Why Don't Technology Managers Want Our Knowledge?

Alan Porter, Director of the Technology Policy and Assessment Center, Professor in the School of Industrial & Systems Engineering, and the School of Public Policy, Georgia Tech

Nils Newman, President, IISC, Inc., Atlanta

Robert Watts, Senior Staff , U.S. Army Tank & Automotive Command, Warren, MI, USA

Prof. Donghua Zhu, Director of the Institute of Forecasting & Development of the Hefei University of Technology, Hefei, China

Cherie Courseault, PhD Candiate, Industrial & Systems Engineering, Georgia Tech

Webb Myers, M.S. student, Public Policy, Georgia Tech

Elmer Yglesias, PhD Candidate, Public Policy, Georgia Tech

Abstract

Rich information resources are amenable to mining for competitive technological intelligence. Yet managers presently underutilize this resource. We explore why this is, propose solutions, and forecast a transition to widespread acceptance by 2010.

Keywords: competitive technological intelligence, innovation indicators, bibliometrics, knowledge discovery in databases (KDD), information products, information utilization, textual data mining, technology opportunities analysis

Introduction

Strategic management demands significant attention be directed to technology as a key competitive force. For most organizations this requires increasing focus on *external* technological developments. Constricting internal R&D budgets, recognition of collaborative advantages, and need to detect competitive threats all demand effective information concerning emerging technologies.

Fortuitously, electronic information resources provide a wonderful resource to mine for such competitive technological intelligence. In particular, informative worldwide web sites and huge databases compiled on various domains of science and technology are becoming widely available. The "pull" of management's need for technological information must be linked with the "push" of these information resources. To help make this link, powerful suites of search, analysis, and visualization tools are coming available. We'll consider these under the general label of "Knowledge Discovery in Databases" (KDD) – a rapidly evolving discipline that blends mathematics, statistics, and artificial intelligence to tease knowledge out of warehoused data [1,2,3]. Such analyses are also known as bibliometrics or textual data mining [4,5].

A cloud intrudes on this sunny picture of increasing "knowledge generation" capabilities. Technology managers are slow to adopt these new capabilities. This paper probes reasons for the resistance and offers ways to improve the analytical approaches to boost utilization.

We focus on experiences in applying one such suite of software tools called the "Technology Opportunities Analysis System" (TOAS) [http://searchtechnology.com]. A team including the authors and colleagues at Georgia Tech, Search Technology, Inc., and IISC has been developing TOAS since 1993 [6]. Such software helps analyze information retrieved from large electronic databases to serve technology management [7]. TOAS extracts information about particular emerging technologies [8] through a process of:

- search & retrieval from publicly accessible abstract databases (e.g., *Engineering Index, INSPEC, U.S. Patents, Business Index*)
- profiling & analysis of the search sets (typically 200 to 20,000 abstracts)
- representation & interpretation of derived knowledge.

Particular knowledge applications can include:

- comparative benchmarking of alternative technologies
- evaluation of competing R&D proposals and programs
- competitive technological intelligence (what technologies is Company C actively pursuing?)
- technology forecasting & assessment (what developments are likely?)

Potential users range across practicing scientists & engineers, information professionals (e.g., technology licensing specialists), product development managers, and strategic planners & policy makers. Issues in assessing knowledge for this 'management of technology' arena are elaborated elsewhere [9].

However, putting such forms of knowledge to use poses real challenges, notably:

- What is the nature of the information products generated?
- What factors affect acceptance?
- How can we produce the right information products when needed?

The following sections address these challenges; the Conclusions section offers suggestions on persuading management to use these sorts of derived knowledge more fully.

What is the nature of the Information Products generated?

Turning to the first of the three challenges, what is special about this knowledge? Several facets stand out. A recent *Harvard Business Review* article asked "What's your strategy for managing knowledge?" [10]. The authors differentiate codified knowledge (reusable knowledge stored in databases) from personalized knowledge (highly customized solutions delivered through person-to-person communication). They suggest treating each type of knowledge differently. However, we extract information from collections of research and development (R&D) abstracts to meet distinct technology management needs. This represents a hybrid form of information exploitation -- "customization of codified knowledge, on demand." This poses peculiar problems.
KDD, using software such as our TOAS, exploits the vast resources of large information collections (e.g., *MEDLINE* presently contains some 12,000,000 abstracts of articles or conference papers – a major portion of the world's medical R&D in a single compilation!). Widely accessible search and retrieval processes put this information "at one's fingertips" in minutes on almost any medical

science & technology topic [c.f., http://www.ncbi.nlm.nih.gov/PubMed]. But technology managers don't want to "know about" who's doing research in general, they want intelligence that helps them reach particular decisions. That requires carefully targeted analyses deriving in part from these vast resource sets.

Second, KDD constitutes a *new form* of knowledge. It derives from patterns emergent in large data compilations. It relies on statistical analyses and text mining that are not likely to be familiar to the target strategic managers and technical professionals. Note that the "knowledge" derives from the information resources per se – not from experts digesting the raw information to generate the new knowledge. This is unfamiliar ground to most technology managers and professionals.

Third, TOAS generates "information products" focused on a particular domain -- technological change. We seek empirical measures to assess technological maturation, contextual influences, and market opportunities. For instance, we compute the percentage of publishing (and/or patenting) in a research domain by industry, as opposed to that by academia or government, as a commercial-readiness indicator [c.f., TOAS analysis of our own 'KDD' field: http://tpac.gatech.edu]. Such derived knowledge is intended to support decision processes better than the raw information can. In other words, the knowledge uses conceptions of how the technological innovation process works to guide the text mining [11].

Fourth, this derived information can be presented in quite different ways. We can consolidate the content into one comprehensive report or "packetize" into sequential issue analyses. We can vary the media used, emphasizing *text* interpretations, *numerical* presentations (e.g., lists of "Top 10" nations, organizations, authors, or journals), and/or *graphical* representations (e.g., "technology maps"). We have experimented with providing reports packaged with CD's containing the topical abstracts and a simplified version of the TOAS software so users can hyperlink from a particular result to check out the underlying information (e.g., to read, say, IBM's abstracts on 'intelligent agents' on the spot).

What factors affect acceptance?

We have learned from experience that acceptance of non-traditional information products and tools is not automatic. Knowledge users pose manifold requirements [12].

Over the past six years we have used TOAS and its predecessor "TOA Knowbot" software to perform, perhaps, 100 analyses for different governmental and business users [c.f., "Technology Opportunities Analysis" at http://tpac.gatech.edu]. Since early 1999, TOAS has become available for use by those users themselves. The experiences have been sobering – these new information products are not readily adopted in support of technology management. For instance, we recently generated three topical profiles for a leading military medical group. We did a number of things "right" (we thought) -- framed the studies based on input from their technical professionals in each area, provided interim results for corrective feedback, and delivered the reports together with the raw abstracts on CD for them to pursue what most met their needs. We also presented two briefings to explain the value of the information and how to get at it. To the best of our knowledge, the reports just gathered dust.

In one of our success stories, we profiled R&D on a particular technology for the technology foresight group of a leading consumer corporation. They became repeat customers and we generated several more such reports. Only then did they come to understand the nature of TOAS-derived information well enough to tell us – whoa! -- totally refocus future reports to answer *our* five key technology management questions. Over time our reports became more useful – sometimes still in unintended ways. For instance, in one case, trying to frame a TOAS study served to identify and bring together diverse corporate interests with a stake in that particular technology.

This corporate group went on to obtain TOAS for their own use, rather than having us perform analyses for them. But that brought additional challenges. The senior staff doesn't engage the data or the analyses themselves. They have an information specialist search the databases and they assign interns to become proficient in TOAS analyses. We'll return to concerns about conveying information through such intermediaries.

Ron Kostoff finds a paucity of successful applications of "management science," including text mining, in technology management [4]. Bibliometrics meet resistance in R&D evaluation processes. Focusing on textual data mining, he finds the analysts badly decoupled from the target users [4, 13].

With support from the National Science Foundation (NSF) [14], we have reviewed experiences and experimented with alternative ways to provide information to business and government. In collaboration with several companies and two government agencies, we have explored specific technology management information needs vis-à-vis KDD capabilities. We have identified a number of issues.

First, target users need the right information for the task at hand. Too often, we have found that our cleverly crafted reports miss the mark. There appear to be deeply rooted preferences among knowledge workers for either *reports* or *answers to questions*. Some organizations emphasize reports of various kinds; others do not. Sometimes, it takes pilot testing to craft information into a useful form. For instance, we have recently been working with a National Institute for Occupational Safety and Health (NIOSH) lab on use of TOAS in support of Government Performance Review Act (GPRA) requirements [13]. In this case we found that the subject domain experts called upon for peer review simply had no clue how they might use our information products. We are currently working to recast these as answers to questions, including:

- How timely is this lab's research?
- How productive are the lab's researchers compared to peers?
- How do the lab's topical emphases map with external researchers (e.g., any gaps of note)?
- Who are the leading external researchers in this area (e.g., as potential reviewers)?
- Who is citing the lab's research (e.g., are labor and industry using the results)?

In technology management, reports vary in emphases, including alerting to emerging technology opportunities, benchmarking organizational performance, or forecasting technological prospects. However, we have learned that "standard" reports rarely contribute to imminent decisions – what busy professionals and managers seem to most need. For instance, we have performed a study concerning which, if any, alternative technologies could compete with a proposed new structural material. The purpose was to help the client decide whether to build a production plant. Posing the questions so TOAS would provide effective answers demanded ongoing communication between the would-be user and the analyst to get this right.

Second, users need the information on a timely basis. A survey of 26 technology professionals and managers collaborating in our NSF project found them typically needing results fast:

- within a day (21%)
- within a week (45%)
- within a month (24%)

But our analyses have typically taken a month or more. We got the message that, to be useful, we need to be able to provide results on demand.

Third, most would-be users have a difficult time accepting KDD "knowledge." The ready availability of electronic databases is relatively new and unfamiliar. Coming to accept knowledge derived therein by data mining is takes experience testing its validity and utility. TOAS-based information products do not have instant credibility. Managers prefer to have experts provide the requisite knowledge. Experts constitute a more familiar font of knowledge. They also offer "blamability" advantages should decisions go awry.

Combining KDD information products with expert opinion is a good response. However, this is easier said than done. In some cases, clients don't want outsiders (experts) to become aware they are probing particular technologies (competitive technological intelligence works in two directions!). Also, analysts adept at using textual data mining tools may not be so comfortable pursuing expert opinion.

Fourth, we believe that involving the user in the KDD process really facilitates acceptance [4]. However, this often proves hard going. Scientists and engineers tend to be uncomfortable with profiling entire R&D domains. But, if they can be engaged, they can develop enriched perspective on how their research links with others', fostering new contacts and collaborations. For managers, digging into these information products calls upon largely unfamiliar skill sets. They tend to designate a junior associate or an information specialist to learn TOAS. This imposes an intermediary in the process, making acceptance of the information products more difficult.

In working with, perhaps, forty private and public organizations, we have uncovered a special fragility. It seems that persons with expertise in dealing with these emerging electronic information resources and KDD processes are still relatively rare. Precious learning on how to adapt the information products to the organization's decision needs can be undone by a single reassignment, because knowledgeable substitutes are not readily available. This supports Kostoff's suggestion that KDD and related analytical tools must become an integral part of the business decision processes to attain high and enduring value [4, 13].

How Can We Produce The Right Information Products When Needed?

Since 1993, we have built up the TOAS tool suite through an iterative process of using the available generation of software to perform analyses, in turn, leading to requests to our programmers to enhance its capabilities. This has resulted in Windows-based software that provides basic list processing, co-occurrence-based matrices, natural language processing to parse texts into analyzable phrases, fuzzy matching for data cleaning, thesaurus generation and enhancement via experience, principal components analyses, and representational capabilities (combining multi-dimensional scaling with a path erasing algorithm). TOAS is thus, first of all, general purpose "KDD" or text mining software.

Combining conceptual understanding (about technological innovation processes) with empirical experiences, we have evolved a set of specialized algorithms that greatly facilitate generation of specific technology management outputs. Namely, we can semi-automatically generate six different technology maps, two innovation indicators, term clusters, four trend model fits, and two versions of a proprietary principal components decomposition. Focusing tool enhancement on one domain (technology management) enables quick generation of particular management-oriented information products. For example, the technology decomposition reflected in [11] took months to perform; the constituent analyses could now be done in a couple of days.

The need to enable lightly trained analysts (e.g., students) to generate particular useful information products has prompted us to pursue even more detailed "templates" combining tacit analytical knowledge and familiar resources [15,16]. We are now programming "scripts" that capture sequences of analytical steps for quick production of rather sophisticated information products. An Office of Naval Research research intelligence group at Edison House (London) is working to identify particularly effective TOAS results (e.g., a particular table of leading researchers on a topic with their affiliation and topical emphases). Then, have this scripted so that Naval Reservists, within a two-week period, can readily generate a report consisting of a number of such "analytical chunks" on other technologies.

Thus, by elaborating upon the TOAS software through domain-specific algorithms and explicit processing scripts, we can quickly generate a number of management-oriented information products on demand. This facilitates an iterative sequence wherein a busy manager or technical professional reviews a particular information product to suggest additional next analyses.

Conclusions

To summarize, we have noted a new generation of "information products" that exploit large quantities of text information on particular topics (in our case, emerging technology topics). We have not detailed the nature of these products here [c.f., http://tpac.gatech.edu; http://www.dtic.mil/dtic/kostoff], but rather have considered their nature vis-à-vis management needs. We next noted the generally poor acceptance of these new information products by technology professionals and managers [17]. Then we pointed to changes we are making to our TOAS tools to enable faster and easier generation of particular forms of information product. We conclude by offering suggestions to enhance the utility of such derived knowledge [18].

1. As analysts, develop strong rapport with the target users to *understand* [19]
- the users' information needs

- their familiarity with different types of information resources and analyses
- the credibility of different information products in their eyes
- the time frame in which they need results.

2. Understand the users' organizational culture.
- Do they prefer answers to specific *questions*? If so, provide those rather than extensive profiles that show off the analysts' KDD capabilities.
- Do they favor formal *reports*? If so, know and use their familiar form. Make the report suit their needs, rather than generating what we (analysts) find easy and interesting.

3. Consider delivering the information products in properly digestible amounts. In particular we have had some success in *packetizing* the information. Rather than providing an overwhelming amount of information all at once, work with the target user(s) to determine a suitable sequence of smaller, sharply issue-focused information chunks. Such *sequencing* also offers the prospect of increasing the user's understanding of the information generation process and ownership in that process. It can begin by answering a straightforward question quickly (e.g., that day), followed up in pursuit of additional issues that surface for the user as his/her understanding of the issue evolves.

4. Adapt the *information form* to the user's needs. At the most basic level, ascertain preferences
- among numerical, text, and graphic emphases
- between verbal and written reports, and
- look for better modes – e.g., MS PowerPoint slide sets (for ready reuse)

However, don't limit just to familiar forms. In our modest survey of technology professionals and managers collaborating in our NSF project, 73% personally do some mining of external information resources these days and 42% were open to trying explorable, interactive reporting (e.g., web-based or CD-based hyperlinked elements).

5. Document the *value added* by use of information mining (KDD, bibliometrics, etc.). For instance, the capability of making a case by providing a relatively comprehensive perspective on the world's (or "Company C's") R&D in a given domain may provide a significant edge. Sometimes, the quantitative results of information mining may gain an advantage in justifying a given program. A few successes in beating internal budgetary competitors may convert many to the potential of these new types of analyses.

6. Integrate various forms of derived knowledge to meet the user's needs. As mentioned, analysts prefer to work in their comfort zone – e.g., textual data miners don't particularly welcome having to solicit and interpret expert opinion.

7. Address the *fragility* issue. Strive to develop organizational capabilities able to sustain interest in information exploitation for the long haul. This may entail a combination of consultant analyses, identification of analytical capabilities within the organization, training, and evaluation of what does and does not work well. At the most basic level, we have learned that one needs a dedicated user set (maybe even a champion) just to obtain access to the information resources and sustain them. wazzu
When all is said and done, will these emerging information products be accepted by all technology professionals and managers? No! However, we offer a bold projection. By the year 2010, the

exploitation of electronic information resources concerning emerging technologies will be effectively universal among technology managers and professionals. It just won't be acceptable for a manager to justify selection of a few projects from many proposals based on intuition. Engineers won't begin a new project without ascertaining the worldwide developmental status on the technology, competitor technologies, and supporting technologies. Competitive technological intelligence, deriving in part from information mining, will be the norm. Strategic technology planning (i.e., technology foresight, technology roadmapping) will build on the use of such tools (though not exclusively).

Is the previous paragraph self-consistent? On the one hand it states that all technology managers will not accept the use of these new information products. On the other, it claims that their use will be almost universal. The apparent inconsistency will resolve itself. Again a prediction – those technology managers unwilling to use these knowledge resources will be phased out. The required skill set for technology managers and professionals will come to require capability to use these information products, implying some facility in using the information analysis tools. Those too stiff or lazy to learn them, will be replaced.

We offer a few bases for this nasty prediction. First, professional norms change fast in the information age in which we live. In the 1980's, Alan Porter was involved in two studies for the Industrial Research Institute on the adoption of computers by industrial scientists and engineers. We were able to project and then observe a remarkable transition. In 1980, on the order of 10% of industrial scientists and engineers regularly used a computer; by 1990, this was pushing 85% [20,21]. In the 1990's we have seen a similar amazing transition to the networking of those computers and electronic materials (the Internet and databases).

In concert with these information technology driven changes have come new educational generations. We can reflect back on generations of students shifting from learning slide rules to learning calculators to learning PC's. Similarly, we assert that those graduating college prior to 1990 rarely learned electronic searching skills (e.g., Boolean database searching, web browsing). In contrast, recent technical and managerial graduates, almost surely have intimate knowledge of how to access electronic information resources. So, during the coming decade we expect the diffusion of electronic searching capability in business to make obsolete those lacking it. Moreover, it will enable the technology managers and professionals to either "do it themselves" or, more typically, to be knowledgeable consumers of KDD-based information products.

Institutional barriers to electronic information access remain considerable. Many firms' security concerns about guarding their internal information spill over to block access to external information – namely, firewalls of various sorts. That is compounded by Internet phobia – concerns that employees will waste time wallowing on the Internet. But the signals that this is changing are strong. Leading "information age" companies like Buckman Industries open two-way information channels [22]. Internet commerce opportunities are exploding, generating interest and movement toward resolution of information security concerns. And, database providers and gateway companies (e.g., Dialog) are increasingly open to negotiating unlimited access contracts as opposed to the former charge-per-abstract model.

The engagement of dominant companies also supports the projection that information mining products will be widely used by technical professionals and managers. SAS offers a data mining suite [http://www.sas.com/software/data_mining]. In 1998, IBM introduced its *Intelligent Miner* software suite to enable others to generate intelligent analytical tools. While ourselves and other domain-specialized players have been developing such tools for a decade or thereabouts, we believe IBM's introduction will initiate many tools over the coming 5-10 years. In addition, in our analysis of the KDD field, we noted an intriguing datum. Key members of the leading research group in KDD suddenly moved from Pasadena (Caltech) to Redmond, WA. And when Microsoft enters a software domain, we know it's going mainstream.

References

[1] Carlisle, J.P., Cunningham, S.W., Nayak, A., Porter, A.L., and Sprague, Jr., R.H., Related problems of knowledge discovery, *Hawaii International Conference on Systems Sciences (HICSS-32), Proceedings on CD – Modeling Technologies and Intelligent Systems Track; Data Mining and Knowledge Discovery Mini-track* (1999).

[2] O'Leary, D.E., "Enterprise Knowledge Management," *Compute. 31 (3),* (March, 1998) 54-61.

[3] Watts, R.J., Porter, A.L., Cunningham, S.W., and Zhu, D., "TOAS Intelligence Mining: Analysis of Natural Language Processing and Computational Linguistics," in J. Komorowski and J Zytkow (eds.), *Principles of Data Mining and Knowledge Discovery (First European Symposium, PKDD'97, Trondheim, Norway),* Springer, (1997) 323-335.

[4] Kostoff, R.N., and Geisler, E. Strategic Management and Implementation of Textual Data Mining in Government Organizations, *Technology Analysis and Strategic Management* (to appear).

[5] Ahonen, H., Heinonen, O., Klemettinen, M., Verkamo, A.I., Applying data mining techniques for descriptive phrase extraction in digital document collections, *Proceedings, IEEE International Forum on Research and Technology: Advances in Digital Libraries (ADL'98)* (1998) 2-11.

[6] Porter, A.L., and Detampel, M.J.: Technology Opportunities Analysis, *Technological Forecasting and Social Change*49 (2) (1995) 237-255.

[7] Feldman, R., Fresko, M., Hirsh, H., Aumann, Y., Liphstat, O., Schler, Y., and Rajman, A., and Reimer, U., Knowledge management: A text mining approach, *Practical Aspects of Knowledge Management (PAKM98)* (1998) 9/1-10.

[8] Watts, R.J., Porter, A.L., Courseault, C., "Functional Analysis: Deriving Systems Knowledge from Bibliographic Information Resources," *Information, Knowledge, Systems Management* (to appear).

[9] Porter, A.L., Roper, A.T., Mason, T.W., Rossini, F.A., and Banks, J.: *Forecasting and Management of Technology*. Wiley, New York, NY (1991).

[10] Hansen, M.T., Nohria, N., and Tierney, T., "What's your strategy for managing knowledge, *Harvard Business Review 77,* (March-April, 1999) 106.

[11]Watts, R.J., and Porter, A.L.: Innovation Forecasting, *Technological Forecasting and Social Change* 56, 25-47 (1997).

[12] Simpson, J., Kingston, J., and Molony, N., Internet-based decision support for evidence-based medicine, *Knowledge-Based Systems 12* (1999) 247-255.

[13] Kostoff, R.N., Network-centric Peer Review (submitted for publication).

[14] Porter, A.L., Carlisle, J.P., Watts, R.J. *Mining Bibliographic Information on Emerging Technologies,* U.S. National Science Foundation, Management of Technological Innovation Project (DMI-9872482), 1998-2000.

[15] Kerber, R., Beck, H., Anand, T., Smart, B., Agrawal, R., and Stolorz, P., Active templates: Comprehensive support for the knowledge discovery process, *Proceedings, Fourth International Conference on Knowledge Discovery and Data Mining* (1998) 244-248.

[16] Staab, S. and Schnurr, H-P. Knowledge and Business Processes: Approaching an Integration. In: OM '99 - Proceedings of the international Workshop on Knowledge Management and Organizational Memory (IJCAI-99). Stockholm, Sweden (1999).

[17] Porter, A.L., Myers, W. and Courseault, C. Engineering and Technology Management: Intelligent or Otherwise?, Portland International Conference on Management of Engineering and Technology (PICMET) (1999).

[18] Taylor, R.M., Steps on the path to knowledge, *Knowledge Management Information Management Event Proceedings* (1999) 33-36.

[19] Nayak, A., Knowledge Discovery from Databases for Technology Innovation Management – A User Centered Approach to Process Design, Portland International Conference on Management of Engineering and Technology (PICMET) (1999).

[20] Porter, A.L., and Rossini, F.A., Current and Future Uses of the Computer: Industrial R & D in the United States, *R & D Management*16 (1986) 279-289.

[21] Rossini, F.A., Porter, A.L., Jacobs, C.C., and Abraham, D.S., "Trends in Computer Use in Industrial R & D, *Research/Technology Management* 31 (Sep-Oct, 1988) 36-41.

[22] Buckman, R.H., Knowledge sharing at Buckman Labs, *Journal of Business Strategy* 19 (Jan-Feb, 1998), 10-16.

Title: Why Don’t Technology Managers Want Our Knowledge?

Alan Porter, Director of the Technology Policy and Assessment Center, Professor in the School of Industrial & Systems Engineering, and the School of Public Policy, Georgia Tech, Atlanta GA, USA.
Phone: 404-894-2330
fax: 404-894-2301
e-mail: alan.porter@isye.gatech.edu

Nils Newman, President, IISC, Inc., Atlanta

Robert Watts, Senior Staff , U.S. Army Tank & Automotive Command, Warren, MI, USA

Prof. Donghua Zhu, Director of the Institute of Forecasting & Development of the Hefei University of Technology, Hefei, China

Cherie Courseault, PhD Candiate, Industrial & Systems Engineering, Georgia Tech

Webb Myers, M.S. student, Public Policy, Georgia Tech

Elmer Yglesias,PhD Candidate, Public Policy, Georgia Tech