of these 6 journals found in the NSF system, 4 are in the *Computers* category, which is the 38th-ranked category in the NSF system. The STS classification system obviously has far greater coverage in the data mining area.

In summary, the NSF and STS classification systems represent two alternative journal classification systems. The NSF categories cover a highly select group of international journals that are classified into a relatively small (127) number of categories. Only the top international journals are included. Proceedings are excluded. The classification of these journals is vetted by experts. The STS journal classification system was our first attempt to overcome two weaknesses in the NSF approach: coverage and aggregation. The effect of these improvements, as shown by the ability to identify where a test-case university was ranked first, is discussed in a subsequent section.

Traditional Counting and Weighting Methods: There are two primary ways that scientific documents have been counted and weighted to determine strengths and ranks. The simplest approach is to count the documents by disciplinary category. This gives a raw score of the amount of current publication activity in each disciplinary category. This approach is used in the current study.

Another approach is to only look at the most highly cited articles (top 1% and top 5% are typical thresholds) and count by institution or country. This can be done within disciplinary categories (Hill, Rapoport, Lehming, & Bell, 2007), or for science as a whole (King, 2004). This approach has the advantage of measuring impact or strength, rather than just activity. However, it has the disadvantage that it is not timely. This indicator measures the strength of what was done several years ago (time lags of 10 years are common in these analyses), rather than current strength. Only documents that were published years ago can be identified using this method because the method requires allowing enough time since publication for papers to accrue large numbers of citations.

A modification of this approach is used in this study. We used a much lower citation threshold than most studies, which resulted in using the top 14.5% of highly cited papers. Further, we limited the citation window to citations from only one citing year, 2006, which gives a view of the cited articles that are of the most importance to current science, rather than those that have been the most important over a longer, and more dated, period of time. We also used a simple count of highly cited papers instead of a sum of the number of citations to these papers, to overcome problems of wide variations in disciplinary citation behavior.

Where is UCSD ranked #1? Table 1 lists the areas where UCSD is ranked first using either the NSF categories or the STS categories. Relative share of publications or relative share of highly cited references, are used to identify areas where UCSD is ranked first. Relative share is defined as the share of publications by researchers at a university divided by the share of publications of the largest competitive university. Using this formula, the leading university will have a ratio greater than one (their share divided by the share of the #2 university). All other universities will have a ratio less than one (their share divided by the share of the #1 university). This statistic is used to indicate how far ahead (or behind) a university is in publication activity.

Relative share was calculated based on current publications and highly cited references. Relative publication share (RPS) shows how far ahead the university is in current activity, but may be inflated because of publications in a regional journal. Relative reference share (RRS) tells us whether the university is the dominant source of publications that others are citing. This reflects a historical rather than current strength.

Both journal-based classification systems identify *Oceanography* as a major scientific strength of UCSD. The home of this activity is the Scripps Institution of Oceanography (SIO), a UCSD facility that is generally considered a world leader in this area of research. The NSF categories suggest that UCSD is 10% ahead of the #2 university, while the STS category suggests the lead is larger at 24%. Both systems suggest that UCSD is roughly tied for leadership when one focuses on highly cited references (.97 and 1.02).

Table 1. Categories where UCSD is ranked #1: Traditional Methods

NSF Categories	**RPS (current)**	**RRS (refs)**
Oceanography & Limnology	1.11	0.97
General engineering	1.13	0.78
STS Categories		
Oceanography	1.24	1.02

UCSD was also ranked first in the smallest NSF category, *General engineering*. We looked further into this category to determine why it did not appear as an area of strength using the STS classification system. A highly comparable STS category, *Oceanographic instrumentation*, was found with the same lead UCSD researcher in this area. However, UCSD did not appear as the top-ranked university in this STS category because it included a very institution-specific journal (Journal of Harbin Engineering University / Harbin Gongcheng Daxue Xuebao) which put the Harbin Engineering University in the top position. We then looked at the UCSD relative share of highly cited references in *Oceanographic instrumentation*, which was 0.91. UCSD was second behind the University of Washington. This was similar to the results using the NSF *General engineering* category in which UCSD was fourth with University of Washington first using relative reference share.

Overall, the NSF and STS journal classification system generated very similar results for UCSD. They both point to strengths at Scripps Institution of Oceanography, suggesting that SIO has two separate strengths, one in oceanography and the other in instrumentation. In both classification systems, the oceanography category is much larger than the instrumentation category, suggesting that oceanography plays the dominant role at SIO.

The effect of including regional journals in a classification system can also be seen. UCSD was not ranked first in the *Oceanographic instrumentation* category because of a regional journal from China. The importance of tracking the relative shares in both current papers and in highly cited papers can be seen. UCSD's leadership position in oceanography and oceanographic instrumentation is more weighted to current than past performance. The following section explores whether a reference-based classification system generates similar results.

Section 2: Science Mapping

Science mapping differs from the traditional approaches described above in that it uses a reference-based classification system instead of a journal-based classification system. First, we describe how the reference-based classification system was generated. We then describe how distinctive competencies for UCSD were determined. Distinctive competencies are defined as areas of science where, overall, an institution (such as UCSD) is ranked first, and the size of the area is greater than 1800 papers over a 5-year period. We use a threshold of 1800 papers because this is roughly the size of the smallest NSF category over 5 years (5x370 papers, see Figure 1).

Reference-based classification system: We use co-citation analysis to generate our classification system. Co-citation analysis clusters references, and then assigns papers to these clusters. There is a long tradition in using this method which has been very well described in the literature (Griffith, Small, Stonehill, & Dey, 1974; Small, 1999; Small, Sweeney, & Greenlee, 1985). Our classification system is based on a co-citation analysis of reference papers in the 2006 Scopus database. Over 2.1 million references were assigned to 40,400 clusters, which we call paradigms. The number of paradigms is emergent, rather than specified. A very low co-citation threshold was used, which results in high coverage of the available literature and a very low disciplinary bias (Klavans & Boyack, 2006). A total of 5.6 million indexed papers from the 2003-2007 publication years were then assigned to the paradigms on a fractional basis using their references. For example, if a paper had 6 references in paradigm A and 4 references in paradigm B, the paper would be fractionally assigned, 0.6 to A, and 0.4 to B. Only those indexed papers with references in the paradigms are assigned to paradigms. This is a majority of papers, but not all. Those not assigned are those with either no references, or only a few references that were not cited above our threshold. The method described above resulted in the identification of 40,400 categories or paradigms for science. Each paradigm is relatively small (on average, only 140 current papers and 54 reference papers). The distribution of current papers is used in the second stage to determine distinctive competencies for a university.

Distinctive Competencies: The first step in identifying UCSD's distinctive competencies was to focus on those paradigms where UCSD has strengths. We found that UCSD had a relative publication share (RPS) of greater than 0.54 in 626 paradigms. The 626 paradigms in which UCSD has a high relative share were then clustered using university-authored paper overlaps between paradigms. In many cases, current papers are fractionally split between two or more of the 626 paradigms. When summed over all UCSD-authored papers, these overlaps are substantial and form a relatedness matrix, which was then used as input for clustering. We used the VxOrd graph layout algorithm and an associated average linkage routine to assign paradigms to groups (Klavans & Boyack, 2006). Thirty-nine separate groups of paradigms were identified in this fashion.

Where is UCSD #1? Table 2 gives a list of the ten distinctive competencies that are above the minimum threshold of 1800 papers and where UCSD was ranked first. The remaining 29 groups represent smaller areas of strength or larger areas where UCSD was not ranked first.

The data on relative publication levels show that, in each case, UCSD is clearly the leader in each of these distinctive competencies. For example, UCSD publishes 51% more articles than its closest competitor in the largest distinctive competency. UCSD also published 97% more highly cited references that its closest competitor in the largest distinctive competency. As stated earlier, RPS is an indicator of current leadership; RRS is an indicator of past success. If RPS is less than RRS, it is usually due to the fact that researchers from other universities are quickly building upon the past discoveries of the leader, thus the current market is becoming crowded by more competitors. By contrast, if RPS is greater than RRS, this implies that the university is building on its own discoveries more quickly than are other universities. In nine of the ten cases, the leadership in current activity (RPS) is less than the leadership in past success (RRS). In only one case (*Oceanography; Instrumentation and Earth sciences*) is the pattern reversed.

Table 2. Categories where UCSD is ranked #1: Distinctive Competencies

Rank (DC#)	# of Papers, Paradigms & Disciplines	RPS (curr)	RRS (ref)	Inter-views	Description
1	23,102; 156; 36	1.51	1.97	6	-omic technology & molecular basis of disease
2	11,830; 76; 34	2.36	3.10	6	Cerebral vascular imaging & Neural functioning
3	8,652; 51; 17	1.75	1.67	2	Oceanography; Instrumentation & Earth sciences
4	8,289; 66; 19	2.31	2.68	1	Non-linear dynamics & Neuroscience
5	6,399; 25; 11	1.45	1.61	3	Optical bioengineering
6	4,135; 31; 17	1.11	1.44	4	Clinical aspects of biorhythm disorders and aging
7	3,733; 39; 14	1.47	2.32	3	Cellular control mechanisms
8	2,595; 15; 11	1.63	2.32	2	Genomic regulation
9	1,893; 12; 11	1.38	2.22	1	Behavior disorder (smoking, drinking, drugs)
10	1,848; 8; 4	1.41	1.86	1	Optics & Plasma flows

All ten distinctive competencies in Table 2 are highly multidisciplinary. There was a distinct pattern to this multidisciplinary research at UCSD. The first five and the tenth distinctive competencies were instances where measurement techniques are being applied to scientific areas.

The largest distinctive competency involves development of -omic technologies (techniques in genomics, proteomics, bioinformatics, metabolomics, etc.) that are then applied to the study of the molecular basis of disease. The second distinctive competency uses cerebral vascular imaging to better investigate neural functioning in multiple neurology subspecialties. The third distinctive competency links remote sensing technology to the investigation of large-scale oceanographic phenomena. The fourth uses modeling of non-linear dynamics from physics to investigate neuroscience. The fifth uses optical bioengineering for machine vision. The tenth uses optics and lasers to investigate plasma flows in plasma physics.

Several other competencies were not comprised of an intersection between technology and science, but were rather multidisciplinary combinations of areas within one or two broad areas of related science. This could be observed in the sixth, seventh, and eighth distinctive competencies. Clinical aspects of biorhythm disorders and aging were mostly within the broader areas of medical research and brain research. Cellular control mechanisms were mostly within the broader areas of infectious disease & medical research. Genomic regulation was mostly within the broad areas of biology and medical research.

A third multidisciplinary pattern is a link between the social sciences and mainstream science and engineering. This is especially important if scientific solutions are going to be applied to social problems that may involve changes in human behavior (such as the energy crisis and corresponding changes in energy use by society). The ninth distinctive competency is an example of this pattern. The focus of the research is on behavioral disorders (cessation of smoking, alcoholism and drug addiction). Over half of the research was in medical and brain research, the remainder in the social sciences and health services.

Methodological Difficulties: One of the most difficult tasks in this project was identifying, analyzing and labeling distinctive competencies. For example, a common problem in clustering is over-aggregation (putting groups together that might best be kept separate). When this problem appears it is usually in the largest and most diverse networks. To illustrate this problem, the following discussion focuses on the problem of identifying, analyzing and labeling the largest distinctive competency from Table 2.

Figure 2 shows a network diagram of the largest distinctive competency at UCSD. Each circle represents a paradigm. Paradigms that are next to each other are highly related. Paradigms that are related, but less so, are distant and have a line connecting them. There were 156 paradigms in this network, with over 23,000 current papers assigned to these paradigms.

Using different clustering thresholds, it would be possible to break the network in Figure 2 into ten to twenty smaller components. This could be a viable alternative for identifying distinctive competencies; it would generate a larger number of network components, each of which is more tightly linked together. However, this alternative approach would significantly reduce the number of components that are greater than the proposed threshold of 1800 papers, and would thus reduce the number of distinctive competencies. We investigated this approach using more stringent clustering thresholds, and found that the 626 paradigms would have formed 163 components; only 5 of these components would have exceeded the proposed threshold of 1800 papers. Further, this approach would not have picked up one of the key aspects of this network:

the relationship between -omic technologies and the molecular basis of disease. Following is a discussion of how this interpretation was generated.

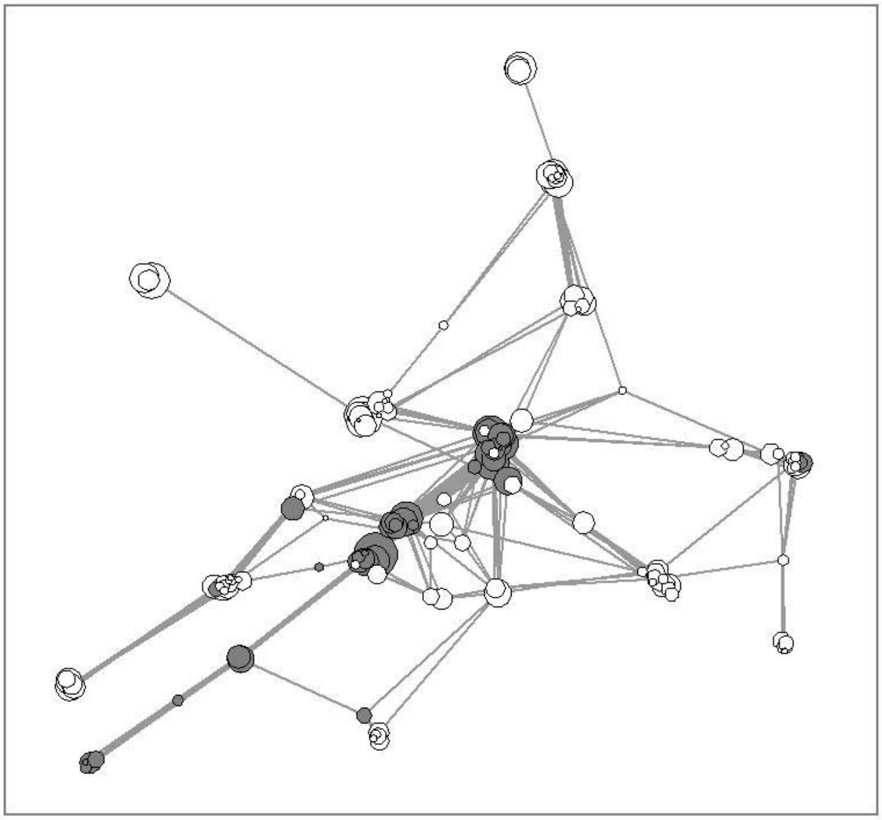

Figure 2: Network diagram of the largest distinctive competency at UCSD.

Our interpretation of the network in Figure 2 (i.e., the relationship between -omic technology and the molecular basis of disease) was based on initial interviews with six scientists at UCSD, follow-up interviews with scientists and administrators at UCSD, and an independent evaluation of the network by an outside expert. The shading in Figure 2 helps illustrate how this interpretation emerged. The shaded paradigms represent work in -omic technologies. The unshaded paradigms represent research related to cellular biology and disease. Three initial interviews were with scientists whose work is focused in the central shaded areas. Two of these three scientists were ranked #1 and #2 in terms of publications in the entire network. The paradigms they identified as their areas of expertise were associated with computational biology, structure modeling, metabolomics, bioinformatics and other –omic- related technologies. The other three interviews were selected to represent different (unshaded) parts of the network. These paradigms dealt with cell function, cell division, and cell membrane biology, as well as more applied research on cardiac, pulmonary and vascular diseases.

As listed in Table 1, we interviewed scientists in all of the distinctive competencies, with the number of interviews corresponding roughly with the size of the competency. Initially, we asked these experts to help generate interpretations of the distinctive competency in which they were active. This was a relatively successful approach when the network had less than 50 paradigms. But this approach did not work for the large network shown in Figure 2. We found from our interviews that areas of expertise are highly localized; each scientist would self-limit his/her area of expertise to a small area on the map. Scientists said they were not knowledgeable about paradigms that were two or three steps away from their area of expertise. For example, for the network shown in Figure 2, no single individual had a detailed knowledge of the entire network. However, each had a detailed knowledge of their own localized area and working knowledge of some of the immediately adjacent portions of the network. Overall, the six initial interviews were enough to establish the connections between the various portions of the full network. And yet it is likely that scientists active in paradigms that are more than two or three steps away from the shaded areas in Figure 2 would not perceive themselves as building on -omic technology.

After conducting interviews with 50 scientists, and testing different clustering thresholds as mentioned above, we decided to use the definition of distinctive competencies as originally specified. Although some of the resulting competencies seem to be larger than optimum, we find that the interpretations of these larger networks are compelling in terms of their multidisciplinary natures. We also understand that these network interpretations likely only refer to the majority (60% to 80%) of the constituent paradigms. In a sense this is analogous to geographical, social, or economic descriptions that are used to characterize a state (such as California) that consists of many cities and counties. The characterization of the state usually refers to the major population centers or broad geographic areas that are more similar. No characterization of a state holds for each city or county.

We continue to actively consider other solutions for the sizes and descriptions of distinctive competencies. One strategy may be to limit network size so that there is a maximum distance (such as three steps) from any paradigm to any other paradigm. This approach more closely aligns with the scope of expertise of individual scientists, and would therefore correspond more closely (in terms of scale) with expert judgment. This solution would generate many smaller distinctive competencies, few of which meet the critical threshold of 1800 papers. The assumption of an 1800-paper threshold may also need to be re-examined, and perhaps reduced.

Another strategy for identifying distinctive competencies would be to allow paradigms to occur in multiple distinctive competencies. This approach reflects an interesting pattern we found in the data. Instrumentation technology was being used in multiple distinctive competencies, but could only be assigned to one. The single assignment of nodes to clusters is a requirement in all clustering algorithms. New clustering algorithms are being developed to overcome this shortcoming, and allow certain nodes to occupy multiple positions. These new clustering algorithms may provide a more accurate method for identifying distinctive competencies from a set of paradigms.

Less costly approaches for analyzing and labeling the distinctive competencies are also being sought. The approach developed for this study (i.e., interviews with scientists, administrators and outside consultants) is very costly. It does generate a standard that can then be compared to less costly and more time-efficient approaches. These alternatives will be investigated in the future.

Section 3: Comparison & Discussion

The differences between the results from the traditional ranking approaches and science mapping are quite stark. The traditional methodology only identifies two areas where UCSD is ranked first, areas associated primarily with SIO, a known world leader in Oceanography. The science mapping technique identifies thirty areas where UCSD is ranked first (ten of which are large enough to be considered distinctive competencies), and several others where it is ranked second or third. First, we compare how these different methodologies characterize the strengths associated with SIO. We then explore possible reasons that the traditional methodology could not identify the additional nine areas in which UCSD has distinctive competencies.

UCSD/SIO is listed as the research leader in Oceanography-related areas using both the traditional and science mapping approaches. The two traditional methods list UCSD as having 309 articles in the two related NSF categories (*Oceanography* and *General engineering*) and 332 articles in the two STS categories (*Oceanography* and *Oceanographic instrumentation*). The science mapping approach shows UCSD with 369 in the *Oceanography, Instrumentation & Earth sciences* distinctive competency, which is 19% and 11% more publications than the two traditional methods, respectively. An analysis of the articles and the journals in which they appear helps illustrate the subtle differences between the different methodologies.

The traditional approach suggests that UCSD has strengths in two independent but related disciplines: *Oceanography* and *Instrumentation*. Both disciplines are directly associated with SIO. Oceanography is the much larger category of the two; if considered together, 92% and 70% of UCSD's publications are in the *Oceanography* category using the NSF and STS classification systems, respectively. Thus, oceanography is the dominant theme of the two.

By contrast, the science mapping approach automatically put the oceanography and related instrumentation research at UCSD together into one grouping, suggesting that this distinctive competency has a much broader multidisciplinary program than can be seen using any discipline-based approach. The majority of this research was also by researchers at SIO. Approximately 50% of UCSD's publications in this distinctive competency are in oceanography. 33% are in a variety of engineering and instrumentation-related journals (including geophysics, fluid mechanics, optics, acoustics, and analytical chemistry). The remaining 17% are in earth sciences and general interest journals (including articles in ecology, climate change, environmental science), and high impact journals such as *Science, Nature* and *PNAS*. This multidisciplinary mix of research was not picked up using the traditional approach. Rather, the traditional approach picks up the two largest components as separate strengths.

We now focus on the second issue – why doesn't the traditional approach identify the other nine discipline-sized areas in which UCSD is the research leader? One possible reason that the traditional approach fails to pick up these additional areas of research leadership is that they are all highly multidisciplinary. All ten distinctive competencies pull from subsections of many disciplines, as shown in Table 1. In fact, the largest distinctive competency, pictured in Figure 2, contains paradigms that belong to 36 different disciplines. The competency associated with SIO and discussed immediately above contains paradigms that belong to 17 different disciplines. Even the smallest distinctive competency contained paradigms from 4 different disciplines.

One of the main reasons that UCSD funded the development of the STS disciplinary classification system was the over-aggregation of the NSF and ISI classification systems. It was assumed that a more detailed classification system would be better at resolving the multidisciplinary nature of many activities. A further assumption was that one might be able to identify the component parts by breaking up each discipline into sub-disciplinary parts or smaller journal clusters. The STS disciplinary classification was successful in that is far more detailed that existing systems. Higher granularity also leads to higher accuracy.

But, in retrospect, we realize that there are theoretical limits to any journal classification system. Those familiar with journals know that almost all journals publish in multiple areas. These areas of research will be called specialties for the sake of argument. One can observe this fact by simply looking at the table of contents of a journal or the article classification system that the journal uses. Journals tend to publish research on dozens of research specialties. It logically follows that any journal cluster (i.e. discipline) will consist of multiple specialties, some of which overlap specialties from other journal clusters. We have concluded that one cannot identify individual specialties unless one abandons the use of journal clusters and works at a much finer level of analysis – the individual papers published and referenced in these journals. The results of this study support these conclusions.

Section 4: Summary

Both the traditional journal-based classification method and science mapping are proposed as objective methods for identifying where a university is ranked first in an area of science. Both rely on well-established methods for analyzing the scientific literature. Claims of research leadership can be substantiated with corresponding lists of research publications and expert interviews.

We see two advantages for using the traditional method for identifying where a university is the research leader. First, the methodology is easily implemented by users at a relatively low cost. This is because the journal classification systems are easy to obtain and literature searches can be done at low cost using web-based data services. Second, this is a tried and true methodology with well-known pitfalls and shortcomings. The fact that the journal categories are overly aggregated simply means that a very high ranking (of #2 or #3) implies that one is likely ranked first in a subset of that discipline. High rankings are, by themselves, useful assessments of the relative strengths of a university.

The traditional methods did not, however, accurately identify the majority of areas in which UCSD is a research leader. Science mapping identified thirty coherent areas, rather than just two, where UCSD was ranked first. Ten of these areas have a literature base that is large enough to be called discipline-sized, and we thus label them as distinctive competencies. We argue that traditional methods that rely on journal-based classification systems cannot identify research leadership when the area is multidisciplinary. We suggest that multidisciplinary research is not idiosyncratic to UCSD. Rather, we suspect that most distinctive competencies in most universities will prove to be multidisciplinary in nature.

Science mapping at the scale of millions of scientific articles is a relatively new technique. While it builds on a 30-year history of reference-based classification systems, many of the algorithms were recently created so that one could identify the distinctive competencies of an institution, region or nation. The methodology is still in its infancy. While its advantages are clear, it is also clear that this new technology can be improved. As mentioned in a previous section, some of the competencies may be over-aggregated, and although we have chosen to stick with our current clustering thresholds, we continue to experiment in that area. In addition, better techniques are needed to label distinctive competencies. The expert researchers we interviewed were able to readily label the smaller competencies, but we have not yet identified an automated text-based approach that can generate labels that are similar to those created by experts.

Despite these shortcomings, we suggest that the benefits of using detailed science mapping to identify research leadership are compelling. By increasing the number of science building blocks by two orders of magnitude (from 554 disciplines to 40,400 paradigms), we see a finer structure of research that captures important aspects of how researchers are organized. Science mapping reveals cases where measurement techniques are linked with new areas of science. It can reveal cases where research on important social issues (such as global warming and environmental degradation) can be strengthened. Most importantly, it can reveal the specific areas in which a university has achieved, or is planning to achieve, research leadership.

In summary, we suggest that science mapping can be a significant benefit to multiple users. Students can use this to select the best programs in their chosen topics of interest. Faculty who participate in multidisciplinary research can provide evidence of their contributions for the promotion and tenure process. Finally, university administrators concerned about attracting students, faculty, and research dollars, can have a more objective view of what is required to be the research leader in an area of science.

Acknowledgments: We are grateful to the Office of Research Affairs at the University of California, San Diego for funding this research and to Vice Chancellor for Research Arthur B. Ellis for helpful discussions. We also thank the 50 UCSD researchers whom we interviewed for their assistance in delineating paradigms in which they hold expertise, and their help in interpreting and labeling networks of paradigms leading to validation of the UCSD distinctive competency maps.

References

Boyack, K. W. (2009). Using detailed maps of science to identify potential collaborations. *Scientometrics, 79*(1), 27-44.

Griffith, B. C., Small, H. G., Stonehill, J. A., & Dey, S. (1974). Structure of scientific literatures. 2. Toward a macrostructure and microstructure for science. *Science Studies, 4*(4), 339-365.

Hill, D., Rapoport, A. I., Lehming, R. F., & Bell, R. K. (2007). *Changing U.S. output of scientific articles: 1988-2003* (No. NSF 07-320). Arlington, VA: National Science Foundation, Division of Science Resource Statistics.

King, D. A. (2004). The scientific impact of nations. *Nature, 430*, 311.

Klavans, R., & Boyack, K. W. (2006). Quantitative evaluation of large maps of science. *Scientometrics, 68*(3), 475-499.

National_Science_Board. (2008). *Science and Engineering Indicators 2008*. Arlington, VA: National Science Foundation (volume 1, NSB 08-01; volume 2, NSB 08-01A).

Small, H. (1999). Visualizing science by citation mapping. *Journal of the American Society for Information Science, 50*(9), 799-813.

Small, H., Sweeney, E., & Greenlee, E. (1985). Clustering the Science Citation Index using co-citations. II. Mapping science. *Scientometrics, 8*(5-6), 321-340.