SciTech Strategies, Inc.

BETTER MAPS • BETTER DECISIONS

Evolution of Distinctive Competencies:

An Analysis for the University of Utah

Draft Final Report

June 2009

Prepared for:

Tom Parks
Vice-President of Research
University of Utah

Prepared by:

Kevin W. Boyack
Richard Klavans
SciTech Strategies, Inc.

Table of Contents

Introduction

Research strengths are the single most important factor in charting the growth and health of the university. A University relies on its research strengths to attract the best students. Universities recruit and retain star faculty in order to build and maintain their strengths. Universities pursue research grants in order to support future research strengths. Universities establish multidisciplinary research centers in order to create strengths that are focused on specific topic areas.

The consequence of these concerted efforts is a research portfolio that defies traditional methods of evaluation. Simply relying on subjective evaluations of strengths by outside experts is not sufficient. Funding sources (whether they are students, grantors, government agencies or corporations) are relying increasingly on the use of objective indicators to determine who is best in a particular area of research. Simply relying on traditional publication counts by disciplinary area also doesn't work well. Research strengths are highly specific and focus on a topic; disciplines are extremely aggregated and have little to do with the way that research strengths evolve.

In this report, we describe how a new measurement technique can overcome some of the shortcomings of expert-based and discipline-based approaches to the evaluation of a university's research portfolio. This new technique was initially developed in 2007 for Art Ellis (Vice Chancellor of Research for the University of California at San Diego). Tom Parks (Vice President of Research) and other representatives from the University of Utah learned about this new technique when they participated in a 'Best Practice' conference conducted by the University Leadership Council of the Education Advisory Board in 2008. The technique was also being incorporated into a web-based evaluation tool with a scheduled market introduction by Elsevier in mid-2009.

We gratefully acknowledge the role that Tom Parks, and his colleagues at the University of Utah, have played in improving this measurement technique. Their insights have played a central role in our ability to understand how research strengths evolve. An objective and accurate method for evaluating a university's research strengths is one step closer to reality.

Initial Project Plan

The University of Utah was interested in both the existing technique for identifying research strengths and the future techniques that were under development at SciTech Strategies, Inc. This project therefore had two intended outcomes.

One outcome was to identify the University's research strengths using a beta version of Spotlight, a research evaluation tool that was scheduled for commercial introduction in June, 2009. Additional data on these research strengths would be made available to the University. Possible shortcomings of Spotlight were addressed as they became apparent.

The second outcome was to examine how these distinctive competences evolved. As stated in Appendix A of the proposal, "Utah will be the first site where we will test our algorithms for identifying patterns of evolution. The testing procedure will primarily be based on interviews with faculty and administration at the university. The goal will be to describe, as accurately as possible, the evolution of the distinctive competencies at the University of Utah over a five year period".

This report proceeds as follows. The following section describes the methodology underlying Spotlight and provides a high-level review of the findings. Three shortcomings were identified: the descriptions of distinctive competencies were difficult to understand; the results may have been inaccurate because of recent faculty changes, and the results could not be used by licensing to identify corporate targets. All of these shortcomings are addressed in this report.

The subsequent section focuses on how we analyzed evolution. This required two sets of interviews. The first set was aimed at the 'great person' argument (distinctive competencies were really projections of the research strategy of the lead researchers). The strategic directions pursued by star researchers have obvious implications for managing the evolution of a distinctive competency. The second set of interviews examined the threads (evolutionary patterns of related paradigms) associated with a distinctive competency. Experts were asked to validate the core threads (those directly resulting in the distinctive competency) and identify which related threads were part of environmental scanning.

The final section highlights major findings from these two sections.

Distinctive Competencies in 2007

Methodology

The steps used to identify the distinctive competencies of the University of Utah are outlined in Figure 1.

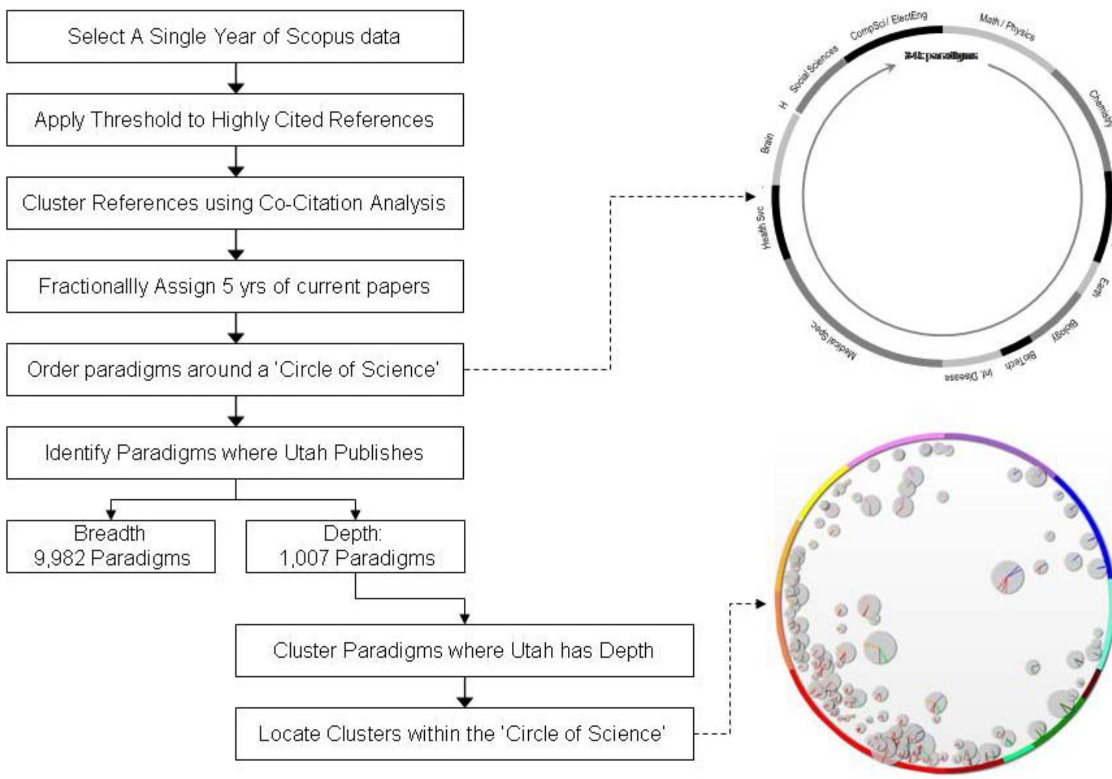

Figure 1: Steps in generating a model of science and identifying distinctive competencies.

<u>2007 Scopus Database:</u> We started with the most current yearly database. This project was initiated in January, 2009, but we could not use the 2008 database (due to publication and indexing lags, indexing for the 2008 publication year would not be completed by Scopus until May 2009). We therefore used the 2007 Scopus database, which had 1.376 million articles (with references) that were published in over 15,000 journals and conference proceedings.

<u>2.1 Million Highly Cited References:</u> The method we used to identify document clusters is called co-citation analysis. The method starts by identifying a set of highly cited references. Standard policy by most researchers that do co-citation analysis is to select the top 1% by field (Chen, 2006; Franklin & Johnston, 1988; Small, 2006). We use a significantly lower threshold so that far more references can be identified. This policy, which is unique to SciTech Strategies, Inc., results in the earlier detection of emerging research fronts and a more accurate description of

the structure of science. We are selecting all references with at least 5 citations from the citing year (the threshold is lower if the reference was published in the past 3 years), and correspondingly use 13% of the available references to specify a structural model of science.

Initial Paradigm Solution: Co-citation analysis involves the clustering of these highly cited references based on co-occurrences (whether citations co-occur in the bibliographies of more recent articles). A modified cosine measure is used to characterize reference:reference relatedness. Ten runs of our clustering algorithm create slightly different cluster solutions (the differences are based on the use of random starting points). We then apply a consensus criteria (six out of ten runs agree that the same references belong in the same cluster) to create an initial paradigm solution of approximately 200,000 reference clusters.

Final Paradigm Solution: Five years of current publications (2003-2007) are then assigned to these 200,000 reference clusters by matching the references in the current papers. Attention is paid to create unambiguous assignments (at least two references must be in the same reference cluster for there to be an unambiguous assignment). There are far more unambiguous assignments of current papers to reference clusters because of our use of a significantly larger set of references.

The assignment of a current paper to more than one reference cluster creates two critical statistics: an indication of the size of the cluster and a cluster:cluster relatedness measure. We then identify all clusters below a minimum size (25 papers), and combine these small paradigms with their most related paradigm. Over 98% of the clusters that were below this minimum size had sufficient relatedness to other clusters to be combined with those clusters to create clusters with a minimum size.

After this aggregation step, these clusters are called paradigms. The final solution consisted of 84,202 paradigms in 2007. The largest paradigm had 847 current papers (an average of 170 papers per year). This is well within the expected range of cluster solutions. (The general rule of thumb has been that any cluster with more than a few hundred current papers in a year is suspect because scientific communities rarely get this large before they break apart into subgroups to focus on different aspects of the research problem.) There were only 26 paradigms that had an average size of more than 100 papers per year.

Circle of Science: Klavans & Boyack recently created a consensus map of science based on analyses of 20 maps of science that had been created over the past 70 years. These maps used a variety of methods, measures and data sources (Klavans & Boyack, 2009). We found that the underlying structure of science forms a circle. Starting (arbitrarily) at the top, the map goes from math to physics, physical chemistry, chemistry and biochemistry, with engineering, geoscience and biology forming a parallel path before it recombines with biochemistry. One can proceed around the circle from biochemistry to infectious diseases, medical specialties, neurology and psychology, with health services forming a parallel path. From psychology one proceeds to the social science, computer science and returns back to mathematics. The humanities branch off of the social sciences.

This procession of fields was used to create an ordering of 554 disciplines around a circle of science. The 84,202 paradigms were assigned to disciplines (based on journal assignments) and consequently located on the edge of this circle of science.

The reason for placing these paradigms along the edge of the circle of science is to illustrate whether a particular research program is multidisciplinary. Research programs (even the research programs that individual researchers pursue) involve dozens if not hundreds of paradigms. Any research program that focuses on the research in a related set of disciplines (i.e. all within physics) will have a location on the circle of science that is close to the edge. Any multidisciplinary program will have an interior position. We will return to this visualization approach when we locate the distinctive competencies of the University of Utah.

<u>Breadth vs. Depth:</u> The next two steps involve splitting Utah's research efforts into two groups. The assignment depends on relative publication share (RPS). RPS is defined as your publication level divided by the publication level of the largest competitive institution. The publication leader will always have an RPS greater than 1.0 (the competitor, by definition, is smaller). The followers will have an RPS less than 1.0 (the follower's publication level divided by the leader's publication level). The RPS threshold for the University of Utah was 0.44. This threshold is based on an equation we have developed for all major universities, and varies from approximately 0.2 (for the smallest university) to 0.6 (for the largest university).

There were 9,982 paradigms where the university had breadth (the university publications were assigned to these paradigms but the RPS < 0.44). One can view these areas of breadth as areas that faculty are scanning (as part of their day-to-day research activities) and collaborating (being included as a secondary author on an article).

There were 543 paradigms where the university had depth (an RPS ≥ 0.44). These are paradigms where the researchers at the university play a significant (and in many cases dominant) role in creating scientific progress. These paradigms, if related, become the basis for the university's distinctive competencies.

<u>Clustering Paradigms with Depth:</u> The 1,007 paradigms where the university has depth are clustered based on the idiosyncratic publication patterns within the university. Many of the university's current papers were assigned to more than one paradigm. These multi-paradigm papers create a unique database of paradigm-paradigm relatedness. No two universities link

the same paradigms in the same way. These university-specific patterns are used to cluster the 'depth' paradigms. 543 (54%) of these paradigms became members of (149) clusters. The remaining 464 paradigms were isolated areas of strength. 70.4% of the university's publications were in areas of breadth; 10.3% of the publications were in isolated strengths and 19.2% were in paradigm clusters.

One interesting consequence of this clustering is that the relative publication share (RPS) for the related set of paradigms rises as paradigms are connected. It's rare that a researcher, who publishes multi-paradigm papers, chooses paradigms with the same set of competitive institutions. It's quite possible to have an RPS of 0.6 in four equally sized paradigms, but because the research leaders in each paradigm are completely different, the RPS of the set of four paradigms could theoretically reach 2.4! In reality, the RPS values will easily exceed 1.0 but rarely exceed 3.0.

We apply three criteria to determine if a cluster of (depth) paradigms are distinctive competencies or emerging competencies. The most commonly used measure is RPS (relative publication share). We also look at RRS or relative reference share (whether the university has the most number of highly cited references in the paradigms). The third indicator, SOA (state of the art), is used to determine if the current papers build on the more recently published references.

<u>Location of Distinctive Competencies and Emerging Competencies:</u> Each cluster of (depth) paradigms, now called distinctive competencies if they meet minimal criteria for size and RPS, RRS or SOA, is located on the circle of science. Information about these distinctive competencies is placed in a web-enabled program (called Spotlight, see (Klavans & Boyack, 2010) for technical details). Emerging competencies (the clusters of depth paradigms that do not meet minimal criteria) are also placed in the program. The user can then explore this database to better understand their portfolio of research strengths.

<u>Examples of Output:</u> Examples of the tabular output in Spotlight are shown below. Figure 2 shows the output of the first 20 DCs (or ECs) along with their size, growth, and leadership statistics. Figure 3 shows the type of detailed summary that is available for each distinctive competency. Additional detail on each competency is available under the "Top Authors", "Top Institutions", "Circle", and "Cluster Map" tabs in this Spotlight view.

We have used Distinctive Competency #6 as an illustration of the capabilities of Spotlight. This distinctive competency is the largest circle shown in Figure 1. Using the face of a clock, DC06 is located at about 2:30 and has a more interior position. The interior position is due to its multidisciplinary nature. Referring back to Figure 1, the paradigms that comprise this distinctive competency are from chemistry (the blue lines within the circle) and medical specialties (the red lines within the circle).

Competencies	Market size	Growth	Article Share	Growth	Rank	RAS	RL	SotA
1 (DC)	1541.5	6.34% ▲	101.6 (6.59%)	1.00% ▲	1	3.15	1.88	1.32
2 (EC)	1079.9	2.77% ▲	23.9 (2.22%)	0.65% ▲	5	0.78	0.63	1.18
3 (DC)	1583.2	10.38% ▲	62.4 (3.94%)	-0.80% ▼	1	1.70	2.42	0.58
4 (DC)	1419.5	-2.10% ▼	93.6 (6.60%)	0.26% ▲	1	3.11	2.59	1.26
5 (DC)	1493.6	8.05% ▲	34.8 (2.33%)	0.37% ▲	1	1.14	0.28	0.77
6 (DC)	1194.7	1.46% ▲	92.3 (7.73%)	1.52% ▲	1	3.94	4.23	0.98
7 (DC)	1088.7	-0.59% ▼	60.0 (5.51%)	1.33% ▲	1	3.56	3.76	3.05
8 (DC)	1049.6	-1.57% ▼	38.9 (3.70%)	-0.26% ▼	1	1.44	1.39	0.90
9 (DC)	799.3	-7.16% ▼	15.7 (1.96%)	-0.06% ▼	3	0.60	1.08	1.54
10 (DC)	1160.2	1.19% ▲	32.4 (2.79%)	-0.41% ▼	1	1.24	0.53	-0.33
11 (DC)	781.2	4.32% ▲	21.8 (2.79%)	0.23% ▲	1	1.17	0.77	0.76
12 (DC)	1024.2	-1.09% ▼	24.8 (2.42%)	-0.19% ▼	2	0.91	0.79	0.55
13 (DC)	929.4	7.72% ▲	25.2 (2.71%)	-0.48% ▼	1	1.43	0.35	-0.27
14 (DC)	889.7	-4.26% ▼	22.5 (2.53%)	0.38% ▲	1	1.02	0.54	1.15
15 (DC)	638.5	10.15% ▲	12.2 (1.91%)	-0.01% ▼	1	1.23	0.00	0.21
16 (DC)	801.0	15.06% ▲	26.8 (3.34%)	0.32% ▲	2	0.83	0.39	0.85
17 (DC)	891.7	0.09% ▲	28.8 (3.23%)	-0.09% ▼	1	1.65	2.24	0.61
18 (DC)	896.1	6.70% ▲	44.9 (5.01%)	-0.07% ▼	1	1.92	1.39	0.63
19 (DC)	1055.7	8.08% ▲	34.8 (3.29%)	-0.48% ▼	1	1.08	0.39	-0.52
20 (EC)	724.2	2.55% ▲	9.9 (1.37%)	0.13% ▲	14	0.39	0.33	0.06

Figure 2: Spotlight output showing aggregate quantities for the top 20 Utah competencies.

Distinctive Competency #6
Main keywords: drug delivery; gene therapy; drug release

Change region Global
How to read this p

General | Top Authors | Top Institutions | Circle | Cluster Map

	Market Size (Global)	University of Utah		Relative Article Share	Reference Leadership	State-of-the-Art
	1194.7	92.3	7.7%	3.94	4.2	1.0
Growth	1.46% ▲	3.91 ▲	1.52% ▲	▲		

University of Utah articles

Fractionalized	92.3	(0.7% of 14152 articles total for this institution past 5 years)
Articles	152	View 152 articles in Scopus
Rank past 5 years	1	
Rank past 2 years	1	
Citation count	774.2	

Top authors from University of Utah

Name	Fractions
Kopecek J.	16.9
Lu Z-R.	15.6
Kim S-W.	14.8
Kopeckova P.	11.9
Parker D.L.	11.8

Go to authors

Top disciplines

Name	Fractions (% of total)	
Pharmaceutical Research	763.4	(63.9%)
Clinical Cancer Research	258.1	(21.6%)
Bone & Osteoporosis	82.4	(6.9%)
Macromolecules & Polymers	53.8	(4.5%)
Leukemia	35.8	(3.0%)

View all

Figure 3: Spotlight output showing detailed summary information for Utah DC6.

Discussion

The following shortcomings in the current methodology were identified as Utah began using Spotlight:

- the descriptions of distinctive competencies were difficult to understand
- the results may have been inaccurate because of recent faculty changes
- the results could not be used to identify corporate targets associated with the DCs

<u>Describing Distinctive Competencies:</u> Spotlight uses state of the art text summarization algorithms to describe a distinctive competency. Unfortunately, the underlying phrases, drawn from keywords and titles, have a very high level of scientific literacy associated with them. One could say that the algorithms are too successful – they describe the area of research using terms that only the scientists in that specific area understand. The researcher in that area recognizes exactly what the research is about. There are no algorithms available that can 'translate' these highly specific phrases into phases that would be more meaningful to research administrators and funding sources (i.e. provosts, students, funding agencies or corporations).

We concluded that the only way to deal with this problem was to ask the top researchers in the distinctive competency to come up with concise, accurate, and differentiating labels for the competencies. This could be implemented very easily. First, we identified the top (one to four) researchers for each distinctive competency. We then provided them with a summarization of the research (using the text approaches mentioned above), and asked them to look at the whole and give us a label for the area that would be meaningful to potential students, funding agencies and corporations. This could all be done via email.

This process was also very successful at validating the distinctive competencies. Researchers were quite willing to translate their passion (which was reflected in the specialized language) into common language. There was one case where the lead researcher for the distinctive competency said- quite clearly- that this was no longer the main focus of his lab. They were still doing that work because companies were still paying them to do it, and thus the area was still a competency, but the focus of the lab had changed with new funding sources. This validation step can also help identify distinctive competencies that were about to be (or had already become) obsolete because of publication lags.

<u>Recent faculty changes</u>: The strengths of a university depend on the researchers who are at the university. The addresses used in the analysis, however, reflect the institution where a researcher did the work that was published, and do not necessarily reflect their current affiliations. If researchers have a change of address, literature databases are not updated to reflect those changes. The adding/dropping of key researchers was a possible source of inaccuracy in the reported results.

We tested this possibility by getting add/drop lists. The administration at Utah provided the names of the top researchers who had been hired in the past few years. They also reviewed the list of the top 100 researchers based on publication counts and identified those individuals who had recently (or were about to) leave the university. The entire analysis of distinctive competencies was then redone with these changes.

The results of these analyses suggested that, in this case, there was little change in the overall results. There was no significant change in the structure: the same distinctive competencies emerged. A few very small distinctive competencies were dropped. Overall, research superstars that left the university did not have a devastating effect on the distinctive competencies of the university. Most were part of research teams, and as long as other members of the team remained, the competency was still present at the university. The research superstars that had recently joined the university did not create new competencies. It appears that the hiring was more focused on significantly strengthening existing competencies.

A separate and detailed comparison of the results of the add/drop calculation with the original calculation are provided in Appendix A.

Licensing: Tom Parks asked us whether the distinctive competencies could be used to help identify potential corporate licensing partners. This is a new application area for Spotlight that was worth exploring. In order to accomplish this, we coded in any organization (whether university, corporation, hospital or government lab) that published at least 250 papers in a 5 year time horizon (or 50 papers per year). Corporate researchers, active in one of Utah's distinctive competencies, can now be contacted to determine if they are interested in funding university research or licensing patents or intellectual property that has resulted from the university's distinctive competency. The detailed data tables included as Appendix B to this report (Excel format) include industrial institutions and authors. These are highlighted in yellow for easy identification.

It's important to emphasize that this only focuses on finding corporate licensing partners for areas where the university has depth. However, from the corporate perspective, it makes sense to place the #1 university in a specific research area on the top of their list of potential research partners.

In summary, the method for identifying the distinctive competencies of Utah was very accurate. Possible sources of inaccuracy were identified and assessed. The lag time (between research and publishing) was rarely a problem because major research programs have momentum, and one can easily find out from the research leaders if there has been recent changes in the level or direction of research. The possible inaccuracy because of researcher mobility was assessed. Recent losses of research superstars had little effect because many of the team members remained at the university. Recent hires of research superstars augmented existed competencies.

Evolution of Distinctive Competencies

We conducted two sets of interviews at the University of Utah. The first set was aimed at gaining insights into the relationship between research superstars, their research strategies and the distinctive competencies they participate in. The second set of interviews was designed to look at how distinctive competencies evolve.

Superstars, Research Strategies, and Distinctive Competencies

What is the relationship between the research strategy of a superstar and the trajectory of research associated with a distinctive competency? This is an extremely difficult question to answer empirically because, as far as we know, there hasn't been an effective way to objectively define either phenomenon. The following methodology describes a new method for identifying the research strategy of a superstar. We also looked at the tendency of star researchers to publish in areas where they have already had success (previous publications are part of the reference structure of the paradigm). We then discuss the implications of these findings.

Methodology

Seven of the most highly published researchers at the University of Utah were asked to participate in a card-sorting exercise. The sample of researchers was drawn from the top 20 researchers, and interviews were arranged based on the availability of the researcher during the time period that we were conducting these interviews.

These seven researchers can be considered superstars. The average performance for the 8,853 authors (professors, staff, and students) at the University of Utah was 3.7 publications and 1.4 references in 2007. (This does not count the researchers at the university that published in previous years but didn't publish in 2007.) In comparison, these 7 researchers had an average of 64.6 publications in 2007 and 48.3 references in the reference structure for the 2007 model.

We identified the 30 paradigms where each researcher published the most, placed descriptions of this research on cards, and asked the researchers to use these cards to describe their research strategy. They were asked to describe their research strategy by first putting the cards into groups (areas of related research that they were pursuing). Any cards that didn't represent an area of research they were pursuing were placed in an 'unrelated' pile. This highly structured task was relatively simple for researchers to accomplish (it took approximately 20 minutes and was reported to be 'fun' because the cards represented what they were most passionate about).

The next step was more unstructured. We asked the researchers to describe the relationships between these research areas. The responses were quite varied, and provided some unique insights into how the researchers viewed their work. One researcher had a bulls-eye view (there was a core research area, some related areas, and then the unrelated pile mentioned above). Another researcher created three piles – one for the lab that they were most passionate about, another for the lab where they do a lot of work but don't think there will be much money to support the research in the future, and finally the ever present 'unrelated' pile. More articulated structures were also provided – making distinctions based on theory, applications, or instrumentation.

Our first, and relatively obvious, hypothesis was that there would be more publications by the researcher (or by the university) in the grouped cards than the ungrouped cards. The grouped cards represent research that the researchers are interested in. The unrelated pile represents publications and references that may have occurred because of past collaborations.

Our second hypothesis was that these star researchers would publish much more in paradigms associated with distinctive competencies. If star researchers are the source of the university's distinctive competencies, then there should be more publications in the distinctive competencies. The expected publication level was 19% (distinctive competencies only account for 19% of the overall publications in the university). We expected that the star researchers would have a much higher fraction of their publications in distinctive competencies than 19%.

The third, and perhaps most interesting, hypothesis had to do with the relationship between current publications and reference papers in a paradigm. We assume that researchers publish articles with the intent that many of these articles will become part of the reference structure of the paradigm. Consider, for example, the average paradigm where there are 14 current articles each year that are citing an intellectual base of 25 references. Assume that the window for citation and for replacing the reference base of a paradigm is about 7 years. This means that the odds of having a current paper become part of the intellectual base is about 25:98 (the 98 papers published over the seven year period will 'replace' the 25 references that define the intellectual base of the paradigm. This ratio will, of course, depend on the window for citation and the rate of reference replacement. But assuming similar citation windows and reference replacement rates, a simple regression of reference papers (as a function of current papers) will indicate the ability of a researcher to get their papers into the intellectual base of a paradigm. The expected ratio for all Utah researchers is 0.756 reference papers per current publication. We expect this ratio to be much higher in the grouped paradigms (versus the ungrouped paradigms). We also expect this ratio to be much higher for the paradigms where the university had depth (a higher relative publication share) than for paradigms where the university had breadth.

Following are the results for each hypothesis.

Hypothesis 1: The top researchers will publish more in the grouped paradigms than in the paradigms that were judged to be unrelated to their research interest.

About 90% of the current papers were in the 221 grouped cards, while the remaining 10% were in the 111 cards that were judged to be 'unrelated' to their research interests. This is really not a surprising finding – it simply means that the assignment of the author's publications to the paradigms is relatively accurate. We selected the 30 paradigms with the most publications by the researcher, and were not surprised that, when they partitioned these paradigms, the ones where they publish more were considered part of their research strategy.

Hypothesis 2: The top researchers will publish more in the paradigms that are part of a distinctive competency than the paradigms that are not.

The seven researchers we interviewed published 61% of their research in the 40 paradigms that were associated with distinctive competencies. It was relatively rare to find a case where the researcher said the paradigm was in a distinctive competency and also said that the paradigm was 'unrelated' to their research interests (this accounted for only 2.3% of the publications in distinctive competencies).

We did find, however, that not every top researcher published in distinctive competencies. One top researcher, in particular, did not have any publications in paradigms that had depth (a high relative publication share). This researcher appears to be a star scanner – publishing in a very wide set of research areas. There were, however, very few cases of the top researchers emphasizing breadth over depth. The tendency was to participate in more depth than breadth paradigms as a consequence of their high productivity.

Hypothesis 3: The relationship between reference papers and current papers will be stronger for paradigms that are in groups, and strongest for paradigms that are in distinctive competencies.

We did simple regression equations for the following groups:

332 Paradigms (entire sample)	Nref	= .025 + .738 * NCurr (ns) (.000)	adj R^2 = .8034
111 Paradigms in the 'unrelated' piles:	Nref	= .056 + .178 *NCurr (ns) (.027)	adj R^2 = .0354
211 Paradigms in the 'related' piles:	Nref	= .128 + .734* NCurr (ns) (.000)	adj R^2= .8024
36 Paradigms in the related & DC pile:	Nref	= 1.01 + .692 * NCurr (ns) (.000)	adj R^2 = .7751

The relationship between publishing and having references in the 332 paradigms in this sample is exceptionally strong. Superstars submit more articles to the paradigms where their prior articles are part of the intellectual base. The coefficient is extremely significant. The overall Rsquare is exceptionally high.

But the relationship between publishing and having references in a paradigm is far weaker in those 111 areas of research that are considered ‘unrelated’. This relationship is insignificant at the traditional .01 level of significance.

The likelihood of a publication becoming a reference is roughly twice as good for superstars as for the general populous of researchers at Utah. The ratio of references to publications for the general populous is 0.38 (1.4/3.7). The ratio of references to publications for superstars is 0.75 (48.3/64.6).

It was interesting to see that the relationship was not stronger for the small sample of (36) paradigms that were in distinctive competencies and were in the ‘related’ pile of paradigms. At present, we’d have to conclude that there really is no difference, as long as the research area is part of the research strategy of the star researcher.

Discussion

The ability to describe the research strategy of an individual using a card-sorting technique is a promising development. The fact that most of the researchers enjoyed the process, and felt that the results captured the way they organized their research activities, indicates to us that the results are of high quality. The fact that the actual configurations of groups were so different is intriguing. This may provide a way, in the future, to better understand the diverse types of research strategies that star researchers pursue.

The results confirm that superstars are at the heart of many distinctive competencies. Their extremely high productivity (in both current publications and getting their publications accepted into the paradigm reference structures) would naturally result in their having more areas of depth than breadth. Having 60% of their publications in depth paradigms is roughly twice the expected value of 29.6%.

Evolution of Paradigms

There are only a few research groups that have examined how paradigms (document clusters based on co-citation analysis) evolve over time. In the 1980’s, one research group was led by Henry Small. Small was one of developers of co-citation analysis (Small, 1973), and applied this method while working at the Institute for Scientific Information (ISI) (Small, 1999, 2006). The other research group was led by Len Simon. Simon was a consultant to ISI, and started a company (the Center for Research Planning) in order to develop a formal model of science. Both groups used the ISI databases. Both used co-citation analysis (clustering highly cited references to form structural elements and then assigning current papers to these reference structures). Both used references to link these structural elements over time. Both observed patterns of birth, stasis, splitting, merging and death as far back as 1990.

The two approaches, however, had quite different purposes. Small, and many others that follow his tradition, focus on the most highly cited references. In the early 1980's, he selected a 1% threshold as a reasonable criteria. Focusing on just the top 1% of the most highly cited references resulted in about 51,000 references that were put into 9400 clusters (Franklin & Johnston, 1988). This 1% criterion has lasted for over 25 years. In his most recent studies (Small, 2006), the criteria has become more restrictive – the references could only be in the top 1% of the most highly cited references that were published in the past 6 years. This resulted in the identification of only 20,000 references that were put into 5,000 clusters. Small, and corresponding approaches that have very high and restrictive thresholds, identify elite groups (one can think of them as the 'hot' areas of research).

Simon, and followers such as Klavans & Boyack, had a different purpose – to identify the overall structure of science. They wanted to include science that was hot, cold, and everything in between. This meant including far more references and clustering them into far more clusters. In the early 1980's, Simon clustered 128,000 references (the most could be handled by contemporary computer hardware and software) and grouped them into 28,000 clusters (Franklin & Johnston, 1988). The most recent study, used in this project, clusters 2.1 million references into roughly 80,000 paradigms. Simon, and the corresponding approaches used by Klavans & Boyack, have much lower thresholds that identify as much of the structure of science as possible.

One can characterize these two approaches as elitist (focusing only on the hot areas of science) and populist (identifying as many structural elements as possible). An early study using the populist approach may help to illustrate the consequence of these different approaches. Klavans and Simon had been asked to evaluate SmithKline Beecham's (SKB) research portfolio in 1990. In the process, they identified the work by Marshall (on a bacterial cause for gastric irritation) as being very 'hot' (building on recently published references that were rapidly being cited). The analysis also pointed out that Marshall's work was a significant threat to the existing GI research program at SKB (there were links between this hot research area and the more common research areas that assumed that gastric irritation was due to the secretion of acids). This threat had not been identified by any of the content experts that SKB hired to do a research portfolio analysis. (Marshall's work was not generally acknowledged as 'hot' until the mid 1990's, and he correspondingly won the Nobel Prize for this work in 2005.) The head of research for SKB validated this result before acting upon it in 1991.

It's doubtful that the elitist approach would have identified Marshall's work as quickly. The study mentioned above used ISI data from 1988, and Marshall's work did not meet the 1% threshold that quickly. The elitist approach would not be able to evaluate the impact of Marshall's work because the target areas (the colder areas of gastric research) would not have been identified. While we acknowledge that the elitist approach can successful identify hot areas of science (albeit not until such work meets the thresholds for 'hot'), we point out that one cannot build an evolutionary model solely on an elite structure that ignores over 90% of the social structure in science.

The populist approach is used in this study. Separate models were created for eight years (2000 to 2007). Separate calculations for the distinctive competencies at the University of Utah were completed. The only question remaining was how to link these structural elements to best show evolutionary patterns. The following section describes our approach and presents the results from interviews with nine researchers that were central to distinctive competencies.

Methodology

The methodology consisted of (a) creating threads (linked paradigms over time) and then (b) a card-sorting technique where interviewees could validate the coherence of the threads associated with a distinctive competency, and then assess whether another set of threads are part of the scanning environment related to the distinctive competency.

Threads: Threads are paradigms that are linked over time. This linkage was based on reference-matching. Reference-matching, independently developed around 1990 by both Small & Simon, took the reasonable approach that paradigms (as defined as a cluster of references) are related if they have overlapping references. Paradigms cannot have overlapping references within a year (this is a consequence of co-citation analysis, which assigns each reference to one and only one cluster). However, paradigms from adjacent model years will have overlapping references.

Ideally, one wants to have a similar reference structure between years. Reference structures are highly similar when you use extremely high citation thresholds (a large fraction of the same highly-cited references will be used year after year). But, when we looked at the actual statistics using the new (and far lower) threshold for including references, we found that the between-year reference structures had low similarity (about 33%). Of the roughly 2 million references in the model in a given year, only 1 million would be in the model for the following year. A different 1 million references would be take their place in that second year. The total number of unique references to properly describe the relationship between any two years is thus around 3 million, and the number in common is only about 1 million.

The low level of similarity between reference structures using our lower thresholds is due to micro-bursts. There would be a burst of citations to a reference paper that would not exist in the prior or subsequent year. These micro-bursts would allow a reference paper to be included (and be a part in creating) the paradigmatic structure during a year. But the reference would not be used in the adjacent years, and would therefore not be available to match paradigms between years.

Our solution to this problem was to 'add in' the dropped (or added) references. Assume, for example, one identifies the 2 million most highly cited references in 2006 and the 2 million highly cited references in 2007. One can then identify the 1 million references that were dropped (they were in the 2006 paradigmatic structure but not the 2007 paradigmatic structure). These 1 million references can be added into the existing 2007 paradigmatic structure (using the same co-occurrence methodology that created the structure). This doesn't change the structure – it only tells us where the dropped references would have occurred given the new paradigmatic structure.

The same procedure can be used to add the 1 million references that were new in 2007 into the 2006 reference structure. This doesn't change the structure – it only tells us where the additional references would have occurred in the older paradigmatic structure.

This results in a 97% similarity in the reference structures between two years. In addition, the average number of references per paradigm increases (to 38 references). This provided a very rich data environment in which to examine year to year overlaps in reference structures. We used a very high year-to-year threshold for paradigm similarity (30%) before clustering. A paradigm with 38 references would have to have 12 references in common with another paradigm before we would claim that the two paradigms were related.

This measure of paradigm:paradigm relatedness (from one year to the next, but not spanning two or more years) was then used to cluster paradigms (from 2000 to 2007) into 82,037 threads. Approximately 30% of all paradigms were isolates (they did not have the minimum 30% overlap with any paradigm in prior years or subsequent years). The percentage of isolates in 2007 was 48% (we expect that 18% of these paradigms will be linked with the paradigms that will be created in 2008). The average thread had 4.37 paradigms and only lasted for 3.6 years (with a minimum of two and a maximum of eight). From these statistics, one can sense that the vast majority of threads consisted of simple linear connections between paradigms from year to year.

Interview Strategy: We conducted nine interviews to learn more about the thread structure for eight distinctive competencies. We identified the star researchers for distinctive competencies and, based on the availability of these researchers, spent about ½ hour in a card-sorting exercise with each researcher to (a) validate and name the distinctive competency and (b) determine which threads were in the scanning environment for that distinctive competency. The threads associated with a distinctive competency were identified from the data (this simply required one to identify the 2007 paradigms in the distinctive competency and then determine which threads these paradigms were a member of). The related threads was determined by (a) determining the reference profile for the entire set of threads in the distinctive competency and then (b) determining which threads shared common references.

Naming (Labeling) of Competencies: The protocol that was used to help a researcher to name the competency in which he or she is a major player is quite simple, and is as follows:

1) Generate a card for each thread that intersects the distinctive competency. Each card should list the top 10 key phrases associated with the thread (high information phrases extracted from titles and abstracts), the top 5 highly cited papers that are the basis for the thread (first author name and title are sufficient), and the top 5 researchers in terms of paper counts in the thread. Card decks for each distinctive competency have been generated for the University of Utah competencies for use by Tom Parks or his designates.

2) Choose one of the top 4 or 5 University of Utah researchers in the distinctive competency and schedule a 10 minute (maximum) time to meet.

3) At the meeting, show the researcher the set of cards and explain the following:

   a. You have been identified as a key researcher in one of the noteworthy (distinctive) competencies at our university using bibliometric methods. A competency is a networked set of topics in which our university has a worldwide leadership position. We need your help to label, or provide a title for, this competency.

   b. Each card represents a topic in science. You can think of each card as representing a partition of science roughly equal in size to a session (or several joined sessions) at a scientific conference. Each card has a list of the key phrases, key cited references, and highly productive authors in that partition of science.

   c. The information on each card should be considered as a gestalt, without worrying about whether a single piece of information fits.

d. These topics (cards) are highly linked here at our university, and form a network of topics in which we have a worldwide strength.

e. Imagine that you are advertising this competency to prospective graduate students, or perhaps to a company that is looking to collaborate with someone in the field. We need to name this competency in a way that will be accurate, sufficiently detailed, and that will attract students and collaborators.

4) Ask the researcher to provide a name for the competency based on the sum of the cards.

The results of those interviews were very clear – all researchers interviewed were able to easily provide a name for the competency based on the cards they were shown. These names are given in the following table.

Table 1: Names of distinctive competencies from interviews.

Competency	Name
DC03	*Stable isotope ecology (light isotopes)*
DC06	*Water-soluble polymers for drug delivery; subcellular targeting*
DC07	*Molecular diagnostics for inherited diseases*
DC08	*Combined "receptor and lipid" signaling*
DC18	*Cardiac electrophysiology and related imaging techniques*
DC29	*Scientific visualization – point sets, isosurfaces, and volume rendering*
DC59 (two interviews, two names)	*Ultra-high energy cosmic rays* *High-energy particle astrophysics*
DC106	*Medical image reconstruction - x-ray emission tomography*

Threads, DCs and Scanning Activities: Figure 4 is an example of the threads associated with distinctive competency #3. In this case, we asked experts at the University of Utah to determine whether there was a coherency to the threads associated with 2007:DC03 (the threads shown at the bottom of the figure that are linked by the shaded box). This was done by creating summaries of the research for each thread (on cards). The key researcher for DC03 labeled this area as 'ecology of isotopes', and indicated that the set of threads in the distinctive competency was a coherent set. The isolate (the paradigm at the bottom of Figure 4 that was not part of a thread) was also considered a coherent member of this set by the researcher.

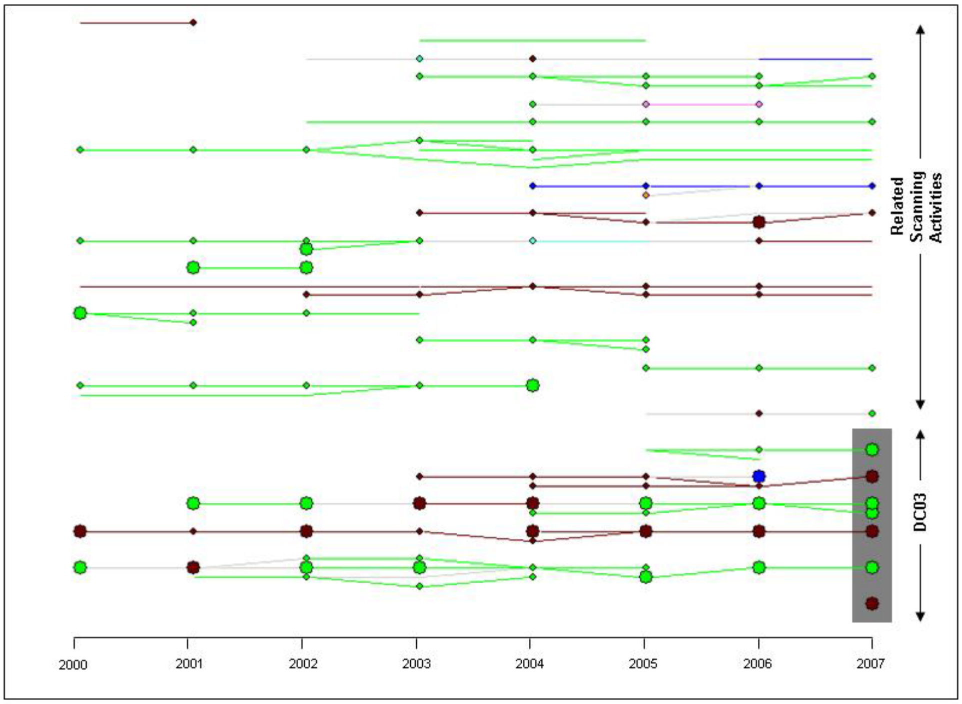

Figure 4: Threads associated with 2007:DC03 (Ecology of Isotopes)

We then asked the researcher to look at the related threads and determine if they were, in fact, related to the distinctive competency. The list of threads are ranked by the degree of relatedness (the most related threads are closest to the distinctive competency, the least related threads are most distant). Only the threads judged by the researcher to be related are shown in Figure 4.

The colors of the circles correspond to the field of research; green is biology, and dark brown is earth science. This research deals with the ecology of isotopes. The distinctive competency had high coherence and the related threads were part of the scanning process for this distinctive competency.

The size of the circles in Figure 4 tell us if the paradigm had depth (represented by a larger circle), breadth (represented by a smaller circle), or whether there was no publication by Utah researchers in the paradigm (the circle is drawn with no radius and appears as a line). As expected, all of the paradigms associated with DC03 in 2007 had depth (this is part of the definition of a distinctive competency). When one goes back in time, one sees that linked paradigms from previous years were also depth paradigms.

Most of the threads are long (an average of 4.6 years for the threads associated with scanning, and 5.1 years for the threads associated with the distinctive competency). Since these threads are based on a specific set of references (those associated with the distinctive competency), we suspect that a thread that stops in one year may actually be connected to a thread that starts in the next year. The only reason these threads aren't connected is that the threshold we initially applied is quite high (30%).

Figure 5 shows the threads associated with 2006:DC06. The expert validated the coherence of the distinctive competency and selected the threads associated with the scanning environment.

This distinctive competency built mostly on chemistry (blue) and a medical specialty (red). It appears that there are many isolates (paradigms without threads), the threads are relatively short, and many stop right before others begin. We strongly suspect that death of a thread followed by a birth in the next year of a related thread is actually a case of a thread that is linked at a lower threshold level.

This distinctive competency has a sparse scanning environment. If (as we suspect) the thread at the top of the figure (depth between 2000 to 2004) is connected to threads that start in 2005 and are part of the distinctive competency, the scanning environment is even sparser. There is a significant difference in the scanning environment of DC03 and DC06.

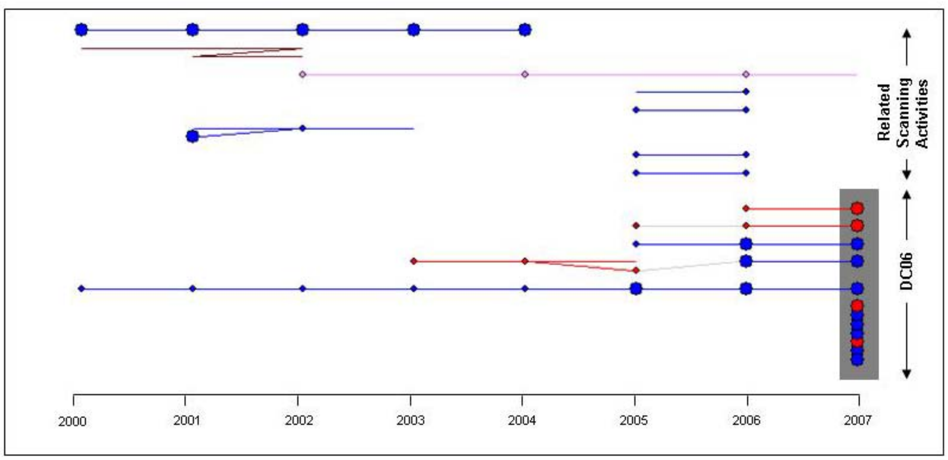

Figure 5: Threads associated with 2007:DC06 (Water Soluble Polymers for Drug Delivery)

Figure 6 shows the thread network associated with DC07. The pattern is different from the previous two examples. This set of networks is almost entirely medical (red), with only a few paradigms dealing with engineering (light blue). Note the very large network in this picture (it is the network that is most related to the reference structure of the distinctive competency. This is the first time we see a relatively larger network in the scanning environment.

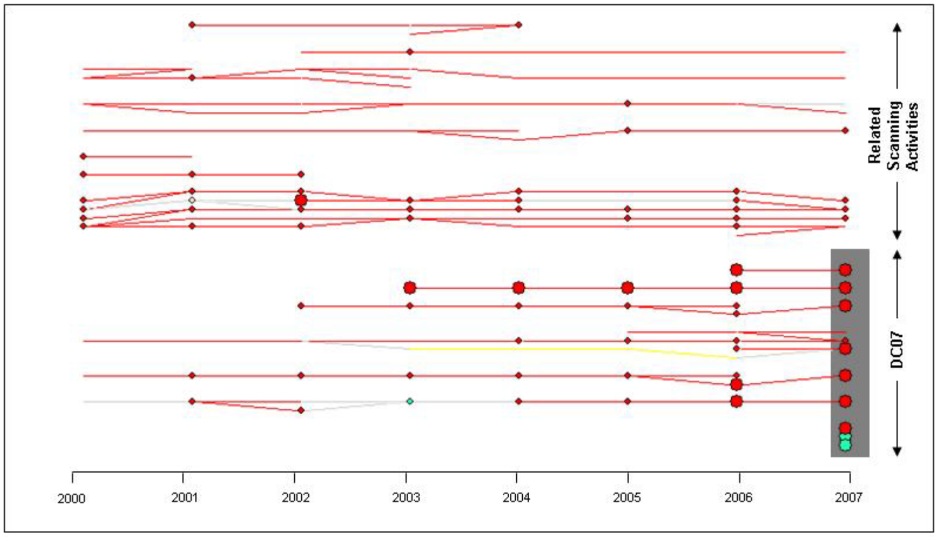

Figure 6: Threads associated with 2007:DC07 (Molecular Diagnostics for Inherited Disease)

One of the threads in DC07 is very interesting – the one in the middle (vertically) of the shaded box that has multiple paradigms in 2007. This thread had very little activity by Utah researchers until 2007. In 2007, only one paradigm had depth, while there were branches to two other 2007 paradigms, one of which was as breadth paradigm for Utah, and the other in which the university did not publish. This is the first example of a case where breadth and absent paradigms are linked to thread associated with a distinctive competency. However, while these non-depth paradigms are in the shaded area, we re-iterate that the distinctive competency only consists of the larger (depth) paradigms.

The next distinctive competency (2007:DC8, Figure 7) is highly multidisciplinary. The research within the distinctive competency deals with medicine (red), engineering (light blue), computer science (pink) and chemistry (blue). The related networks (at the top of the figure) are also from a broad set of disciplines. This is another case where there are non-depth paradigms in the shaded area, and where there are isolates (depth paradigms) associated with the distinctive competency.

The four examples in Figure 8 reveal similar patterns between the threads in a distinctive competency and those associated with scanning activities. By contrast, DC18 and DC29 (Figure 8) have more isolates in the distinctive competencies (depth paradigms that were not linked to prior paradigms). DC59, which deals with ultra-high energy cosmic rays, has an exceptionally large network. Note that the entire set of threads in this chart are purple (the color associated with physics). The last example is of a more multidisciplinary area (DC106- Medical Image Reconstruction).

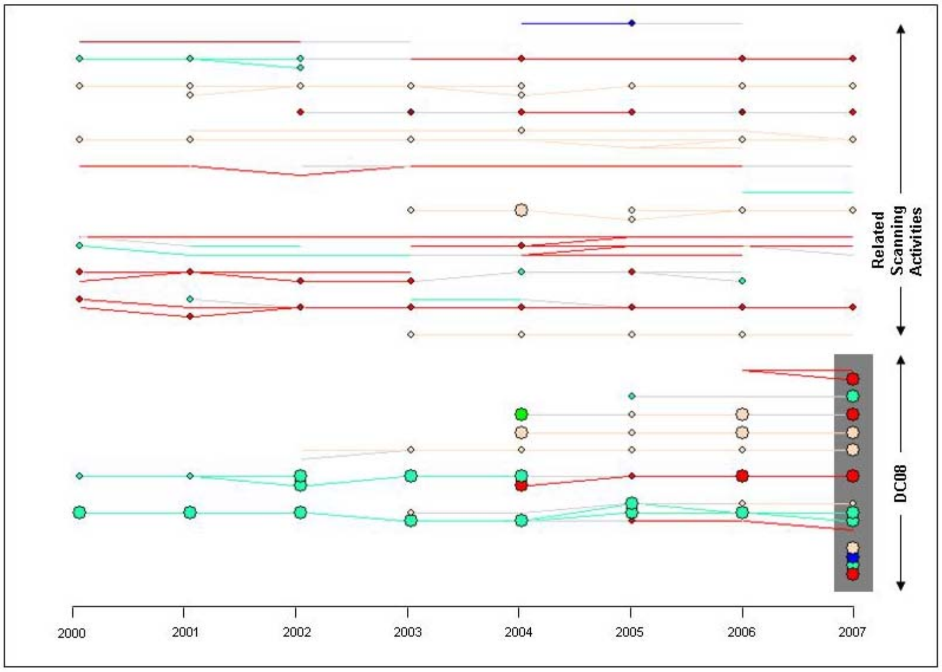

Figure 7: Threads associated with 2007:DC08 (Combined Receptor and Lipid Signaling)

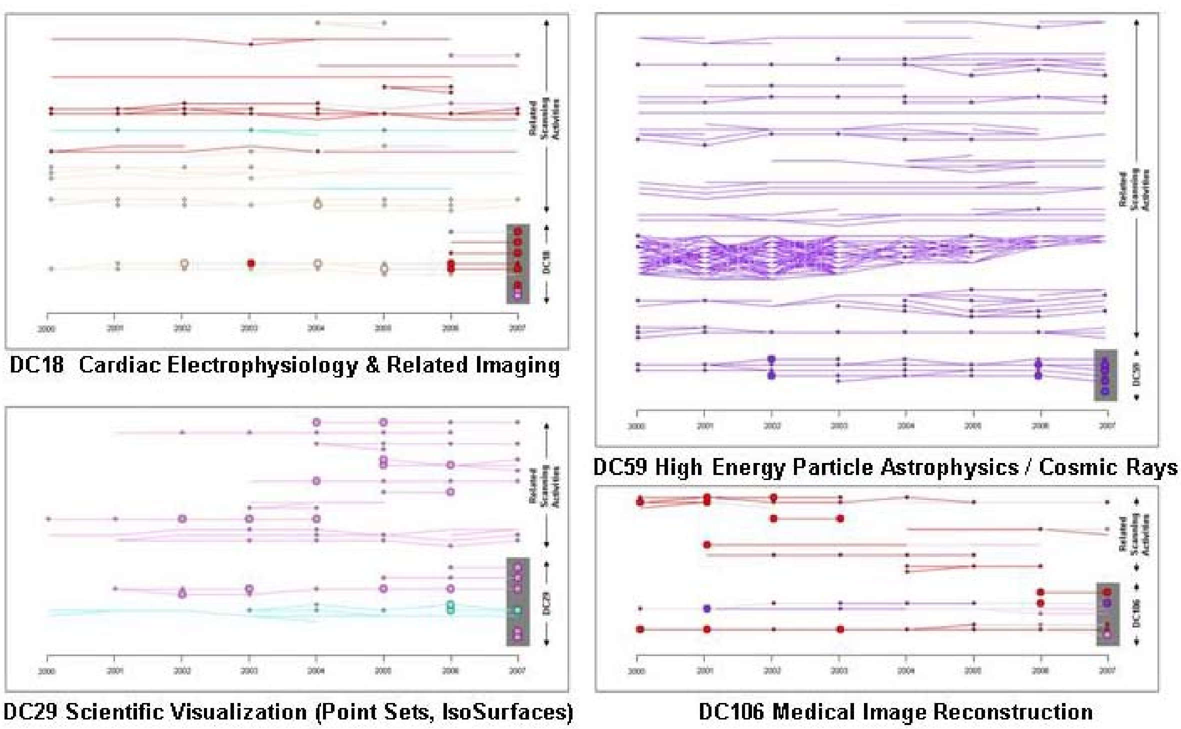

Figure 8: Threads associated four Distinctive Competencies

Overall, the interviews with the star researchers had the same results. The threads associated with distinctive competency formed a coherent set. The interviewees could easily label these threads (they were intimately involved with this research). The interviewees had little problem deciding which of the related threads (shown at the top of the visualizations) were related to the distinctive competency. The visual displays of these evolutionary patterns show differences in the distinctive competencies and their scanning environments. More importantly, they also show that each of these distinctive competencies have a history, have persisted in some form for many years, and can be expected to continue into the near future provided that the key researchers remain at the University of Utah.

Discussion

The identification of threads appears to be a highly accurate way to track the evolutionary paths of distinctive competencies. The decision to 'add in' the missing references and to correspondingly generate large reference structures for each paradigm, creates a more accurate measure of paradigm:paradigm relatedness. This approach is also consistent with the (populist) assumption that more references are better (they better represent the social structure of research).

The interviewees considered each distinctive competency as a coherent group and were able to label these sets of threads from the data we provided. They were able to evaluate which threads were related.

What is most intriguing, however, is the actual patterns shown in the eight examples. These patterns raise a host of issues that are worth following up in the future. Following are a highlight of three possibilities.

First, the ability to identify the scanning environment of a distinctive competency is a major leap forward. As mentioned previously, some of the threads in the scanning environment should actually be part of the distinctive competency environment. This is a correction that can be easily made. Once these corrections are made, we will be able to identify the breadth paradigms that are associated with a distinctive competency. Right now, the ratio between depth and breadth is about 1:4 (20% of the articles published by a university are part of a distinctive competency). The breadth paradigms can be partitioned further. Perhaps the ratio would be 1:1:3 (the analysis has not been done yet). The analysis can be done using the distinctive competencies and related threads from hundreds of universities to determine what it means to have different ratios.

Second, the evolutionary patterns shown in these eight examples suggest alternative methods for monitoring the health of a distinctive competency. At present we simply show the growth in publications in a distinctive competency as an indicator of health. But this assumes that the same reference structure would persist and totally ignores the role of the scanning environment. We now know that the reference structures are unstable, and there is a wide variation in the types of scanning environments for each distinctive competency. It may be more important to track what happens in the scanning environment. The richness of the scanning environment (and the ability of researchers to draw from this environment) may be a much better early indicator of problems.

Third, the existence of 'isolates' (paradigms that are not related over time) and thread statistics (birth, stasis, splitting, merging and death) provide a whole new lexicon for understanding how distinctive competencies evolve. Isolates may be examples of the paradigmatic change that was originally proposed by Kuhn (1970). The fact that a discipline consists mostly of isolates may mean that the research in that area is 'pre paradigmatic'. We may find that very large networks (such as the one in DC59) is associated with 'big science' (Price, 1963). The objective measures proposed here may provide a mechanism for understanding and evaluating the broader question of how science evolves.

We look forward to working on these research issues in the future.

Conclusions

The underlying purpose of this project was to identify the research strengths of the University of Utah. This was accomplished by utilizing recent technology (Spotlight), analyzing possible weaknesses in the technology, and taking part in the development of the next generation of the technology.

The Spotlight technology is a far more accurate approach for identifying research leadership than anything currently available. By identifying paradigms from the citation patterns in the literature, partitioning these paradigms into areas of depth and breadth, and then further clustering the depth paradigms, we were able to more accurately identify the underlying strengths of the university. Each expert we talked with has validated these results.

The possible weaknesses were also addressed. Having better descriptions of a distinctive competency simply requires one to ask the star researchers in that area. The concern that the loss of star researchers would degrade a distinctive competency did not bear out – individuals could leave but, as long as there was a team in place, the competency survived. The concern that strengths are understated because of recent hires proved to be correct. Yet, the recent hires didn't add new competencies – they reinforced existing ones. And finally, the possibility of using these data to identify corporate partners was explored.

Most of this report, however, has focused on frontier issues. We've just scratched the surface in our understanding of the role of star researchers in a distinctive competency. Our ability to extend the analysis to the scanning environment associated with a distinctive competency and to the structural changes in threads over time will provide a whole new set of insights into how competencies evolve.

Our knowledge about how distinctive competencies evolve is still meager at best. We really have little knowledge about the predictive indicators of the health or illness of a distinctive competency. So far, we've stood in the present and looked back into the past. Our next challenge will be to stand in the past, armed only with what was known at the time, and determine if there are any indicators that can anticipate the future.

References

Chen, C. (2006). CiteSpace II: Detecting and visualizing emerging trends and transient patterns in scientific literature. *Journal of the American Society for Information Science and Technology, 57*(3), 359-377.

Franklin, J. J., & Johnston, R. (1988). Co-citation bibliometric modeling as a tool for S&T policy and R&D management: Issues, applications, and developments. In A. F. J. van Raan (Ed.), *Handbook of Quantitative Studies of Science and Technology* (pp. 325-389). North-Holland: Elsevier Science Publishers, B.V.

Klavans, R., & Boyack, K. W. (2009). Toward a consensus map of science. *Journal of the American Society for Information Science and Technology, 60*(3), 455-476.

Klavans, R., & Boyack, K. W. (2010). Toward an objective, reliable and accurate method for measuring research leadership. *Scientometrics, forthcoming*.

Kuhn, T. S. (1970). *The Structure of Scientific Revolutions* (2nd ed.). Chicago: University of Chicago Press.

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

Small, H. (1973). Co-citation in the scientific literature: A new measure of the relationship between two documents. *Journal of the American Society for Information Science, 24*, 265-269.

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

Small, H. (2006). Tracking and predicting growth areas in science. *Scientometrics, 68*(3), 595-610.