

Business Intelligence: The Next Frontier for Information Systems Research?

Panel at WITS-2001¹

Abraham Bernstein

NYU – Stern School of Business
44 West 4th Street, Room 9-76
New York, NY 10012, USA
(212)998-0803

bernstein@stern.nyu.edu

<http://www.stern.nyu.edu/~abernste>

Benjamin Grosof

MIT – Sloan School of Management
50 Memorial Drive, Room E53-317,
Cambridge, MA 02142, USA
(617)253-8694

bgrosof@mit.edu

<http://www.mit.edu/~bgrosof>

Foster Provost

NYU – Stern School of Business
44 West 4th Street, Room 9-71
New York, NY 10012, USA
(212)998-0806

fprovost@stern.nyu.edu

<http://www.stern.nyu.edu/~fprovost>

Abstract

Business Intelligence is the use of high-level software intelligence for business applications. Advances over the past decade, in processing power, connectivity, and intelligent technologies have engendered a resurgence of interest in the use of software intelligence for business. Increasingly, customers interface directly with a business's systems where the level of effective intelligence affects them directly ("businesses can't afford for these systems to be stupid"). Business systems also increasingly are interacting automatically with other business systems, which also raises issues of competitive and cooperative intelligence. Additionally, these direct interactions (and the associated collection and use of data) lead to novel opportunities for information systems.

How is the Information Systems research community responding? The field of Information Systems is in a unique position to capitalize on the general interest in business intelligence. In particular, *IS* research spans *both business* expertise and expertise in the *technologies*. Such breadth and synergy are (or is it only that they seem?) relatively rare in the other research communities that impinge on Business Intelligence.

This panel will discuss the importance of Business Intelligence applications for business today and the implications for Information Systems research. Is the spanning of business and technology truly unique? Or are we fooling ourselves? Will we be "competing" with researchers from other disciplines? If so, is that good or bad? If the 90s was the decade of economically oriented IS research, will the 00s be the decade of technologically oriented IS research?

We propose a panel discussion, comprising several senior members of the technologically oriented IS research community, to discuss these and related matters.

Keywords

Business Intelligence, IS research, technology, future directions

¹ Panel Description. In: Proceedings of the Workshop on Information Technologies and Systems (WITS '01) (<http://www.busi.mun.ca/parsons/wits2001/>), held New Orleans, LA, USA, Dec. 15-16, 2001, in conjunction with the International Conference on Information Systems (<http://isds.bus.lsu.edu/icis2001/>). Section on Panel Chairs did not appear in Proceedings, due to lack of space.

Subject Matter

Business Intelligence is a nebulous term, as used in the popular business literature. For the purposes of this panel, we define it broadly as *the use of high-level software intelligence for business applications*. More specifically, business intelligence can be defined by the collection of cutting-edge technologies that help to make systems more intelligent.² This includes:

- representation, communication, execution and retrieval of business policies, rules, and processes
- data mining and visualization
- machine learning and knowledge discovery
- information retrieval
- competitive intelligence/analysis
- dynamic pricing
- agents and the Semantic Web
- recommendation and reputation systems
- automated contracting, brokering, and negotiation

Over the past five years, the appearance of applications requiring or benefiting from business intelligence has accelerated. For example, electronic markets for the buying and selling of goods and services over the Web is a fast-growing, multi-billion-dollar segment of the world economy. Knowledge-based techniques for product recommendation, auctions, need identification, vendor selection, negotiation, agent communication, ontologies, business rules, and information integration are of rising interest, in part due to the rise of XML, and have started having practical impact on real Web e-markets. As more knowledge-based pieces of e-commerce have developed, issues are arising of how to put them together into overall functioning markets. In tandem, intra-enterprise aspects of e-commerce are a driver for knowledge-based techniques in e-business more generally. To be sure, although e-commerce has been most visible lately, intelligent technologies have much farther-reaching impact. Knowledge representation and information retrieval technologies are crucial to knowledge management, for example.

Recent significant progress in commercially deployed business intelligence includes:

- shopping agents, including Web portal services, that use knowledge representation, decision analysis, machine learning, and information retrieval techniques.
- recommender services, e.g., e-storefronts that use collaborative filtering.
- data mining by sellers to learn customer buying patterns
- auction services
- XML-encoded domain ontologies and business-message languages
- online product/service catalogs, e.g., techniques to aggregate catalogs
- customer relationship management: categorization and routing of e-mail, case-based associative retrieval, and suggestions generation
- Web information retrieval and information integration, including using NLP, text analysis and machine learning

² The philosophical debate over the meaning/scope of intelligence and whether systems indeed can be intelligent is beyond the scope of this panel. We define “more intelligent” as increased capabilities that previously were the domain of human knowledge workers, or analytical capabilities that previously were not even possible. Under this definition, for example, information retrieval technologies clearly are intelligent technologies.

Other recent significant progress in business intelligence includes theory, and in many cases emerging industry standards, for:

- auctions, economic decision-making, markets, negotiations, and contracts, including from the viewpoints of resource-bounded intelligence, game theory, distributed AI, negotiation, probabilistic and uncertain reasoning, and decision analysis.
- XML-encoded business rules, ontologies, and business-message/agent-communication languages – in short, the business Semantic Web – including for negotiation and knowledge interchange
- intelligent information integration, including from relational, object-relational, textual, and XML databases distributed on the Web and/or intranets.

Exciting research areas also include (but are not limited to):

- intelligent agents for BI, e.g., with rule-based or probabilistic reasoning.
- knowledge representation to describe goods and services, e.g.: terms and conditions, contractual agreements
- buyer and seller decision making, including pricing and bidding
- brokering and matchmaking
- reputation, recommendation, and other third-party services
- promotions, advertising, and navigation of buyer attention
- intelligent presentation of information, e.g., customized to buyer interests
- opportunities and timings of business intelligence techniques, e.g., relative to other software techniques and relative to evolution of (real-world) markets
- BI-relevant aspects of business processes especially in B2B

Problems/Questions

Information Systems researchers are in a unique position to benefit from this attention to intelligent technologies for business applications. A lesson learned (the hard way) over years of research in Artificial Intelligence (and elsewhere) is that deep knowledge both of the technology and the application area is crucial to the success of intelligent applications. Although of course there are exceptional individuals, other fields of business research typically do not have deep understanding of the technologies and other fields of technical research typically do not have deep understanding of business. It would seem that business intelligence is an ideal topic for Information Systems research.

The first focus of the panel will be a critical look at this seemingly exciting position. Does the argument hold water? Or is there a fundamental flaw? Is business intelligence indeed as important as the current popularity makes it seem? Or is it a passing fad?

The second focus of the panel will be, under the assumption that business intelligence is an area of special interest to information systems researchers, what is the best strategy for IS research? What are the Strengths, Weaknesses, Opportunities, and Threats of BI-research for the IS-community? Are there specific actions that IS researchers should take to build a coherent scientific community?

Panelists

Steven Kimbrough – University of Pennsylvania, The Wharton School

Dr. Kimbrough is a Professor at The Wharton School, University of Pennsylvania. His main research interests are in the fields of electronic commerce (and formal languages for business communication),

knowledge and information management, and computational rationality. His active research areas include: computational approaches to belief revision and nonmonotonic reasoning, formal languages for business communication, evolutionary computation (including genetic algorithms and genetic programming), and information discovery in unstructured and semi-structured data bases (e.g., text). He was principal investigator for the U. S. Coast Guard's KSS (knowledge-based decision support systems) project, and co-principal investigator on the Logistics DSS project, which is part of DARPA's Advanced Logistics Program. He is currently participating in the NSF-funded project "Executing Genetic Algorithms Using DNA Genetic materials."

Stuart Madnick – MIT, Sloan School of Management

Professor Stuart Madnick has been a faculty member at M.I.T. since 1972. He has served as the head of MIT's Information Technologies Group for more than ten years. During that time the group has been consistently rated #1 in the nation among business school information technology programs (U.S. News & World Reports, BusinessWeek, and ComputerWorld). He is also an affiliate member of MIT's Laboratory for Computer Science, a member of the research advisory committee of the International Financial Services Research Center, and a member of the executive committee of the Center for Information Systems Research.

Dr. Madnick is a prolific writer and is the author or co-author of over 250 books, articles, or reports including the classic textbook, *Operating Systems* (McGraw-Hill), and the book, *The Dynamics of Software Development* (Prentice-Hall). He has also contributed chapters to recent books: *The Corporation of the 1990s: Information Technology and Organizational transformation* (Oxford University Press) and *Information Technology in Action* (Prentice-Hall).

His current research interests include connectivity among disparate distributed information systems, database technology, software project management, and the strategic use of information technology. He is presently co-Director of the PROductivity From Information Technology (PROFIT) Initiative and co-Heads the Total Data Quality Management (TDQM) research program.

He has been active in industry, making significant contributions as a key designer and developer of projects such as IBM's VM/370 operating system and Lockheed's DIALOG information retrieval system. He has served as a consultant to many major corporations, such as IBM, AT&T, and Citicorp. He has also been the founder or co-founder of several high-tech firms, including Intercomp (acquired by Logicon), Mitrol (acquired by General Electric's Information Systems Company), and Cambridge Institute for Information Systems (subsequently re-named Cambridge Technology Group), and currently operates a hotel in the 14th century Langley Castle in England.

Dr. Madnick has degrees in Electrical Engineering (B.S. and M.S.), Management (M.S.), and Computer Science (Ph.D.) from MIT. He has been a Visiting Professor at Harvard University, Nanyang Technological University (Singapore), University of Newcastle (England), and Technion (Israel).

Daniel O’Leary – University of Southern California, Marshall Business School

Daniel E. O’Leary, professor of accounting at USC’s Marshall business school, has combined academics and professional experience throughout his career. He has worked as a management consultant in information systems and as a vice president of finance and computer operations for a national advertising agency. He is a Certified Public Accountant, Certified Management Accountant, Certified Information Systems Auditor and holds a Certificate in Data Processing. He has consulted with industry and government on artificial intelligence in business. His teaching includes database systems, business process reengineering, virtual organizations and artificial intelligence. He is the editor of International Journal of Intelligent Systems in Accounting, Finance and Management, among other publications.

EDUCATION: Bowling Green University (B.S.);University of Michigan (MBA); Case Western Reserve University (Ph.D.)

Edward (Ted) Stohr – Stevens Institute of Technology

Professor Stohr holds a Bachelor of Civil Engineering degree from Melbourne University, Australia, and M.B.A. and Ph.D. degrees in Information Science from the University of California, Berkeley. He is currently a Research Professor at Stevens Institute of Technology, Hoboken, New Jersey. His research interests are centered on the problems of developing computer systems to support work and decision-making in organizations. He has been the editor of two books on decision support systems and has published articles in many leading journals.

Prior to coming to Stevens, Professor Stohr was a Professor at the Stern School of Business at New York University and prior to that, he was an Associate Professor at the Graduate School of Management, Northwestern University. For the period 1984-95 Professor Stohr served as Chairman of the Information Systems Department at Stern. Until recently, he was Director of the Center for Research in Information Systems at the Stern School. In 1992, Professor Stohr served as chairman of the executive board of the International Conference on Information Systems (ICIS). A founding member of the Association for Information Systems (AIS), Professor Stohr served on the program committee of the first AIS conference held in Pittsburgh, August 25-27, 1995.

He is on the editorial boards of a number of journals including the Journal of Information Systems Research, Information Systems Frontiers, the International Journal of Decision Support Systems, and The Journal of Management Information Systems. He is also Editor-at-Large for the Association of Information System’s electronic journals, Communications of the AIS and Journal of the AIS. He was guest editor, for special issues on Decision Support Systems of the Journal of Management Information Systems, 1986 and 1987, for a special issue of Management Information Systems Quarterly on IS Curricula and Pedagogy in 1995, and is currently a co-editor for a special issue of Information Systems Frontiers on workflow automation. Professor Stohr has also acted as consultant to a number of major corporations.

Panel Chairs

Abraham Bernstein (<http://www.stern.nyu.edu/~abernste>) is an Assistant Professor for Information Systems at the NYU Stern School of Business's Department for Information Systems. His current research focuses on various aspects of supporting dynamic (intra- and inter-) organizational activity in domains such as knowledge discovery, mobile/pervasive computing, and multi-agent systems. His work bases on both social science (organizational psychology/sociology) and technical (computer science, artificial intelligence) foundations. Mr. Bernstein has a Ph.D. from MIT, where he has played a key role in the development of the Process Handbook. Prior to joining the Center for Coordination Science at MIT he worked for Union Bank of Switzerland first as a research scientist at the corporate research center for information technology (UBILAB) and then as a project manager for IT-projects, where he worked on a variety of research issues like HCI for complex tasks, document management, workflow management and data warehousing. Mr. Bernstein also holds a Diploma in Computer Science (comparable to a M.S.) from the Swiss Federal Institute in Zurich (ETH). He is a member of ACM and has publications in international conferences and journals.

Benjamin Grosf (<http://www.mit.edu/~bgrosf>) is Douglas Drane Assistant Professor in Information Technology at the MIT Sloan School of Management. His research is to create and study knowledge-based information technology for e-commerce applications, including high-level business/agent communication, information integration, contracts/negotiation, trust, product descriptions, business rules/policies, and Web services. The pioneer of inter-operable XML business rules and of their application to contracting, he co-leads the RuleML emerging standards effort on inter-operable XML/RDF rules. He is PI currently for a project in the DARPA Agent Markup Language (DAML) initiative, to create Semantic Web technology and explore its business applications. Formerly he was a senior research scientist at IBM T.J. Watson Research Center (12 years there), where most recently he conceived and led IBM CommonRules (V3.0 currently on IBM alphaWorks) and co-led its application piloting for rule-based XML agent contracting in EECOMS, a \$29Million NIST industry consortium project on manufacturing supply chain management. His notable technical contributions also include fundamental advances in rule-based security authorization, conflict handling for rules, rule-based intelligent agents, and integration of rules with machine learning. He is author of over 30 refereed publications, two major software releases, and a patent. His recent service includes co-chairing the AAAI (National Conference on AI) Workshops on AI in E-Commerce (1999) and Knowledge-Based E-Markets (2000). His background includes 2 years in software startups, PhD in Computer Science (specialty AI) from Stanford University, and BA in Applied Mathematics from Harvard University.

Foster Provost (<http://www.stern.nyu.edu/~fprovost>), Associate Professor of Information Systems at NYU's Stern School of Business, studies data mining, knowledge systems, and machine learning, and their alignment with real-world problems. He has applied intelligent technologies to a variety of business problems, including fraud detection, network diagnosis, and customer contact management. His research contributions include notable technical advances regarding robust modeling in the face of imprecision in the business

environment, profiling/monitoring on-line activity, modeling with sensitivity to economic considerations, and the identification of the types of problems for which different learning techniques are preferable.

In 2000 and in 2001 Professor Provost won IBM Faculty Awards for outstanding research in these areas. This past year he co-chaired the program of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD-2001). He was elected as a founding board member of the International Machine Learning Society, is an editor of the journal Machine Learning and a member of the editorial board of the Journal of Machine Learning Research. His editorial essay on the scientific contributions of applied research papers has been recommended reading for authors in the machine-learning community. His special issue of the international journal Data Mining and Knowledge Discovery, on Data Mining for E-Commerce Applications, now is available as a book.

Based on extensive research, we survey applications and algorithms to help bridge the technology gap. Engineering and construction is behind the curve in implementing artificial intelligence solutions. Based on extensive research, we survey applications and algorithms to help bridge the technology gap. DOWNLOADS. Open interactive popup. Owners and contractors can employ a recommender system approach (supervised learning) that uses cluster behavior production to identify the important data necessary for making a recommendation. 10. Artificial intelligence | The next frontier for investment management firms. Adoption and implementation risks. The journey to successfully adopting and implementing AI will involve strategic and operational risk considerations which early adopters, niche innovators, and large-scale players will need to carefully evaluate and address. Clearly define your AI strategy, articulating your business goals and understanding how AI will be utilized as part of your. Doug Dannemiller Research Leader, Investment Management Deloitte Center for Financial Services, Deloitte United States ddannemiller@deloitte.com. Contributors In addition, we thank the following people from Deloitte United States for their contributions to this paper: Liz Bock, Chaithanya Mangalampalli, and Vincenzo Rispoli. Mark Walker of IDC, a research firm, says that in the past three or four years multinational companies have adopted a "completely fresh approach". They have "a lot more skin in the game: investing in local people, so there's proper knowledge transfer, investing in country offices." Companies are in it for the long term now, rather than quitting after a bad quarter or two. It can leapfrog straight to the tech frontier, without worrying about adapting old systems to cope with the data it creates. At the Catholic University of Eastern Africa, which will eventually be the new lab's home, a Kenyan asked Mr Kelly, "Will you bring Watson to Africa?" Yes, he replied, if someone suggests a problem for it to solve.