

PORTER, A.L., AND ZHANG, Y. (2015). "TECH MINING OF SCIENCE & TECHNOLOGY INFORMATION RESOURCES FOR FUTURE-ORIENTED TECHNOLOGY ANALYSES," IN GLENN, J.C. AND GORDON, T.J. (EDS.) *FUTURES RESEARCH METHODOLOGY VERSION 3.1*, THE MILLENNIUM PROJECT, WASHINGTON, DC ONLINE AT THE *GLOBAL FUTURES INTELLIGENCE SYSTEM IN FUTURES METHODOLOGIES*: [HTTPS://THEMP.ORG/](https://themp.org/)

TECH MINING of Science & Technology Information Resources for Future-oriented Technology Analyses

By
Alan L. Porter and Yi Zhang

[I. History of the Method](#)

[II. Description of the Method](#)

[III. How To Do It](#)

[IV. Strengths and Weaknesses of the Method](#)

[V. Use in Combination with other Methods](#)

A. Outputs

B. Uses

[V. Frontiers of the Method](#)

[Appendix](#)

Acknowledgments

This paper draws heavily on the work of colleagues associated with development of “Technology Opportunities Analysis” at Georgia Tech. We particularly note the “TF” (technology forecasting) contributions reflected herein of Nils C. Newman and Robert J. Watts, and contributions to development of our text mining approach and software by Paul Frey, Scott Cunningham, Donghua Zhu, Douglas Porter, and Brian Minsk. Webb Myers and David Schoeneck have advanced the analytics. Long-term colleague, Fred Rossini, heavily influenced our futures perspective. We thank Joe Coates, Ted Gordon, and Peter Bishop for their reviews on the initial version. On this update, Yi Zhang enriches treatment of text analytics and combining tech mining with technology roadmapping. Work on this revision benefits from partial support of US National Science Foundation (Award #1064146 – “Revealing Innovation Pathways: Hybrid Science Maps for Technology Assessment and Foresight”). The findings and observations contained in this paper are those of the authors and do not necessarily reflect the views of the National Science Foundation.

I. HISTORY OF THE METHOD

Tech Mining is short for “text mining of science & technology information resources.” The key premise underlying it is that *intelligence* (as in Central Intelligence Agency) is a prime requirement for effective technology management and Science, Technology & Innovation (ST&I) policy-making. Organizations operating in competitive and/or collaborative environments must track information on external ST&I developments.

ST&I intelligence underpins *Future-oriented Technology Analysis (FTA)*. FTA is a growing community centered on the ongoing series of conferences hosted by the Institute for Prospective Technological Studies (IPST) -- <https://ec.europa.eu/jrc/en/event/site/fta2014/>. It encompasses technology forecasting, assessment, foresight, and roadmapping.

So, how do we gather ST&I intelligence? We seek external information sources to mine. The earliest roots of this method lie in traditional literature reviews. “Technology monitoring” reviews the literature for a given technology of interest, usually scanning a broad range of sources. It seeks to digest sizable amounts of raw information to identify developmental patterns and key events (breakthroughs). Closely related, “environmental scanning” aims to provide early warning about important changes in various environments (e.g., competitive, policy or natural contexts; see the *Environmental Scanning* paper by Gordon and Glenn).

The advent of widely accessible, electronic information resources qualitatively changes these intelligence gathering efforts. One can now process huge quantities of information easily and effectively. “Bibliometrics” – counting publications or other bibliographic items – can help track scientific and technological developments; tech mining goes further.

Tech Mining is a form of “Content analysis” that has blossomed with the availability of electronic datasets and analytical software. One examines patterns in term usage to infer emphases in given domains. Text mining empowers powerful variations on content analysis. This detective work is not restricted to technology analyses – for instance, intelligence agencies may analyze news feeds, tweets, or blogs to track foreign political leaders’ shifting interests.

Data mining seeks to extract useful information from various forms of data, but it usually emphasizes analyses of numerical data (e.g., linking your credit card purchases to your demographics). Tech Mining especially exploits text data sources of various sorts -- keying on *structured text* sources (e.g., abstract records from managed databases that separate information into fields such as “author,” “publication date,” and “keywords”), with moderate facility on semi-structured text (tagged) sources, and limited utility with unstructured texts. It complements efforts to exploit “big data,” but those generally deal with unstructured text and other data forms. Tech mining might be said to key on “large data.” Specialties in text understanding are rapidly evolving with distinctive emphases. For our purposes, we just note emerging areas that concern text mining: “Computational Linguistics,” “Natural Language Processing,” and “KDD” (Knowledge Discovery in Databases).

This paper emphasizes the development of text mining tools to analyze ST&I *emergence*. Such intelligence extraction efforts need not be restricted to a particular subject matter. However, focusing on topical ST&I changes can lead to particularly useful “innovation indicators”

(discussed later and elaborated in the book, *Tech Mining*, by Porter & Cunningham – see “Selected Literature” at the end).

Who invented text mining? We don’t know. Talmudic scholars have tabulated content patterns in the bible for ages. Counting scientific publication activity dates back to the pioneering work of Derek Price in the early 1960’s (see Appendix). Content analysis dates back decades in the social sciences. Modern Tech Mining would count among its pioneers, Henry Small, Tony van Raan, Michel Callon, and Loet Leydesdorff (see Appendix).

What did we do before text mining? The benchmark in some regards is the traditional “literature review.” Therein, a researcher summarizes the most pertinent (e.g., 20 or so) literature sources pertinent to a given study. In an age of electronic information resources, we can still search for those few “nuggets” of really key reviews, models, and developments. What is different in Tech Mining is that “all” related literature is analyzed. In one study we analyzed the development of excimer lasers for machining. We located some 12,000 pertinent references. Instead of winnowing those down to find a handful of key items, Tech Mining elicits useful patterns from the entire 12,000 sources. For instance, one can track relative emphases on particular topics, over time, by key research groups.

II. DESCRIPTION OF THE METHOD

Tech Mining extracts useful intelligence from electronic text sources. This information can serve needs to know about current activities (e.g., profiling who’s doing what in a target area). It can also serve FTA interests in several ways:

- identifying R&D emphases that portend future developments
- providing time series for trend extrapolation and growth modeling
- generating “innovation indicators” that speak to the prospects for successful technology applications [The *Tech Mining* book identifies some 13 enduring Management of Technology (“MOT”) issues, that spawn about 40 questions, which some 200 innovation indicators help answer.]

Effectively answering MOT or ST&I policy questions is why one does Tech Mining. So, the first vital step is to spell out the key questions to be addressed. Knowing those will guide choices about which data sources to use and which analyses to pursue. Focus is critical; otherwise Tech Mining can wallow in too much data, with too many interesting – but not directly useful – analyses to pursue.

That said, let’s consider how one “mines” text, particularly field-structured S&T information resources. Let’s take an example case. Bob Watts was analyzing technology opportunities for the U.S. Army’s Tank-Automotive Research, Development & Engineering Center (TARDEC) in the mid-1990’s. He began to explore the potential for innovative applications of ceramics in automotive engines (especially tank engines). Reviewing local history uncovered earlier bad experiences. TARDEC had invested in ceramics R&D in the 1980’s without significant return on investment. The Army had stopped such funding. Nonetheless, Bob saw promise and pursued his investigation.

Searching the publication abstract database *EI Compendex* (also known as *Engineering Index*) yielded several thousand abstracts. In beginning the analysis, one likely wants to list the leading organizations researching the topic. Knowing *WHO* is engaged is vital – e.g., it may suggest potential partners. Another list might give the prominent “keywords” and/or title or abstract noun phrases (extracting via Natural Language Processing – NLP) descriptive of research emphases in these articles – i.e., we now add a *WHAT* dimension. Further analysis could then cross these two lists as a *MATRIX* showing which topics particular organizations mention frequently. That could help the Army spotlight particular organizations doing work of interest. For instance, General Motors might be working on structural ceramics and silicon nitride, whereas Sandia National Lab was addressing thin film ceramics and silicon carbide.

Many other combinations could provide insights into the emergence of a technology. To give one example, one could cross “keyword usage” by “publication year” to suggest which topics are “hot” (increasing emphasis recently). In this way we are combining *WHAT* & *WHEN* aspects – perhaps to help the Army set its R&D priorities. One might add a fourth dimension to the mining. For instance, in desktop text mining software, *VantagePoint*, you could pop open a detail window to see in which countries (*WHERE*) particular topics are pursued most strongly. [We’ll pick up this case story later.]

Which information resources you choose to mine determines what sort of intelligence you can elicit. Obviously someone checking levels of public concern over an emerging technology’s possible environmental hazard would select different resources to tap than someone investigating recently funded research.

We find it useful to distinguish six main types of information resources (Table 1). Many technology managers fail to take advantage of rows A and B; they rely almost exclusively on tacit judgment. This is folly. Those who obtain empirically based knowledge, in addition to using human expertise, will dominate those who don’t. Others equate electronic information resources with googling the internet (row B). That also misses a vital resource – row A. Tech Mining emphasizes extracting useful intelligence from databases (both publicly available and organization-confidential). Databases represent someone having compiled, filtered, and formatted text (or other) data, making this of higher value than the unfiltered mess of the internet.

Table 1. Six Information Types

Medium \ Message	1) Technology	2) Context
A) Databases	Research funding, publication & patent abstracts, citations	Business, market, policy, popular opinion
B) Internet	Technical content sites	Company sites, blogs, etc.
C) People	Technical experts	Business experts

Analyzing a spectrum of databases that span different stages along a technological innovation can help benchmark how mature the technology is. One expects S&T activity to migrate along this continuum, over time, as an innovation progresses, but recognize that technological innovation rarely progresses linearly:

- Fundamental research [reflected by research project funding (e.g., NSF Award or NIH CRISP database) and publication abstract databases (e.g., Science Citation Index; PubMed)]
- Applied research & development [reflected, for instance, by engineering oriented publication abstract databases (INSPEC, EI Compendex)]
- Invention [consider patent databases such as Derwent World Patent Index and PatStat (from the European Patent Office for academic use)]
- Commercial Application [e.g., new products databases, Business Index, marketing data sources]
- Broader Contextual Factors and Implications [e.g., Lexis-Nexis, Factiva, Congressional Record]

Having commended such databases as extremely valuable repositories, we need to say that they don't tell the whole story. Database content lags – it takes time for most sources to make information available (e.g., publication processes may take a year or more; patents can take several years), so the internet is more current. Not all sources are covered by databases; so one should seek enriching internet content as well. And database access – especially unlimited use licenses vital for Tech Mining – is costly (with relatively few exceptions – e.g., PubMed is free). We like to analyze database content (typically using one or a few key sources on a given study) first; then augment with internet searches (e.g., to check websites of key research organizations). A notable frontier for Tech Mining is the improving software able to find, compile, and format internet content (c.f., www.q12.com/). Also, engaging experts to review the empirical findings and help interpret those (Table 1, Row C) is critical to generate credible technical intelligence and subsequent FTA.

III. How To Do It

Here's a 10-step approach to Tech Mining.

1. Spell out the focal intelligence questions and decide how to answer them
2. Get suitable data
3. Search (iterate) & retrieve ~abstract records
4. Import into text mining software [e.g., Thomson Data Analyzer (TDA), VantagePoint]
5. Clean and consolidate the data - Clumping
6. Analyze
7. Visualize
8. Integrate with internet analyses and expert opinion
9. Interpret and summarize findings; communicate those (possibly multiple ways)
10. Standardize and semi-automate where possible

We'd like to illustrate via three case examples: A) The Army's analysis of the potential new use of thin-film ceramics in engines; B) National benchmarking of R&D initiatives and C) . We've introduced (A) already, so let's step through the Tech Mining method using that, then later we can enrich our perspective with (B).

Step 1 – The Army wanted to assess whether thin-film ceramics held potential for its tank engine applications. If so, then it sought guidance on how best to pursue such development. In particular, might they want to partner with particular R&D organizations?

Step 2 – We’ve introduced a few of the hundreds of available databases. The Army (TARDEC) had previously determined which of these to license for its technical intelligence and research needs.

Step 3 – Search entails several key elements. One must formulate the inquiry – what do you want to know, about what topic? Properly bounding the inquiry is vital – too broad is bad -- profiling all R&D relating to “ceramics” would be a herculean task and likely counterproductive. Too narrow is bad -- our interests contrast against traditional librarian efforts to find a few “just right” publications. We want a middle ground -- a large swath of related research – e.g., ceramic development associated with engine applications – to generate an R&D landscape.

We often find it effective to conduct a “quick and dirty” initial search; download a sample of resulting abstract records; and quickly analyze those using our text mining software (*VantagePoint*). By browsing the resulting term lists (i.e., keywords), preferably with benefit of expert review, we can sharpen the search phrasing. Then we redo the search and retrieve the actual record set to be analyzed in depth. This process can be done in a few minutes given desktop access to the searchable database(s).

Step 4 – In the Army ceramics study, Bob Watts downloaded his search results into *TechOASIS* software (see Appendix) for further analyses. Such software is tuned to analyze field-structured text records.

Step 5 – Data cleaning might be called “clumping.” We want to consolidate author or organization name variations (especially if we’re combining search results from multiple databases). One handy thesaurus aggregates organizations by types: universities; companies; government or non-governmental organizations; and hospitals. Then, a handy macro (a little software script) shows the extent of industrial R&D on a topic in a pie-chart as a useful innovation indicator.

We have devised a topical Term Clumping process (Zhang et al., 2014a) to:

- a) consolidate related terms via general purpose and special thesauri
- b) clean up terminology (e.g., to combine singular and plural forms of a terms, and to overcome small name discrepancies) via fuzzy matching routines
- c) identify more informative terms or remove common terms (depending on the data set) via Term Frequency Inverse Document Frequency (TFIDF) analysis
- d) consolidate highly related terms via correlation rules (e.g., co-occurrence of terms in many records indicates possible relationship)
- e) cluster topics for further analyses via term clustering approaches (e.g., Principal Components Analysis, Latent Dirichlet Allocation, K-Means, etc.).

Step 6 – Analyses can take many paths and forms. These should remain mindful of Step 1 – to address the essential questions driving the analysis. Text mining software does basic operations,

such as counting which author organizations show up most frequently in your set of records. It then allows easy exploration – e.g., via a detail window to examine a given organization’s topical emphases over time. The software can also perform more advanced functions, like extracting noun phrases from abstracts, clustering related entities, and mapping research networks (based on co-authoring, co-citation, or whatever).

Many such operations can be done without text mining software. To begin, many website database search engines enable you to tally occurrences of particular terms, slice search sets by year, and so forth (c.f., SciFinder; EI Village, Web of Science). Once you’ve downloaded search results, you can use common literature tools (e.g., EndNote) to handle the records; MS Excel to help count up frequencies; and so forth. However, if you frequently want to perform such analyses, or want to get at more intricate innovation indicators, we’d recommend you investigate software aids (see Appendix).

Step 7 – Visualizing aims to help represent your findings effectively. Many forms are possible, including:

- Social network analyses – e.g., *VantagePoint* offers cross-correlation maps that show interconnections among countries, organizations, or individuals who co-author within a set of documents
- Geographical mapping – to locate R&D or other activity concentrations
- Science or patent overlay maps – to locate R&D activity among disciplines or patent categories (c.f., Kay et al., 2014).

An important instrument for technology visualization, Technology Roadmapping (TRM), combines qualitative and quantitative methodologies to indicate technology development dynamics on multiple levels. TRM visualizes relations among elements (e.g., R&D activities, system components, target markets) over time. It can help show evolving roles of technology push, market pull, science policy, etc., over time periods. TRM, in turn, can interface with other modeling, such as Technology Delivery Systems, Semantic TRIZ, Triple Helix and Forecasting Innovation Pathways (FIP). [See Zhang et al., 2013, 2014b, 2014c; Guo et al., 2012; Robinson et al., 2013.]

Step 8 – We favor an iterative process that takes updates & enriches Tech Mining results by checking internet resources. E.g., if several organizations appear to be trend-setters, check their websites for current directions and possible novel approaches. The FIP approach, just mentioned, tries to engage topical experts throughout the analyses, but especially via workshops to digest, challenge, and interpret empirical (Tech Mining) findings, and to anticipate likely innovation pathways forward and attendant issues.

Step 9 -- Consider how your key users like to receive intelligence. Accordingly, you may well want to judiciously combine: Oral and written reporting; electronic and paper reporting; text, tables, and figures – possibly with video or animations; and interactive exchange (e.g., workshop).

Sometimes, as analysts, we neglect to explain. We wallow in the data so deeply that we forget to point out the implications of our tallies and figures. And, to paraphrase a Russian colleague, “decision-makers don’t want pretty pictures, they want answers.”

Step 10 – Consider sacrificing some flexibility of the strategy in which every study starts from Ground Zero. We advocate an “RTIP” – Rapid Technology Intelligence Process – premised on standardizing and automating Tech Mining analyses. There are multiple advantages to this:

- Technology managers and professionals become familiar with the Tech Mining outputs. They learn how to get the answers they need from them, thereby making better decisions informed by empirical knowledge.
- Analyses are done much faster. One company told us how a series of “quarterly” competitive technical intelligence reports could now be done in about 3 days. Suddenly, that makes data-based intelligence accessible to pressing decisions.
- Analyses are cheaper. By semi-automating data cleaning and analysis steps (i.e., writing software scripts), analysts can generate more value, but use fewer resources.

IV. STRENGTHS AND WEAKNESSES OF THE METHOD

If you noticed, we mentioned how Tech Mining can answer four types of questions – who, what, where, and when? The other two of the so-called “reporter’s questions” – how and why? – almost always require expert opinion to infer processes (how?) and reasons (why?).

No FTA method stands well alone, including Tech Mining. Basic tabulations require thoughtful interpretation. “Innovation indicators” require complementary expert opinion to test observations and relationships generated from the information resources.

Do you need expert help to mine technological information resources? Ron Kostoff, a leader in technological text mining (see Appendix), argues that heavy involvement of substantive experts in the process is essential. This yields advantages in assuring sensible interpretation, proper treatment of terminology, and “buy-in” from those experts who become actively engaged.

The information products generated can be no better than the sources mined. So, if vital elements of technological development are not treated in your sources, you will not uncover them. For example, if you base your analysis on patents on an emerging technology for which the developers do not generally seek a patent, you will be misled.

Certain analyses can become complicated. “Transparency” enables users to understand results. Relationships discovered via Tech Mining are based on co-occurrences. These entail terms appearing in the same documents more frequently than expected from their overall popularity. It is important to recognize that clustering and mapping of such relationships generally do not have a single “right” way to go. Representations, unless invalid (i.e., misrepresentations due to incorrect statistical analyses or such), prove their value by being used – that is, by depicting relationships so that a user gains new insights.

Tech Mining is still new to technology managers and professionals, as well as futurists. As such it faces credibility issues. If one makes decisions relying heavily on derived knowledge from

text analytics, whom can you blame if the decision goes awry? Many decision makers feel more comfortable leaning on expert opinion, even if it is parochial. We suggest combining Tech Mining results with expert opinion to take advantage of the strengths of each.

A good way to get started at Tech Mining is to contract for a couple analyses. If those usefully inform MOT decision-making, then consider developing in-house capabilities. To do so, one first needs access to suitable electronic information resources. One then wants the tools to search, retrieve, and analyze selected records. The Appendix notes several software options.

Training is needed to perform Tech Mining effectively. Analysts familiar with treating “text as data” (e.g., the “Internet generation” who have been students since about 1995) can quickly grasp the concepts and generate useful findings with a couple of hours of training. Less analytically oriented persons, or those residing on “the wrong side of the digital divide,” may never get it. Training workshops, supported by ongoing access to advice, offer a reasonable middle ground for motivated analysts to pick up how best to use given software.

Costs are driven by database charges. Many universities already license a number of R&D databases useful for technology foresight, so there is no additional cost involved. Simple analyses can be performed using the basic analytical capabilities of the search engines (e.g., *Web of Knowledge*). To license the software illustrated in this paper, *VantagePoint*, costs some thousands of dollars (but varies by organizational type and needs). We see potential for university librarians to shift roles toward becoming “Research Navigators” who don’t just point clients to pertinent references, but generate “research profiles” to help researchers identify best opportunities, research administrators manage university priorities, and technology licensing professionals assess potential “open innovation” matches [e.g., companies situated to benefit from university Intellectual Property (IP)].

Putting information products into comfortable, easily grasped form is critical. We recommend blending numerical tabulations with graphical depictions and text interpretations – tailored to your audience’s preferences. An interesting option is to combine a report with a file containing the raw abstracts and the mining software. That can enable users (who are willing to learn to use the software) to dig down into the data themselves as the report triggers questions. For instance, imagine you’re reading a profile of ceramic engine developments and note that Sandia National Lab is researching thin film ceramics – you can instantly browse through their abstracts to see what they are doing. Note that to get to the full articles requires either a full text database, live links, or a visit to the library.

In observing what makes for effective utilization of text mining outputs, we distinguish users (and organizations) who prefer reports from those who prefer answers to their questions. For instance, a governmental agency (e.g., the Army) interested in ceramic engine applications, would most likely want a well-substantiated report. In contrast, we have found that many company technology managers want to know that well-founded analyses underlie results obtained, but they want their immediate question answered directly (no lengthy report).

Tailoring information products carefully to the key users’ needs is critical. An interesting approach is “packetizing” findings to avoid information overload. For instance, instead of

holding off until a 100-page report is ready, dole out small amounts on particular aspects that you feel the user wants and can readily digest. This can also help engage the user in the analytical process, facilitating buy-in.

Building familiarity with text mining outputs and how to use them effectively takes effort. Relationships between analysts and users must be strong. Factors that enhance prospects include:

- facilitating direct links between the analysts and users (interposing additional people to perform information searches, qualify statistical analyses, review results, etc., is deadly)
- engaging multiple analysts (and multiple decision makers) to develop a robust, learning network
- directing attention to successes so that the organization appreciates the value gained.

V. USE (IN COMBINATION WITH OTHER METHODS)

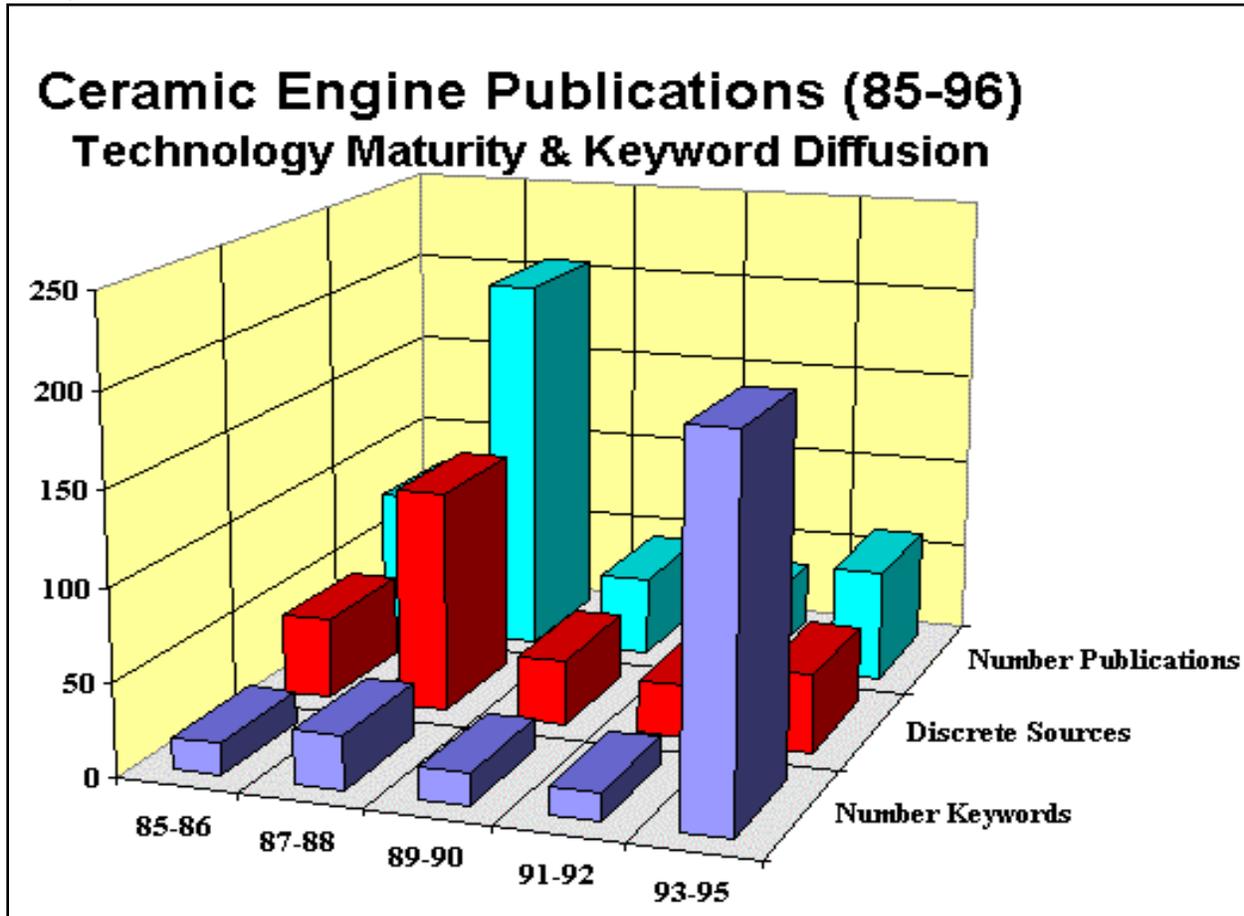
This section illustrates how Tech Mining can be used to inform Future-oriented Technology Analyses. Three case studies are offered. Throughout, keep in mind that text mining needs to be combined with other methods, particularly expert opinion.

A. Investigating the Potential of an Emerging Technology (thin-film ceramics) to meet Army Engine Needs

We earlier introduced the analytical challenges prompting this series of analyses on ceramic R&D. Figure 1 gave the key innovation indicator that alerted to maturation of this technology. The message in Figure 1's time-slices has several parts. Consider, first, the back row showing "Number of Publications." In the 1987-88 time slice, this reached 200 or more, and then crashed to ~35 in 1989-90. In the most recent time period, 1993-95, this has begun to recover slightly. The second row, "Discrete Sources," shows a similar pattern. In essence, following an initial burst of ceramics enthusiasm, funding and research effort had cooled [e.g., a so-called "hype cycle"].

The front row, "Number of Keywords," tells a markedly different story. This indicates that the richness of R&D discourse about ceramic engine applications had sharply accelerated in the recent time period. That was the signal for action.

Figure 1. Keyword Behavior Change as an Innovation Indicator (from Watts & Porter, 1997)



Watts presented these results to experts who confirmed that, indeed, the ceramics research arena had gotten more specialized and was moving toward practical applications. Emerging key terms implied serious research underway on particular materials (e.g., silicon nitride); special properties (e.g., fatigue testing); modeling; and, in particular, on thin film features. TARDEC senior management was convinced by this storyline that thin-film ceramics now could help the Army meet its engine needs.

The second Tech Mining phase pursued the prospects for thin film ceramic coating applications in the automotive context. TARDEC sought to identify the research leaders. Rather to the surprise of the mechanical engineers working on engine improvements, they were not from their community. Instead, it was the semiconductors sector doing the cutting edge R&D on thin ceramic films for chip design and development. Bob's Tech Mining work was able to span these traditional bounds by searching in EI Compendex. He identified particularly exciting R&D activity underway at Sandia National Lab and at a company. Those researchers had not considered coating tank engine parts!

To wrap up the story, TARDEC made contact with those research organizations and funded two large projects to explore coating applications. Adaptation of vapor deposition to pistons and

turbine blades proved workable. In 2004, a major production plant opened to coat used Abrams tank turbine blades for extended life. The ceramic coating offered additional advantages in terms of higher operating temperatures and reduced air pollution. So, Tech Mining had provided two key insights: 1) this research domain had significantly matured to offer real promise, and 2) potential development partners operated in a distant sector with which TARDEC mechanical engineers did not routinely interact.

B. Benchmarking National Research Initiatives

This Tech Mining case also involves the public sector (private sector work is almost always in confidence). Let's use this to illustrate variations in the "10-step" approach.

Step 1) Here, a Canadian Government group sought to benchmark a number of research clusters in other countries as comparisons to assess their own initiatives – i.e., to do research evaluation. We draw illustrations here from their analysis of U.S. photonics research.

Step 2) We searched in Science Citation Index (SCI)
[a companion study investigated patenting, using the Derwent World Patent Index.]

Step 3) Formulating a viable search posed a challenge because there is so much R&D relating to photonics. After an initial trial, with expert feedback to refine what should be included, we settled on a multi-modular strategy: we retrieved abstract records that included 1) certain target terms (e.g., photonic, optoelectronic); we also retrieved records that had 2) both a "photo or optical" term stem AND a term relating to electronics, telecommunications, instrumentation, materials, coatings, or lenses. Given the timeframe of the cluster initiatives of interest, we retrieved data from 1999-2008. We tried using classification information too – SCI Subject Categories – but did not like the results.

Steps 4 & 5) We imported the data into *VantagePoint* and extracted the desired fields for further analysis (e.g., authors, author organizations, key terms, title phrases, publication year, times cited). We cleaned and consolidated these fields.

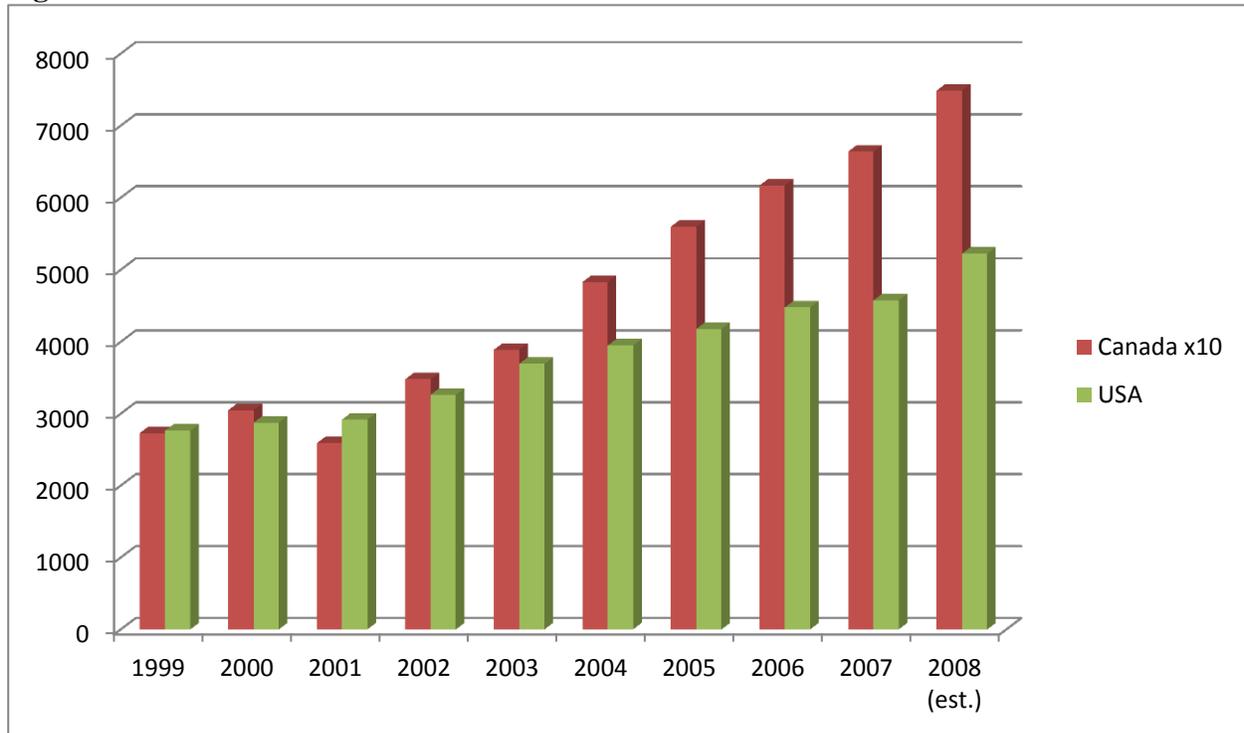
Steps 6 & 7) Analysis and Visualization – we elaborate on these below.

Steps 9 & 10) The Canadian Government group could provide expertise and help interpret.

Step 10) We take advantage of a number of existing macros to generate these results.

Figure 2 shows the overall research trend, comparing Canada to the U.S. Of special interest for the benchmarking comparison is that Canadian photonics research increases relative to the U.S. over this decade.

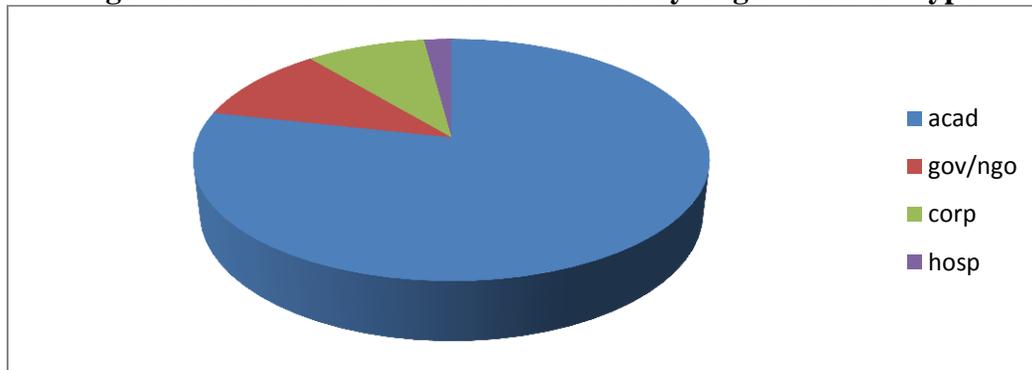
Figure 2. Photonics Research Trend: US and Canada



Note: In that SCI continues indexing additional publications well after the calendar year ends, we multiplied the 2008 tally by 1.2 to get the 2008 (est.) values.

The remaining illustrations give a current profile based on US photonics research in 2008. As mentioned, we find that approximating the percentage of R&D from industry is a telling innovation indicator (albeit one with important nuances). Figure 3 shows the sector breakout. Interestingly, for the benchmarking, the Canadian corporate publication rate (7%) is considerably less than that of the U.S. (12%). We also investigated the extent of cross-sector collaboration. In the U.S., for instance, 11% of the papers with an academic author also had an author with a corporate affiliation. Such evidence of research networking could be important for cluster development.

Figure 3. US Photonics Research in 2008 by Organizational Type



In investigating the effects of the photonics initiatives, we also broke out research publications to enable comparisons by region and by organization. These bibliometric analyses often enable one

to set up quasi-experimental design comparisons to help infer cause and effect. For instance, if one wants to assess the impacts of university U setting up a major new photonics research center in Year X, you can design a set of comparisons (e.g., of research publications and citations received):

- comparing U’s time series before and after Year X (probably setting aside Year X+1 to allow for effects to show up) – does research activity escalate after the center is formed?
- comparing U with similar universities before (hypothesizing similarity) and after (hypothesizing difference)

We also wanted to compare topical emphases between the US and Canada. To do that, we explored information at different levels of detail. For instance we extracted the prominent key words and title phrases prevalent in sets of the articles. We also pulled out information on the Subject Categories into which SCI classifies the publication journals. We then aggregated those into Macro-disciplines and compared the US and Canadian publications for 2007-08. Table 2 shows the leading ones, and we see interesting differences (e.g., relatively more Canadian photonics work in Physics).

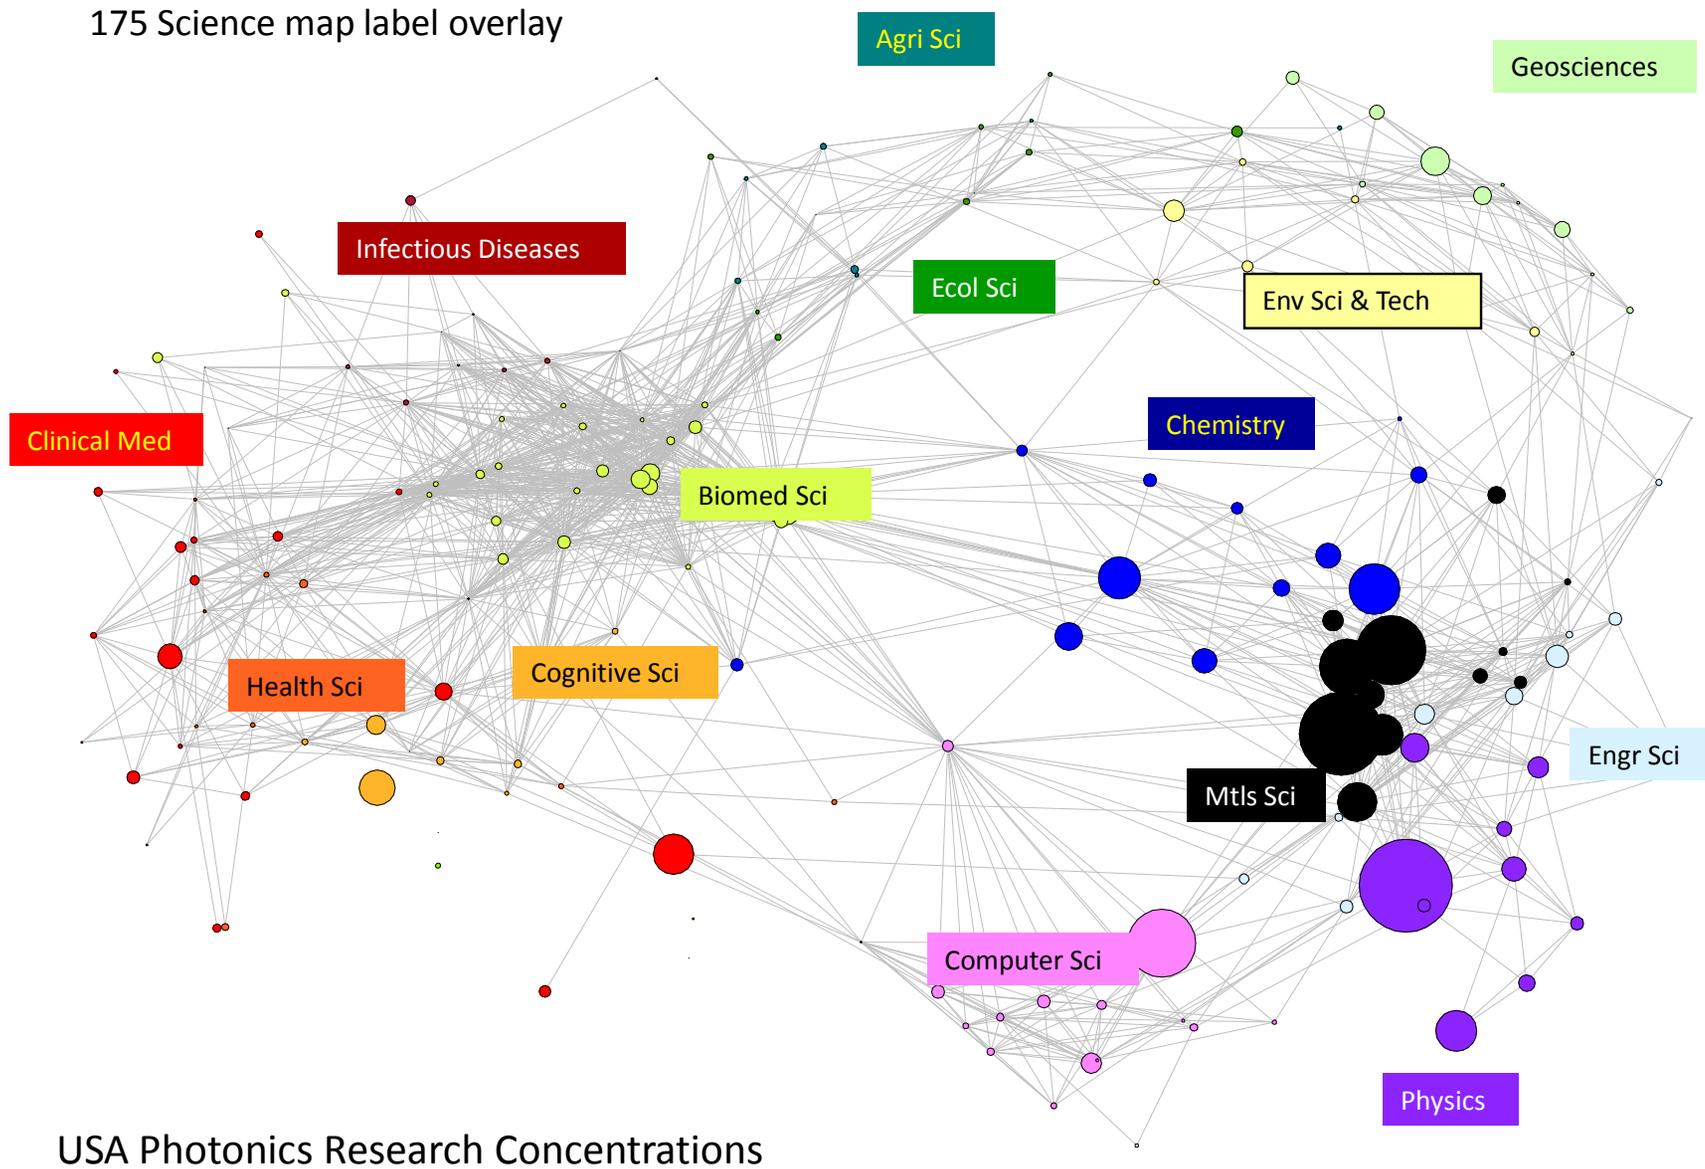
Table 2. National Photonics Publication Macro-Disciplines, SCI, 2007-08

	# Records	8715	1251
#	Macro-Disciplines	USA	Canada
3844	Physics	38.4%	44.9%
3722	Mtls Sci	38.2%	33.3%
1489	Computer Sci	14.4%	19.5%
1467	Chemistry	15.2%	12.8%
843	Clinical Med	8.8%	7.1%

To give a visual feel for the research fields engaged in this photonics work, Figure 5 presents a science overlay map. Against a background of intersecting arcs to represent the 175 Subject Categories covered by SCI, we overlay the photonics publication activity.¹ Here we quickly grasp that photonics research reaches into many fields. So, if the Canadian Government wants to bolster R&D leading to economic returns, funding and other support should not be concentrated in a single field. Furthermore, one may want to take action to promote cross-field collaboration. Figure 5 suggests involving Materials Scientists, Physicists, Chemists and Computer Scientists particularly.

¹ This mapping process categorizes articles indexed in Web of Science according to the journals in which they appear (Rafols and Meyer, forthcoming; Leydesdorff and Rafols, 2009). The Subject Categories are then grouped into “macro-disciplines” using a form of factor analysis (Principal Components Analysis) based on degree of association. Those Macro-Disciplines become the labels in the figure. The photonics research concentrations appear as nodes on these maps. Larger nodes reflect more publications.

Figure 4. US Photonics Concentrations – Overlay on a Base Map of Science



This investigation also sought to compare particular institutions to see how those engaged in research clusters have fared. Tech Mining allows us easily to profile numbers of organizations by applying a macro. We indicate which fields of information we want to see, for which institutions. For fun, Table 3 profiles photonics research activity for three leading US “institutes of technology.”

Table 3. Profiles of 3 Universities: 2007-08 Photonics Research (indexed in SCI)

Organization	Macro-Disciplines	Select Terms	Countries	Collaborating Organizations	h
MIT [273]	Physics [131] Mtls Sci [110] Chemistry [48] Computer Sci [46]	crystal [60] lithograph* [30] nano [29] quantum [25] fiber [24]	Germany [18] UK [16] Japan [14] Singapore [12]	Harvard Univ [15] Caltech [14] Stanford Univ [9]	14
Caltech [253]	Physics [145] Mtls Sci [43] Computer Sci [32] Geosciences [31]	crystal [31] lens [21] quantum [21] nano [17] solar cell [17]	France [42] UK [31] Germany [25] Canada [14]	NASA [27] Univ Arizona [22] UCLA [15] MIT [14]	14
Georgia Tech [194]	Mtls Sci [113] Physics [68] Computer Sci [49] Chemistry [27]	nano [48] crystal [44] lithograph* [25] fiber [23]	China [13] France [11] Germany [9] Taiwan [6]	NEC Labs Amer [12] Emory Univ [10]	13

Note: The last column gives an “h-index” – how many publications received at least this many citations.²

This research profile breaks out for each organization:

- Top Macro-Disciplines reflected by their publications – this can help spot differences in emphases of their photonics research – contrast Caltech’s physics emphasis with Georgia Tech’s materials science.
- Top Key Terms – this may help locate who is researching a particular topic of interest.
- Countries – We see that Caltech seems more internationally networked than is MIT, and that Caltech’s collaboration is heavily with France, the UK, Germany & Canada.
- Collaborating Organizations – MIT shows strong ties with neighboring Harvard; Caltech, with NASA (Jet Propulsion Lab is a joint endeavor); & Georgia Tech, with a firm, NEC. This may give leads on potential interest in partnering.
- The h-index is a measure of the extent of citation of the organization’s publications.²

C. Competitive Technical Intelligence for Dye-Sensitized Solar Cells (DSSCs)

We have analyzed Dye-Sensitized Solar Cells (DSSCs) since 2008 (c.f., Guo et al., 2012; Zhang et al., 2014a, 2014b, 2014c). We searched and retrieved abstract records from two global databases – the Science Citation Index Expanded of Web of Science (WoS) and EI Compendex. Several topical experts have counseled us, and we have updated the datasets over time. Here, we

² “Times Cited” provides a rough measure of paper impact (or quality) – i.e., how often do others reference a given paper. The distribution of citations is highly skewed – i.e., a very few papers receive many citations. We examine citation intensity for the selected organizations’ papers using the “h-index” – the number of papers (N) with at least N citations each. This metric discounts the disproportionate weight of highly cited papers. The h-index was developed by J.E. Hirsch and published in *Proceedings of the National Academy of Sciences of the United States of America* 102 (46): 16569-16572 November 15, 2005.

present selected Tech Mining for a 5784-record data set covering the time span from 1991 (DSSCs introduced in that year by O'Regan and Gratzel) through 2011.

The Term Clumping process (Zhang et al. 2014a) for the DSSCs data has been completed and we have obtained a review by seven experts of the Top 300 consolidated terms, grouped into 8 topical factors from 90,980 phrases and terms. We present the term clumping stepwise results in Table 4, to show the progression from raw to usably “clumped” topical terms & phrases. Table 5 presents the resulting 8 topical factors and their most related items – these become fundamental materials for further semantic TRIZ study and Technology Roadmapping (Zhang et al., 2014b).

Table 4. Term Clumping Stepwise Results

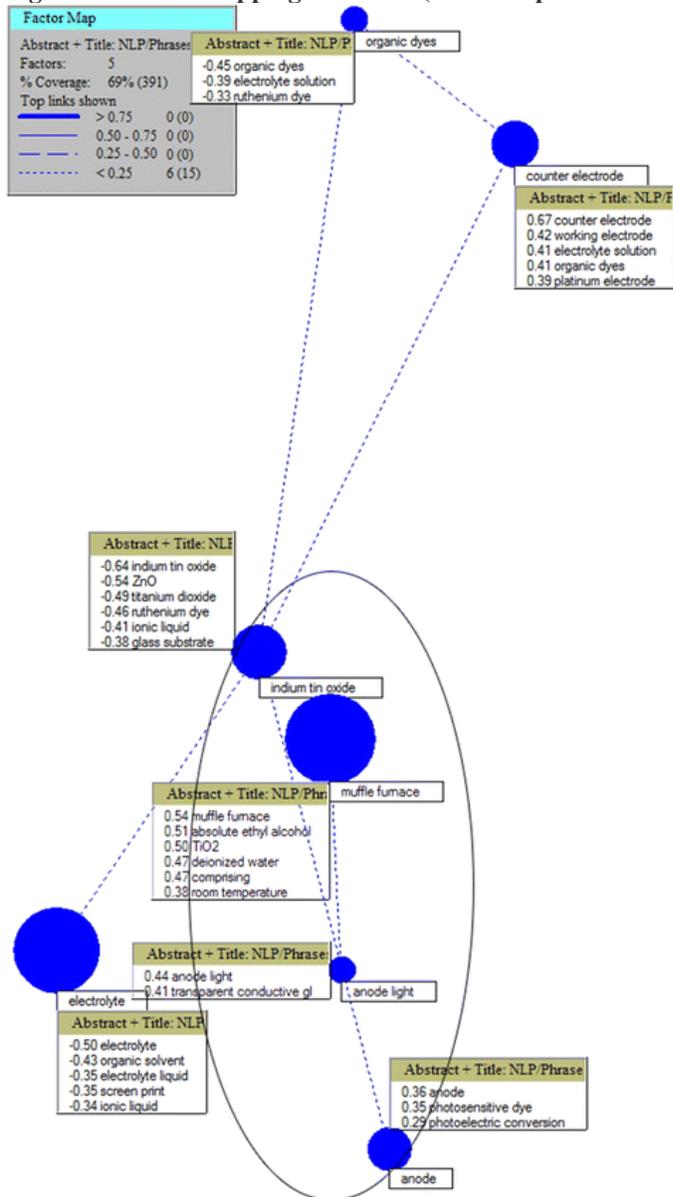
<i>Field selection</i>	<i>Number of Phrases and Terms</i>
Phrases with which we begin	90980
<i>Step a. Applying Thesauri for Common Term Removal</i>	63812
<i>Step b. Fuzzy Matching</i>	53718
<i>Step c. Combining</i>	
Combine_Terms_Network.vpm (Optional)	Not Applied Here
Term_Clustering.vpm	52161 to 37928*
<i>Step d. Pruning</i>	
Remove Single terms**	15299
Fuzzy Matching	14840
<i>Step e. Screening</i>	
Term Frequency Inverse Document Frequency (TFIDF)	14840 (with the Sequence of TFIDF) to 14740***
Combine_Terms_Network.vpm (Optional)	8038
<i>Step f. Clustering</i>	
Term Clustering Analysis (Optional)	3432/2375
Principal Components Analysis (PCA)	11 Topical Clusters

Table 5. Topical Clusters and Related Items

<i>Number</i>	<i>Topical Factors</i>	<i>Related Items</i>	<i>Rank</i>
1	Photoelectric property	Photoelectric property, Hydrothermal method, Higher conversion efficiency	6
2	Sol gel	Sol gel, Sol gel process	3
3	Electron donor	Electron donor, Electron acceptor, Molecular design	7
4	Ruthenium sensitizers	Ruthenium sensitizers, Ruthenium complex, Efficient sensitizers, Absorption spectrum, Charge transfer sensitizer, Red shift, Density functional theory, High molar extinction coefficient	4
5	Electric resistance	Electric resistance, Sheet resistance, Internal resistance	2
6	Modulated photocurrent spectroscopy	Modulated photocurrent spectroscopy, Electron diffusion coefficient, Electron traps, Recombination kinetics, Photo-injected electron, Electron diffusion length	5
7	Photo-induced electron transfer	Photo-induced electron transfer, Electrons transit, Interfacial electron transfer, Rate constant	1

We generated a Factor Map (Fig. 5) via *VantagePoint*, based on term co-occurrence, using principle components analysis (PCA). We asked several Ph.D. students who focus on DSSCs research at the Beijing Institute of Technology for informed suggestions. The four topical factors in the large oval in Fig. 5 are to be consolidated, and the other three factors are retained.

Figure 5. Factor mapping of DSSCs (for the top 50 terms)



Since the traditional ST&I text analysis only focuses on the terms and related topical factors, at this stage, we introduce the nearby verbs and related sentences seeking more complete semantic understanding. We link the terms and sentences with a Subject – Action – Object (SAO) model,

and then, map to Problem and Solution (P&S) patterns (Zhang et al. 2014b). As shown in Figure 6, it is easy to transfer the topical factors/clusters to the “Subjects” or “Objects,” while their nearby “Verbs” are the “Actions” that connect them in the “Related Sentences.” Generally, in the ST&I text, the “Subject/Object” could map to the “Problem,” since most of them evolve with some materials, techniques, or descriptions concerning issues. The “Actions,” and their related semantic structures, can be described as “Functions,” and can be considered as operating on the problems, aiming at “Solutions.” If we track back to TRIZ theory, the “contradiction between object and tool” and the “ideal final result” would be comprehensive explanations for the Problem and its Solution respectively.

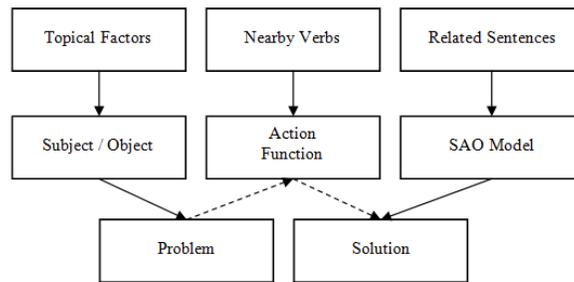


Figure 6. Mapping with Topical Factors, SAO Model and P&S Patterns

As an example, take “conversion efficiency” as a system component for semantic TRIZ study. We process the content of the 5784 records relating to DSSCs with *GoldFire Innovator* software (using titles and abstracts). *GoldFire* retrieved related verbs near the “conversion efficiency” terms and combined them as Subject-Action-Object (SAO) structures -- the same as “P&S” patterns. Table 6 illustrates sample results.

Table 6. P&S Patterns of Semantic TRIZ on “Conversion Efficiency”

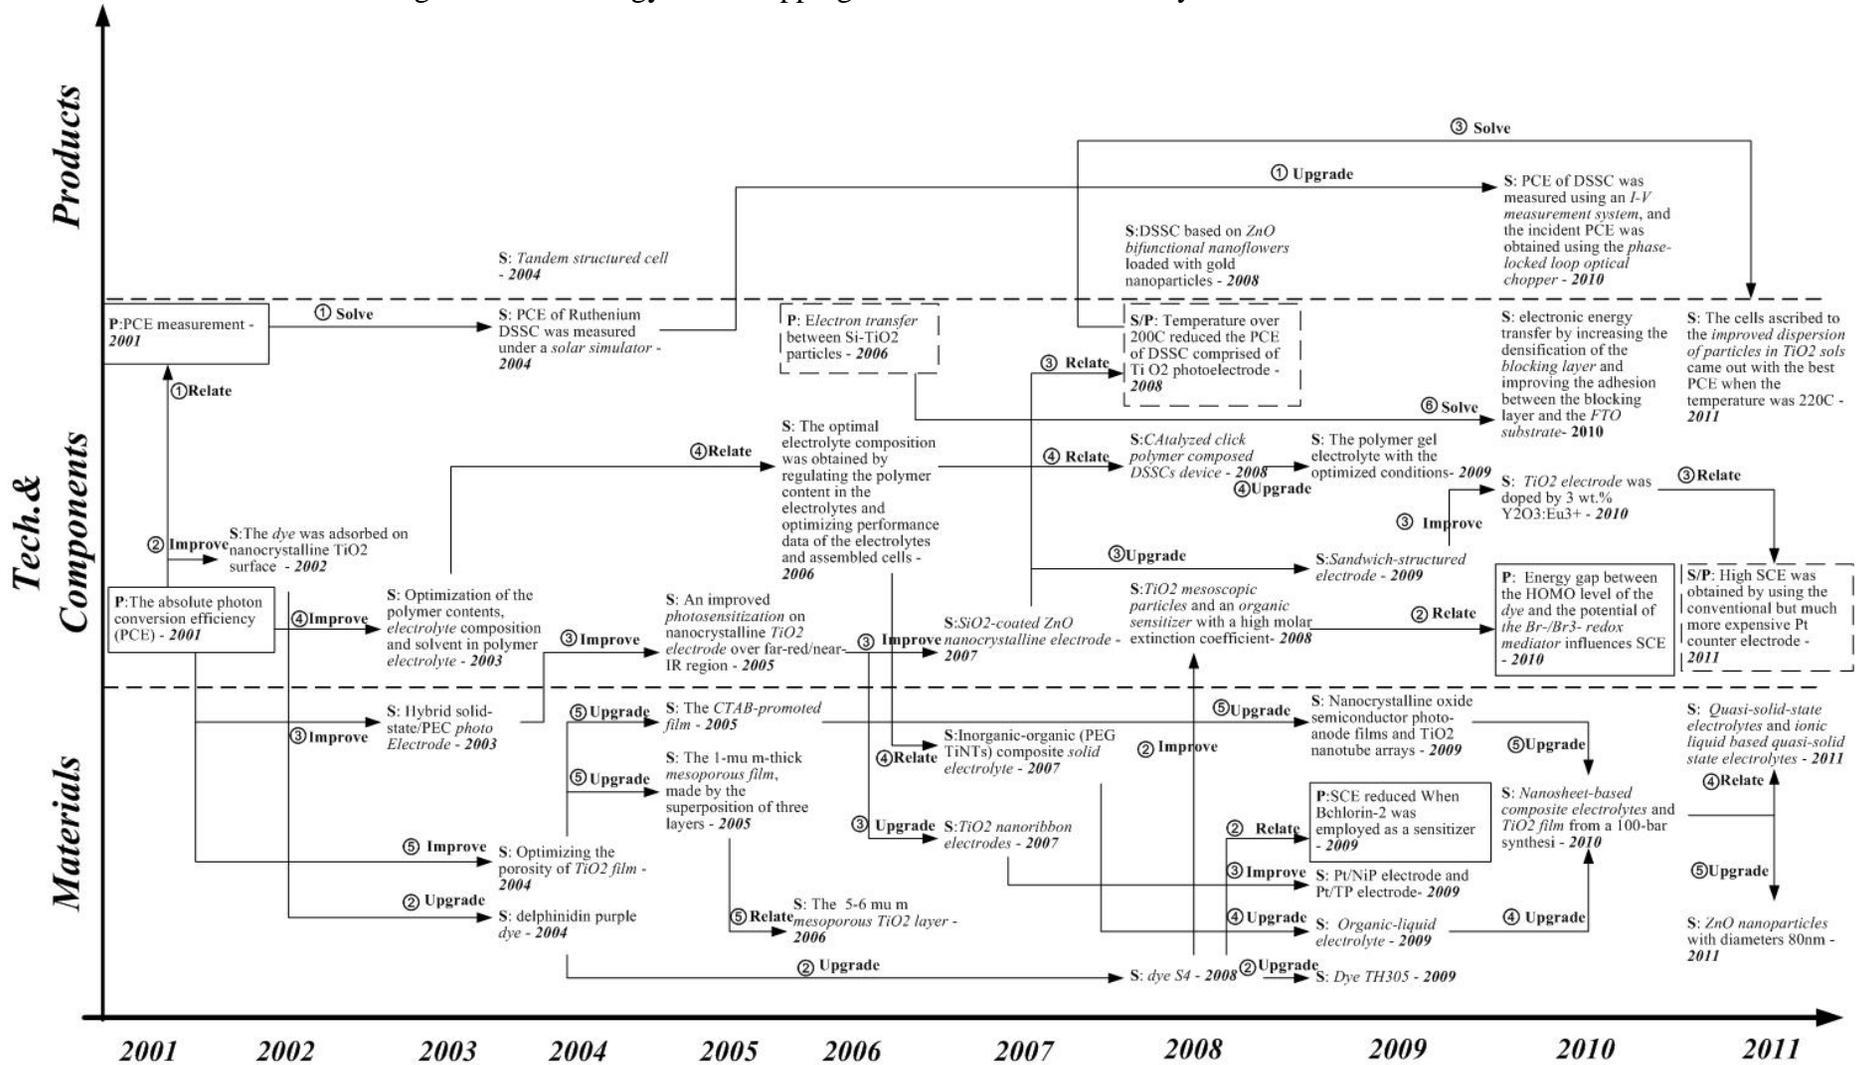
<i>P&S Patterns</i>	<i>Year</i>	<i>Type</i>	<i>Level</i>
<i>The 1-mu m-thick mesoporous film</i> , made by the superposition of three layers, showed <i>enhanced</i> solar conversion efficiency by about 50% compared to that of traditional films of the same thickness made from randomly oriented anatase nanocrystals.	2005	Solution	M
The dye-sensitized solar cells, comprised of TiO2 photo electrode, deposited at substrate temperature of 200C, show maximum photoelectric conversion efficiency ; however, further enhancement of <i>sputtering temperature</i> drastically <i>reduces</i> the efficiency .	2002	Problem	T&C
As a result, the <i>tandem structured cell exhibited</i> higher photocurrent and conversion efficiency than each single DSC mainly caused from its extended spectral response.	2004	Solution	P
The photo-to-electricity energy conversion efficiencies of ruthenium-dye-sensitized solar cells (DSC) <i>are measured</i>	2004	Solution	T&C

Figure 7 presents a resulting TRM for DSSCs (from Zhang et al., 2014b). Key actions in generating this, starting with Tech Mining results, were 1) to identify the topical factors as Subject/Object and the key nearby verbs as solution-oriented Action -- e.g., Solve, Relate, Upgrade, Improve, etc.; 2) to extend the basic SAO model to a variation-based SAO set via a semantic training function; 3) to retrieve the related sentences from the full text via *GoldFire Innovator*; and 4) to refine the results via expert aid.

Referring to Figure 7, we summarize five trends and each of them emphasizes one sub-system of DSSCs:

- 1) Trend 1 is “Dye”; actually, we did not compare the exact value of PCE, but the trend seems to indicate an upgrade of dye and some sensitizer issues;
- 2) Trend 2 is a “Photo Electrode” issue. In the past 10 years, different kinds of materials are used as the electrode, such as TiO₂, ZnO, and Nano materials. Experts also emphasize the importance of TiO₂ here;
- 3) Trend 3 concentrates on “Electrolyte.” One interesting thing that we explored was the changing form of electrolyte -- “solid electrolyte,” “liquid electrolyte,” to “gel electrolyte.” We also found that optimization of polymer content of electrolyte was always a hot topic;
- 4) Trend 4 relates to “film” research. Improvement on the thickness of film and the engagement of nano materials appear to be two main directions;
- 5) Trend 5 seems to be an exception that focuses on the “PCE measurement”; different components or devices were engaged from 2004 to 2010;

Figure 7. Technology RoadMapping for Conversion Efficiency in DSSCs



D. Uses

Effective uses for Tech Mining are still being discovered. This section aims to stimulate your thinking of how this tool set could help meet your information needs.

Competitive Technological Intelligence (“CTI”) has grown aggressively over the recent two decades, as companies and agencies see the need to know what others (companies, universities, government agencies) are doing to develop particular ST&I capabilities. Text mining of results from topical searches in extensive R&D databases is an essential part of CTI. This is becoming ever more critical in an era of “Open Innovation” in which R&D is no longer used exclusively in-house (see References).

Technology Roadmapping (see Appendix) also is growing in use as a vital FTA method. Industry associations [c.f., www.sia-online.org/], government agencies [c.f., <http://www1.eere.energy.gov/industry/aluminum/partnerships.html>], and individual companies seek to lay out prospects for development of families of related technologies (systems and components) over the next several generations. Particularly compelling are efforts by companies, such as Motorola, to map expected developments in technologies together with product development trajectories. Tech Mining can be a key contributor to understanding external competitive environments [c.f., <http://mati.ncms.org/>].

Research evaluation is essentially retrospective in focus, but conveys value to foresight activities by helping to manage R&D programs. Tech Mining has a role to play in comparing and assessing the research outputs of various organizational groups. Certainly, one would not base R&D management or policy solely on counting publications (bibliometrics) or patents, but innovation indicators can provide deeper insights into developmental paths and prospects. Another handy application is to use rapid text mining to inform proposal evaluation processes. Confronted with an unfamiliar proposed research project, a program manager can quickly (and cheaply) identify how much research is ongoing on the topic, what research groups are most prolific, and see how the proposer compares. The manager can also get contact information on potential reviewers knowledgeable on the topic in question.

National foresight studies can use Tech Mining to compare national R&D outputs with those of benchmark nations, or with their country’s economic development targets. This can help identify gaps that may merit attention. Our photonics case exemplifies this sort of analysis.

You may have noted the heavy emphasis on utilization issues in this chapter. We choose to raise these because of our experiences in doing Tech Mining – awareness of these hurdles is a first step toward overcoming them. Furthermore, many of these same issues pertain to futures research in general (c.f., Gordon & Glenn, 2000). We probably all need to work as much on effective presentation of results as on futures research per se.

“Who” should perform Tech Mining? Today, this is probably best done by those with a serious, ongoing stake in generating empirical technology intelligence or foresight. These include patent analysts, CTI specialists, tech forecasters, and information specialists who see the need to provide analyses, not just deliver source materials. Earlier, we suggested an opportunity for

librarians (university, corporate, research institute) to be “research navigators” to provide R&D patterns and trends intelligence to researchers, developers, tech transfer agents, and R&D administrators. Such information can identify technology opportunities, make a case for research initiatives (e.g., in writing proposals), and facilitate innovation management.

Beyond that, there is potential for “anyone” with interests in R&D, ST&I emergence and assessment to “do it yourself.” Desktop software makes the analyses quite doable, but with a learning curve. We can imagine students enriching a term paper with R&D landscaping, or corporate product developers tuning investments via Tech Mining intelligence regarding potential partners and competitors.

We encourage all of you to use the tools at hand to profile the ST&I context. Those tools will be advancing considerably over the coming few years. Check out what profiling/mining capabilities you might have already – e.g., searching *Web of Science*, using the statistical analyses provided by the web interface, allows you to list and count instances by author, keyword, date, etc. Or, even more basically, you can tabulate interesting information right as you search with simple search engines. For instance, if you are searching *INSPEC* for “fuel cells,” you could successively search on “fuel cells and 1995,” “fuel cells and 1996,” etc., to then plot the R&D activity trend. By having the idea of mining these resources, you’ll be amazed at what intelligence you can generate using your available tools.

VI. FRONTIERS OF THE METHOD

Tech Mining is a young method. Most promising is the advance in interpretable innovation indicators that benchmark technological progress. Our *VantagePoint* software was commercially introduced in 2000; improvements in statistical algorithms and representation modes continue at a strong pace. Most promising is the advent of “macros” (*Visual Basic* scripts) that control sequences of steps in both *VantagePoint* and Microsoft software (such as *MS Excel* and *Powerpoint*). Such macros enable automated production of preferred information products in desired form on the desktop.

An intriguing development entails applying a series of text mining analyses to discover inherent cross-relationships not obvious to researchers in individual research domains (see Appendix for Swanson and Kostoff notes). Such Literature-Based Discovery can leverage the expansive repertoire of R&D results to enhance one’s own research and new product development.

As we complete this chapter in early 2015, attention to “big data & analytics” elicits potential to apply Tech Mining tools to multiple data resources and to draw on “crowd sourcing” via the internet to provide expert opinion. Attention to analytics also suggests value in combining multiple modes to present findings via dynamic and interactive visualizations, reinforced with numeric (tabular) information and text interpretations.

APPENDIX

Software resources pertinent to Tech Mining evolve rapidly. Here's one comparative review of many software tools, in terms of doing patent analyses: Yang, Y.Y., Akers, L., Klose, T., and Yang, C.B. (2008), Text Mining and Visualization Tools – Impressions of Emerging Capabilities, *World Patent Information*, Vol. 30 (#4), December, 280-293.

We are seeing increasing involvement of database vendors in provision of analytical tools to use with their data. For instance, Thomson Reuters now offers several of these software aids (Aureka, ClearForest, Thomson Data Analyzer – TDA).

Here's a short list of software especially helpful in extracting intelligence from structured text resources:

Anavist works with Chem Abstracts, and is offered by STN:
www.cas.org/products/anavist/index.html.

Aureka offers particularly attractive visualizations of relationships among patents:
http://www.micropatent.com/0/aureka_online.html.

Matheo software from France: <http://www.matheo-software.com>.

TDA desktop text analytics -- -- <http://thomsonreuters.com/en/products-services/intellectual-property/patent-research-and-analysis/thomson-data-analyzer.html>.

Temis is European-based text analytics: <http://www.temis.com>.

VantagePoint is presented at <http://theVantagePoint.com>.

Selected literature and pertinent Tech Mining approaches:

The Georgia Tech Technology Policy and Assessment Center's approach, called "Technology Opportunities Analysis," is treated on <http://tpac.gatech.edu>.

Tech Mining papers are available at: www.theVantagePoint.com;
<http://www.vpinstitute.org/pages/contactus.html>; and at <http://stip.gatech.edu/>.

M. Brenner, B., Technology intelligence at Air Products: Leveraging analysis and collection techniques, *Competitive Intelligence Magazine*, Vol. 8 (3), 6-19, 2005.

[a great description of how a major company integrates multiple tools to deliver a range of technical intelligence results]

Chesbrough, H.W. *Open Innovation: The New Imperative for Creating and Profiting from Technology*, Harvard Business School, Cambridge, MA (paperback edition), 2006.

Cunningham, S.W., Porter, A.L., and Newman, N.C. Tech Mining Special Issue, *Technology Forecasting and Social Change*, Vol. 73 (8), 2006.

Gordon, T.J., and Glenn, J.C., *Factors Required for Successful Implementation of Futures Research in Decision Making*, Army Environmental Policy Institute, Alexandria, VA, 1999. <http://www.millennium-project.org/millennium/other-studies.html>.

Guo Y, Xu C, Huang L, and Porter A. Empirically informing a technology delivery system model for an emerging technology: illustrated for dye-sensitized solar cells. *R&D Management*, 42(2): 133-149, 2012.

Huston, L., and Sakkab, N., Connect and Develop, *Harvard Business Review*, March, 58-66, 2006. [a compelling success story for Open Innovation at P&G]

Kay, L.; Newman, N.; Youtie, J.; Porter, A.L., and Rafols, I., Patent Overlay Mapping: Visualizing Technological Distance. *Journal of the American Society for Information Science and Technology (JASIST)* 65 (12), 2432-2443, 2014.

Kostoff, R., “Database Tomography” and extensive literature reviews on S&T analyses, including Literature-Based Discovery and Technology roadmapping processes: <http://www.dtic.mil/dtic/search/tr/> [guided search for KOSTOFF retrieves ~50 of his studies]

O’Regan, B., and Grätzel, M., A low-cost, high efficiency solar-cell based on dye-sensitized colloidal TiO₂ films, *Nature*, 353(6346) ,737–740, 1991.

Porter, A.L., QTIP: Quick Technology Intelligence Processes, *Technological Forecasting and Social Change*, Vol. 72, No. 9, 1070-1081, 2005.

Porter, A.L., and Cunningham, S.W., *Tech Mining: Exploiting New Technologies for Competitive Advantage*, New York, John Wiley, 2005.

Porter, A.L., Kongthon, A., and Lu, J-C., Research Profiling: Improving the Literature Review, *Scientometrics*, 53, 351-370, 2002.

van Raan, A., and colleagues at the University of Leiden – bibliometric mapping: <http://sahara.fsw.leidenuniv.nl/cwts/>

Robinson D K R, Huang L, Guo Y, and Porter A., Forecasting Innovation Pathways (FIP) for new and emerging science and technologies. *Technological Forecasting and Social Change*, 80(2): 267-285, 2013.

Roper, A.T., Cunningham, S.W., Porter, A.L., Mason, T.W., Rossini, F.A., and Banks, J., *Forecasting and Management of Technology*, 2d edition, New York: John Wiley, 2011.

Watts, R.J., and Porter, A.L., Innovation Forecasting, *Technological Forecasting and Social Change*, 56, 25-47, 1997.

Leydesdorff, L., and Rafols, I., A Global Map of Science Based on the ISI Subject Categories. *Journal of the American Society for Information Science and Technology*, 60(2), 348-362, 2009.

Rafols, i., and Meyer, M., Diversity and network coherence as indicators of interdisciplinarity: case studies in bionanoscience. *Scientometrics*, 82, 263-287, 2009.

Zhang Y, Guo Y, Wang X, Zhu D, Porter A. A hybrid visualisation model for technology roadmapping: bibliometrics, qualitative methodology and empirical study. *Technology Analysis & Strategic Management*, 2013, 25(6): 707-724.

Zhang Y, Porter A L, Hu Z, Guo Y, and Newman N., “Term clumping” for technical intelligence: A case study on dye-sensitized solar cells. *Technological Forecasting and Social Change*, 85: 26-39, 2014a.

Zhang Y, Zhou X, Porter A L, and Gomila J. M. V., How to combine term clumping and technology roadmapping for newly emerging science & technology competitive intelligence: “problem & solution” pattern based semantic TRIZ tool and case study. *Scientometrics*, 101(2), 1375-1389, 2014b.

Zhang, Y., Zhou, X., Porter, A.L., Gomila, J.M.V., and Yan, A., Triple Helix Innovation in China’s Dye-Sensitized Solar Cell Industry: Hybrid Methods with Semantic TRIZ and Technology Roadmapping, *Scientometrics*, 99:1, 55-75, 2014c.