
Meeting the Responsibility to Explain AI

***Bruce Buchanan
University of Pittsburgh (Emeritus)***

***Reid Smith
Marathon Oil Corporation***

***with substantial contributions from
Joshua Eckroth
The Ohio State University***

***and thanks to many people, including Mike Hamilton,
David Leake, Hannah Frost, Stephen Muggleton,
Alan Bundy, Ed Feigenbaum, Edwina Rissland, Nils Nilsson***

What Problem Are We Addressing?

Helping meet the AAI responsibility to:

- **Explain AI goals and methods and convey research results to the lay public**
- **Contribute to the education of young students in science, technology, engineering, and mathematics**
- **Preserve the history of the field**

- **News**
- **Overviews and background information**
- **Course syllabi**
- **Videos & podcasts**
- **Hardware & software tools, systems, ...**
- **History of AI, including people, organizations, ...**
- **Foundational (“classic”) articles and books**
 - **Items integrated and presented in context**

Topics

- ▶ [AI Overview](#)
- ▶ [Applications](#)
- ▶ [Cognitive Science](#)
- ▶ [Education](#)
- ▶ [Ethics & Social Issues](#)
- ▶ [Games & Puzzles](#)
- ▶ [History](#)
- ▶ [Interfaces](#)
- ▶ [Machine Learning](#)
- ▶ [Natural Language](#)
- ▶ [Philosophy](#)
- ▶ [Representation & Reasoning](#)
- ▶ [Robots](#)
- ▶ [Science Fiction](#)
- ▶ [Speech](#)
- ▶ [Systems & Languages](#)
- ▶ [Vision](#)
- ▶ [Web & AI](#)

Collections

- ▶ [Interviews](#)
- ▶ [Tributes](#)
- ▶ [News](#)
- ▶ [Classics](#)

Search

Search found 25 items

Publication

The Theory That Would Not Die (Review)

... Review of The Theory That Would Not Die: How **Bayes'** Rule Cracked the Enigma Code, Hunted Down ... the way **Bayes** began his search for an answer to the inverse probability problem. Probability theory was in its infancy in **Bayes's** time and, McGrayne writes, applied primarily to gambling: the odds ...

Grossman, Wendy M
ZDNet
2011



[History > People, Programs, Institutions](#)

Link

Bayes, Thomas

... Thomas **Bayes (1702-1761)? Bayes** Theorem 18th-century theory is new force in computing. By Michael Kanellos. CNET News.com (February 18, 2003). "Thomas **Bayes**, one of the leading mathematical lights in ... logic, a whatis definition from TechTarget. "**Bayes** first proposed his theorem in his 1763 work ...



[History > People, Programs, Institutions](#)

Sort by

- [Relevancy](#)
- [Title](#)
- [Type](#)
- [Author](#)
- [Date](#)

Filter by topic

- ▶ [AI Overview \(3\)](#)
- ▶ [Applications \(5\)](#)
- ▶ [History \(3\)](#)
- ▶ [Natural Language \(4\)](#)
- ▶ [Representation & Reasoning \(10\)](#)

Filter by year

- ▶ [1763 \(1\)](#)
- ▶ [1982 \(1\)](#)
- ▶ [1991 \(2\)](#)
- ▶ [1999 \(1\)](#)
- ▶ [2001 \(1\)](#)
- ▶ [2003 \(2\)](#)
- ▶ [2007 \(2\)](#)
- ▶ [2008 \(1\)](#)
- ▶ [2010 \(1\)](#)
- ▶ [2011 \(2\)](#)

Filter by topic

- › AI Overview (3)
- › Applications (5)
- › History (3)
- › Natural Language (4)
- › Representation & Reasoning (10)

Filter by year

- › 1763 (1)
- › 1982 (1)
- › 1991 (2)
- › 1999 (1)
- › 2001 (1)
- › 2003 (2)
- › 2007 (2)
- › 2008 (1)
- › 2010 (1)
- › 2011 (2)
- › 2012 (1)

Filter by panel

- › Educational Resources (1)
- › General Readings (14)
- › Good Starting Places (1)

Filter by tag

- › Bayes (1)
- › book review (1)
- › classic (4)
- › collection (1)
- › courses (1)
- › Edinburgh (1)
- › massive open online courses (1)

Index of Machine Learning Courses

- [Advanced Methods in Decision Support and Data Mining](#), Jose Pena & Tomas Kocka, Aalborg University
- [Machine Learning](#), Altay Guvenir, Bilkent University, Turkey.(Spring '02)
- [Introduction to Machine Learning and Neural Networks](#), Dan Ventura, Brigham Young University.
- [Machine Learning](#), Carlos Guestrin and Tom Mitchell, Carnegie Mellon University.
- [Machine Learning Theory](#), Avrim Blum, Carnegie Mellon University. (Fall 1995)
- [Topics in Machine Learning](#), Rick Caruana, Cornell University (Spring 2003)
- [Introduction to Machine Learning](#), Amnon Shashua, Hebrew University of Jerusalem.
- [Computational Learning Theory and Neural Networks](#), Martin Anthony, The London School of Economics and Political Science
- [Machine Learning](#), Dan Melamed, New York University
- [Machine learning](#), Vasant Honavar, Iowa State University.
- [Advanced Topics in Artificial Intelligence and Cognitive Modelling](#) (a *special topics* course which sometimes covers machine learning), Vasant Honavar, Iowa State University, (Spring 2001).
- [Computational Models of Learning](#), Vasant Honavar, Iowa State University.
- [Machine Learning](#), Steven Salzberg Johns Hopkins University (Spring 1996).
- [Machine Learning and Inductive Inference](#) Hendrik Blockheed, Katholieke Universiteit Leuven, Belgium
- [Machine Learning](#), Ron Rivest, Massachusetts Institute of Technology. (Fall 1994)
- [Machine Learning](#), K. Korb, Monash University, Australia
- [Intelligent Systems I](#), Oxford University, UK. (requires registration to access full course materials)
- [Intelligent Systems II](#), Oxford University, UK. (requires registration to access full course materials)
- [Machine Learning](#), Carla Brodley, Purdue University. (Fall 1995)
- [Machine Learning Course](#), Zdravko Markov, Sofia University
- [Machine Learning](#), Andrew Ng, Stanford University
- [Statistical Machine Learning](#), Jacob Goldberger, Weizman Institute (Fall 2002/03)
- [Machine Learning](#), Michael Cox, Wright State University.
- [GSLT: Machine Learning](#), Joakim Nivre, et. al, Vxj Universitet, Sweden
- [Knowledge Acquisition and Machine Learning](#), Katharina Morik, University of Dortmund, Germany. (In German)
- [Machine Learning](#), Ricardo Vilalta, University of Houston
- [Data Mining](#), Mark Last, University of South Florida. (fall 1999 - files are password protected)

Topics

- › [AI Overview](#)
- › [Applications](#)
- › [Cognitive Science](#)
- › [Education](#)
- › [Ethics & Social Issues](#)
- › [Games & Puzzles](#)
- › [History](#)
- › [Interfaces](#)
- › [Machine Learning](#)
- › [Natural Language](#)
- › [Philosophy](#)
- › [Representation & Reasoning](#)
- › [Robots](#)
- › [Science Fiction](#)
- › [Speech](#)
- › [Systems & Languages](#)
- › [Vision](#)
- › [Web & AI](#)

Collections

- › [Interviews](#)
- › [Tributes](#)
- › [News](#)
- › [Classics](#)
- › [Podcasts](#)
- › [Videos](#)
- › [Course materials](#)

Search

Search found 14 items

Link

IEEE collection of course syllabi

... ACM and IEEE Computer Society are jointly sponsoring the development of a Computing Curricula volume on Computer Science, planned for publication in 2013 (cs2013.org). The Steering Committee of CS2013 Curriculum is seeking exemplars of courses and curricul ...



[Education](#)

Link

Probabilistic Graphical Models

... In this class, you will learn the basics of the PGM representation and how to construct them, using both human knowledge and machine learning techniques; you will also learn algorithms for using a PGM to reach conclusions about the world from limited and ...

Koller, Daphne
Stanford



[Representation & Reasoning](#) > [Bayesian Inference](#)

Link

Artificial Intelligence in Natural Resource Management

... Homepage for Professor H. Randy Gimblett's course at the School of Renewable Natural Resources University of Arizona. As stated in the Fall 2002 syllabus: "The purpose of this course is to explore agent based modeling to aid in automating the de ...

U Arizona

Sort by

- [Relevancy](#)
- [Title](#)
- [Type](#)
- [Author](#)
- [Date](#)

Filter by topic

- › [AI Overview](#) (2)
- › [Applications](#) (2)
- › [Cognitive Science](#) (2)
- › [Education](#) (1)
- › [Interfaces](#) (1)
- › [Machine Learning](#) (1)
- › [Natural Language](#) (2)
- › [Representation & Reasoning](#) (2)
- › [Robots](#) (1)

Filter by year

- › [1988](#) (1)
- › [2002](#) (1)
- › [2004](#) (3)
- › [2007](#) (2)
- › [2009](#) (1)
- › [2012](#) (2)

Filter by panel

- › [Educational Resources](#) (14)

Topics

- › [AI Overview](#)
- › [Applications](#)
- › [Cognitive Science](#)
- › [Education](#)
- › [Ethics & Social Issues](#)
- › [Games & Puzzles](#)
- › [History](#)
- › [Interfaces](#)
- › [Machine Learning](#)
- › [Natural Language](#)
- › [Philosophy](#)
- › [Representation & Reasoning](#)
- › [Robots](#)
- › [Science Fiction](#)
- › [Speech](#)
- › [Systems & Languages](#)
- › [Vision](#)
- › [Web & AI](#)

Collections

- › [Interviews](#)
- › [Tributes](#)
- › [News](#)
- › [Classics](#)
- › [Podcasts](#)
- › [Videos](#)
- › [Course materials](#)

Search

Search found 617 items

Link

Alspace: Tools for Learning Artificial Intelligence

... "These applets are designed as **tools** for learning and exploring concepts in artificial intelligence. They are part of the online resources for Computational Intelligence: A Logical Approach. If you ... Knoll, under the guidance of Cristina Conati, Peter Gorniak, Holger Hoos, Alan **Mackworth**, and David ...



[AI Overview](#)

Link

Alspace-Tools for Learning Computational Intelligence

... input values." From Alspace-**Tools** for Learning Computational Intelligence: "Here are some applets that are designed as **tools** for learning and exploring concepts in artificial intelligence. They ...



[Machine Learning > Neural Networks](#)

Link

Alspace: Neural Networks

... These **tools** are for learning and exploring concepts in artificial intelligence. They were ... direction of Alan **Mackworth** and David Poole. They are part of the online resources for Artificial ... **tools** under the terms of use. Feedback is welcome. ... **Mackworth**, Alan and David Poole.

Sort by

- [Relevancy](#)
- [Title](#)
- [Type](#)
- [Author](#)
- [Date](#)

Filter by topic

- › [AI Overview \(35\)](#)
- › [Applications \(185\)](#)
- › [Cognitive Science \(41\)](#)
- › [Education \(25\)](#)
- › [Ethics & Social Issues \(27\)](#)
- › [Games & Puzzles \(36\)](#)
- › [History \(27\)](#)
- › [Interfaces \(75\)](#)
- › [Machine Learning \(34\)](#)
- › [Natural Language \(55\)](#)
- › [Philosophy \(14\)](#)
- › [Representation & Reasoning \(83\)](#)
- › [Robots \(45\)](#)
- › [Speech \(14\)](#)
- › [Systems & Languages \(12\)](#)
- › [Vision \(25\)](#)
- › [Web & AI \(9\)](#)


Filter by person


- › [Backus, John \(1\)](#)



Model AI Assignments

From

EAAI-2013: The Fourth Symposium on Educational Advances in Artificial Intelligence 

Bellevue, Washington, USA (collocated with AAI-13 )

July 15-16, 2013

What is the Model AI Assignments Session?

The Model AI Assignments Session seeks to gather and disseminate the best assignment designs of the Artificial Intelligence (AI) Education community.

One must learn by doing the thing; for though you think you know it, you have no certainty, until you try. - Sophocles

Recognizing that assignments form the core of student learning experience, we invite AI educators to submit draft assignment materials that exemplify an approach to teaching AI topics at all levels.

Todd Neller, Gettysburg

2013

Found in [Educational Resources](#)











[AI Overview](#) > [Computer Science](#)
[Education](#) > [Intelligent Tutoring](#)



homework

Videos and Podcasts

| | | | | | |
|---|---|--------|---------------------------|------------------|------|
|  | How To Survive A Robot Uprising (aka Robopocalypse) | 4 min | Robot Ethics | | 2012 |
|  | Machine Learning for Robotics | 57 min | Machine Learning, Robots | videlectures.net | 2012 |
|  | The quest for a theory of vision: from the level framework (revised) to the invariances of the ventral stream | 55 min | Cognitive Science, Vision | videlectures.net | 2012 |
|  | Inside Watson's Brain: The Hardware Story | 5 min | Jeopardy! | IBM | 2012 |

| | | | | | |
|---|---|--------|---|----------------|-------------|
|  | Talking to Machines | 60 min | Emotion, Discourse & Dialogue, Turing Test, Humanoid Robots | Radiolab | May 21 2011 |
|  | European Robotics - Part 2 | 26 min | Robots | Robots Podcast | May 7 2011 |
|  | European Robotics (Part 1) | 27 min | Robots | Robots Podcast | Apr 22 2011 |
|  | The Dark Side of Watson | 5 min | AI Overview, Ethics & Social Issues | NPR | Feb 20 2011 |

Science Fiction Classics

Publication

Nineteen Eight-Four



By Orwell, George, 1949



classic, scifi books

Publication

The War of the Worlds



By Wells, H. G. , 1898



classic, scifi books

Publication

R.U.R. (Rossum's Universal Robots)



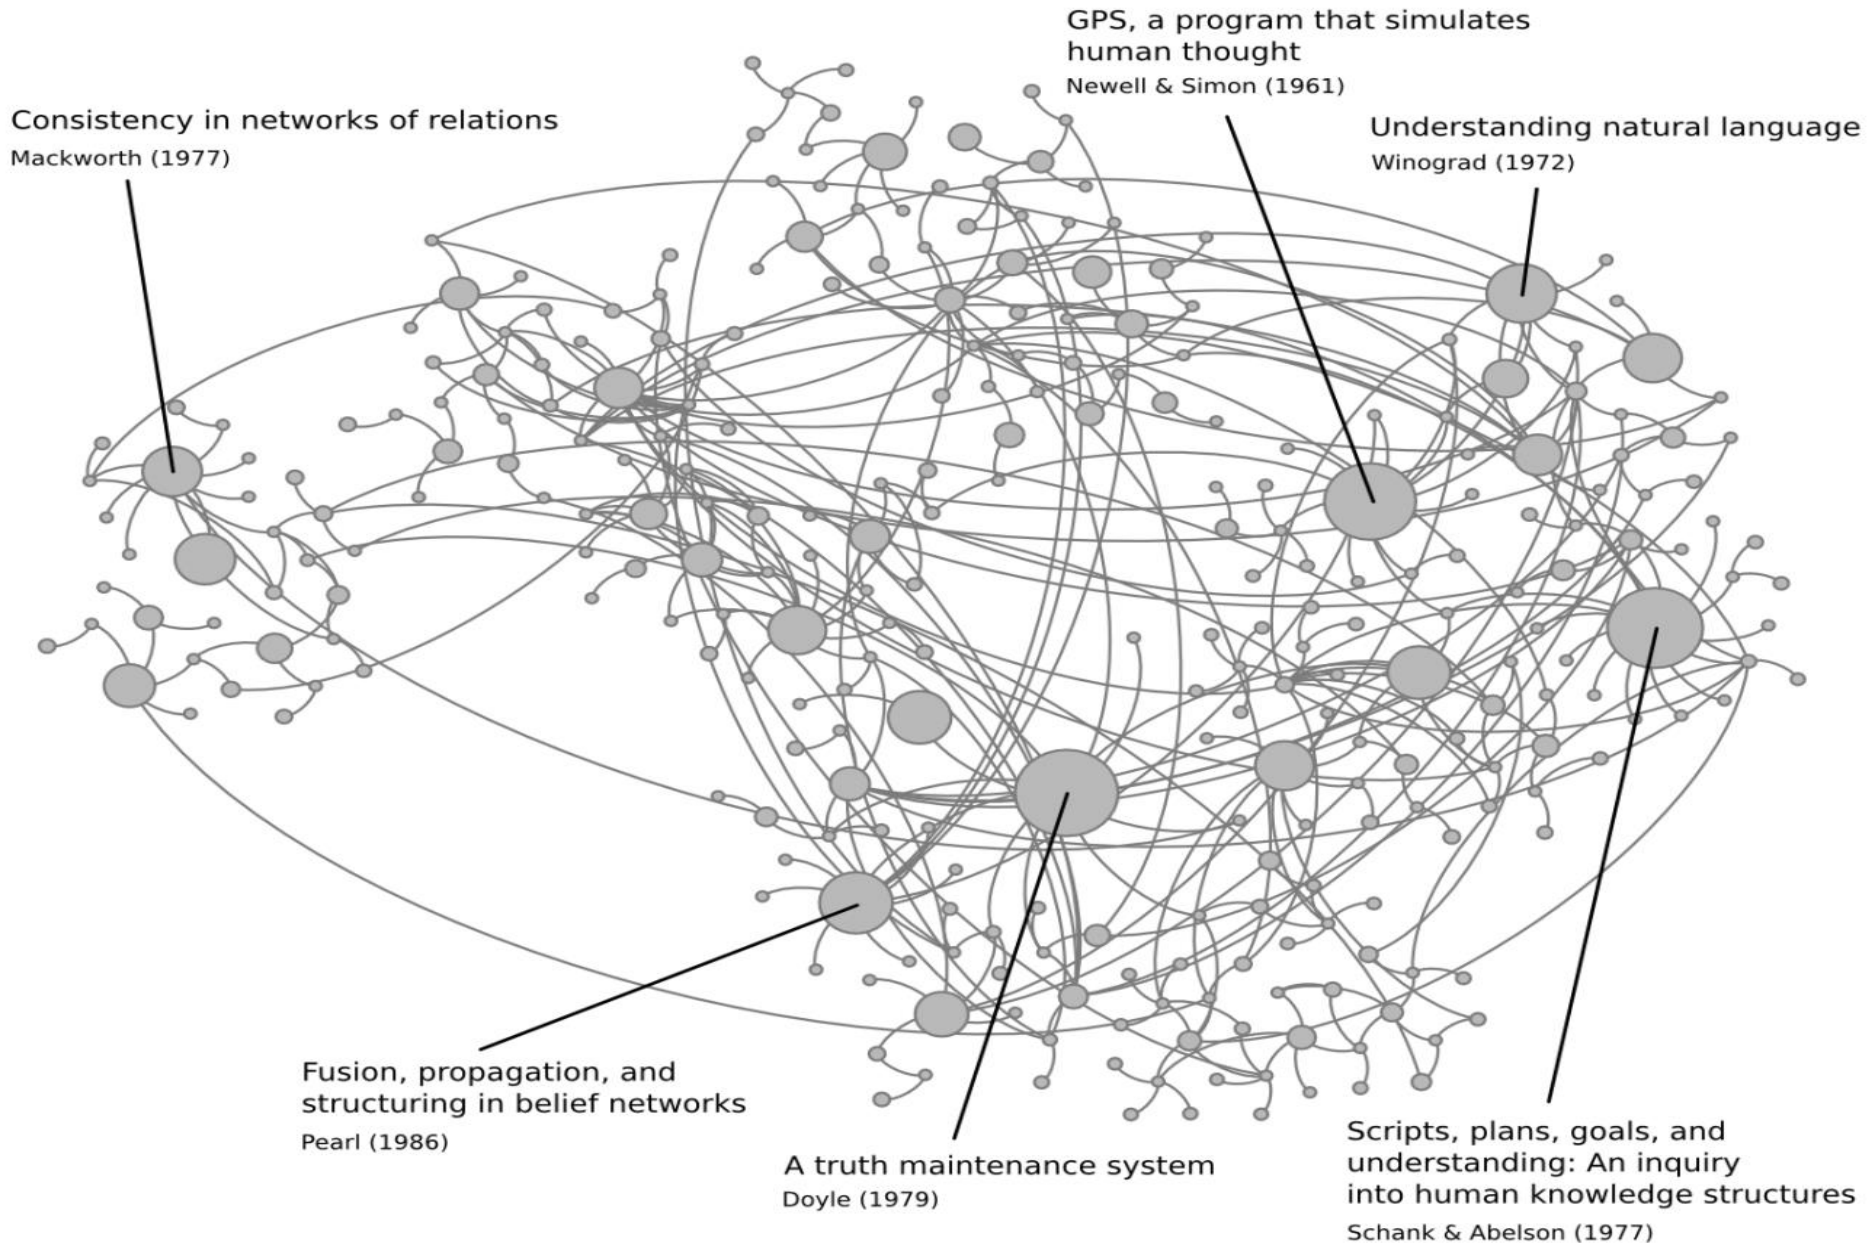
Karel Čapek wrote with intelligence and humor on a wide variety of subjects. He is perhaps best known as a science fiction author, who wrote long before science fiction became established as a separate genre. He can be counted as one of the founders of classical non-hardcore European science fiction, which focuses on possible future (or alternative) social and human evolution on Earth, rather than technically advanced stories of space travel.

By Capek, K., 1920



classic, scifi books

Network of Citations in the Classics



Google Analytics: Jan 10 - Jul 10, 2013

- **26,663 visitors** – 30% outside North America
- **41,182 visits** – 226 per day, 4 min per visit, 4% from mobile devices
- **97,825 page views**
- **1000 referral sites**
- **Top Content**

| | |
|------------------------|------------------------|
| Home: 26.3% | Search: 5.5% |
| News: 4.0% | Brief History: 3.7% |
| Machine Learning: 2.8% | AI Overview: 2% |
| Applications: 1.3% | Videos: 1.2% |
| Classics: 1.2% | Computer Science: 0.9% |

Automation: AI in the News

- **Find**
 - Start with list of credible sources with RSS feeds.
- **Classify**
 - Topic hierarchy (20 main topics, 180 topics & subtopics in two-level hierarchy)
 - Train classifiers from corpus of pre-classified documents
- **Select**
 - Use strength of classification from classifiers + merit of source + diversity
- **De-duplicate**
 - Remove stories describing the same event published within the same week
- **Summarize**
 - Select three sentences for summary
- **Publish and Archive**
 - Weekly e-mail, RSS feed, AITopics web site ... Facebook, Twitter, Google Plus ... mobile devices
 - Search, including tags and facets

Technologies Used for Automation

- **Document collection and extraction**

- RSS feeds for time-based periodicals: journals, newspapers, ...
- Goose: text and image extraction from web pages
- NLTK and Apache OpenNLP: parsing, tokenization, ...
- Tesseract: OCR for PDF
- Apache Tika: text and metadata extraction for PDF, MS Office, ...
- New implementation of Frantzi et al. C-value / NC-value method

- **Document summarization and classification**

- OTS: text summarization
- Weka: suite of machine learning algorithms for classification

- **Document management and retrieval**

- Drupal: content management system
- Apache Solr: faceted search


- **Select Overview & Review Articles**

- Experiment with random set of 255 articles from AI Magazine
 - 54 selected as overviews based on machine learning classifier
 - $F1 = 0.88$


- **Analyze information in the repository**

- Example: 25 years of IAAI papers
 - Multi-Word Term Extraction
 - Metadata Extraction (author, company, ...)
 - Classification (industry, AI technology, task, ...)
- There is a lot more to be done
 - We can provide data for experimentation

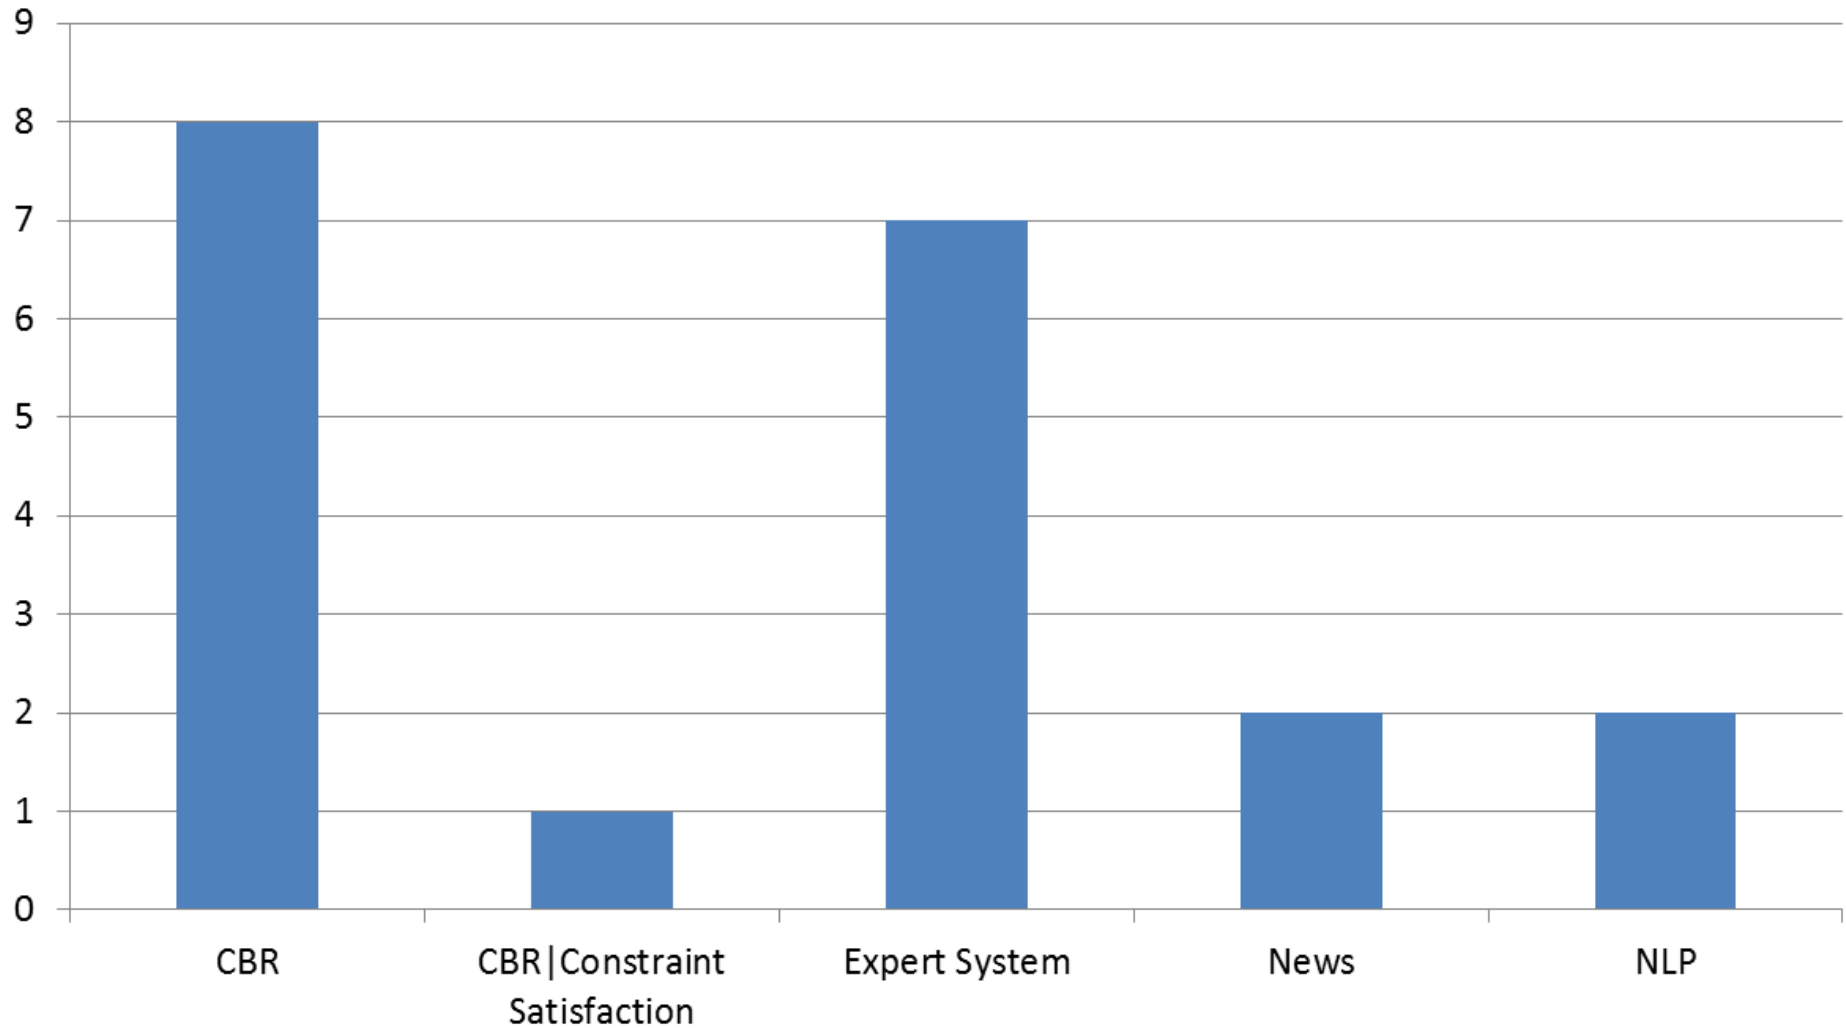
Terms Extracted from 25 Years of IAAI Papers




25 Years of Deployed Applications by Industry



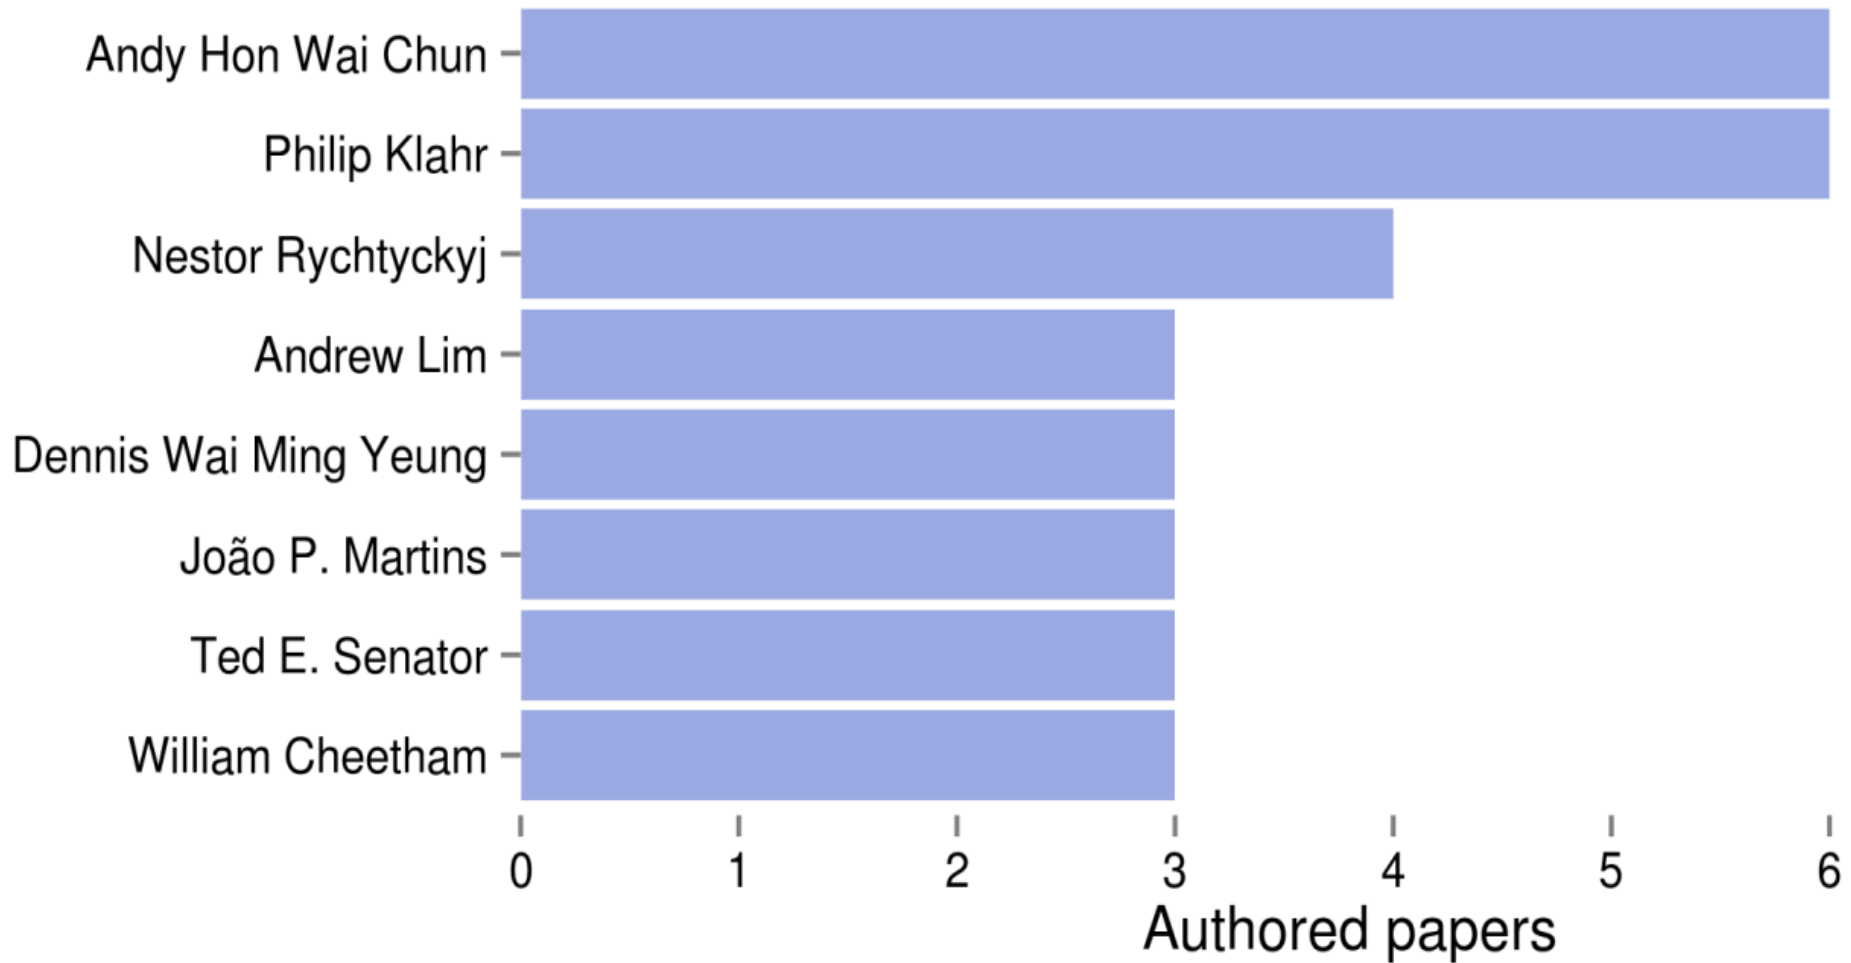
Deployed Applications by AI Technology




25 Years of Deployed Applications by Task



Top Deployed Application Authors



Top Emerging Application Authors



What Is the Larger Context?

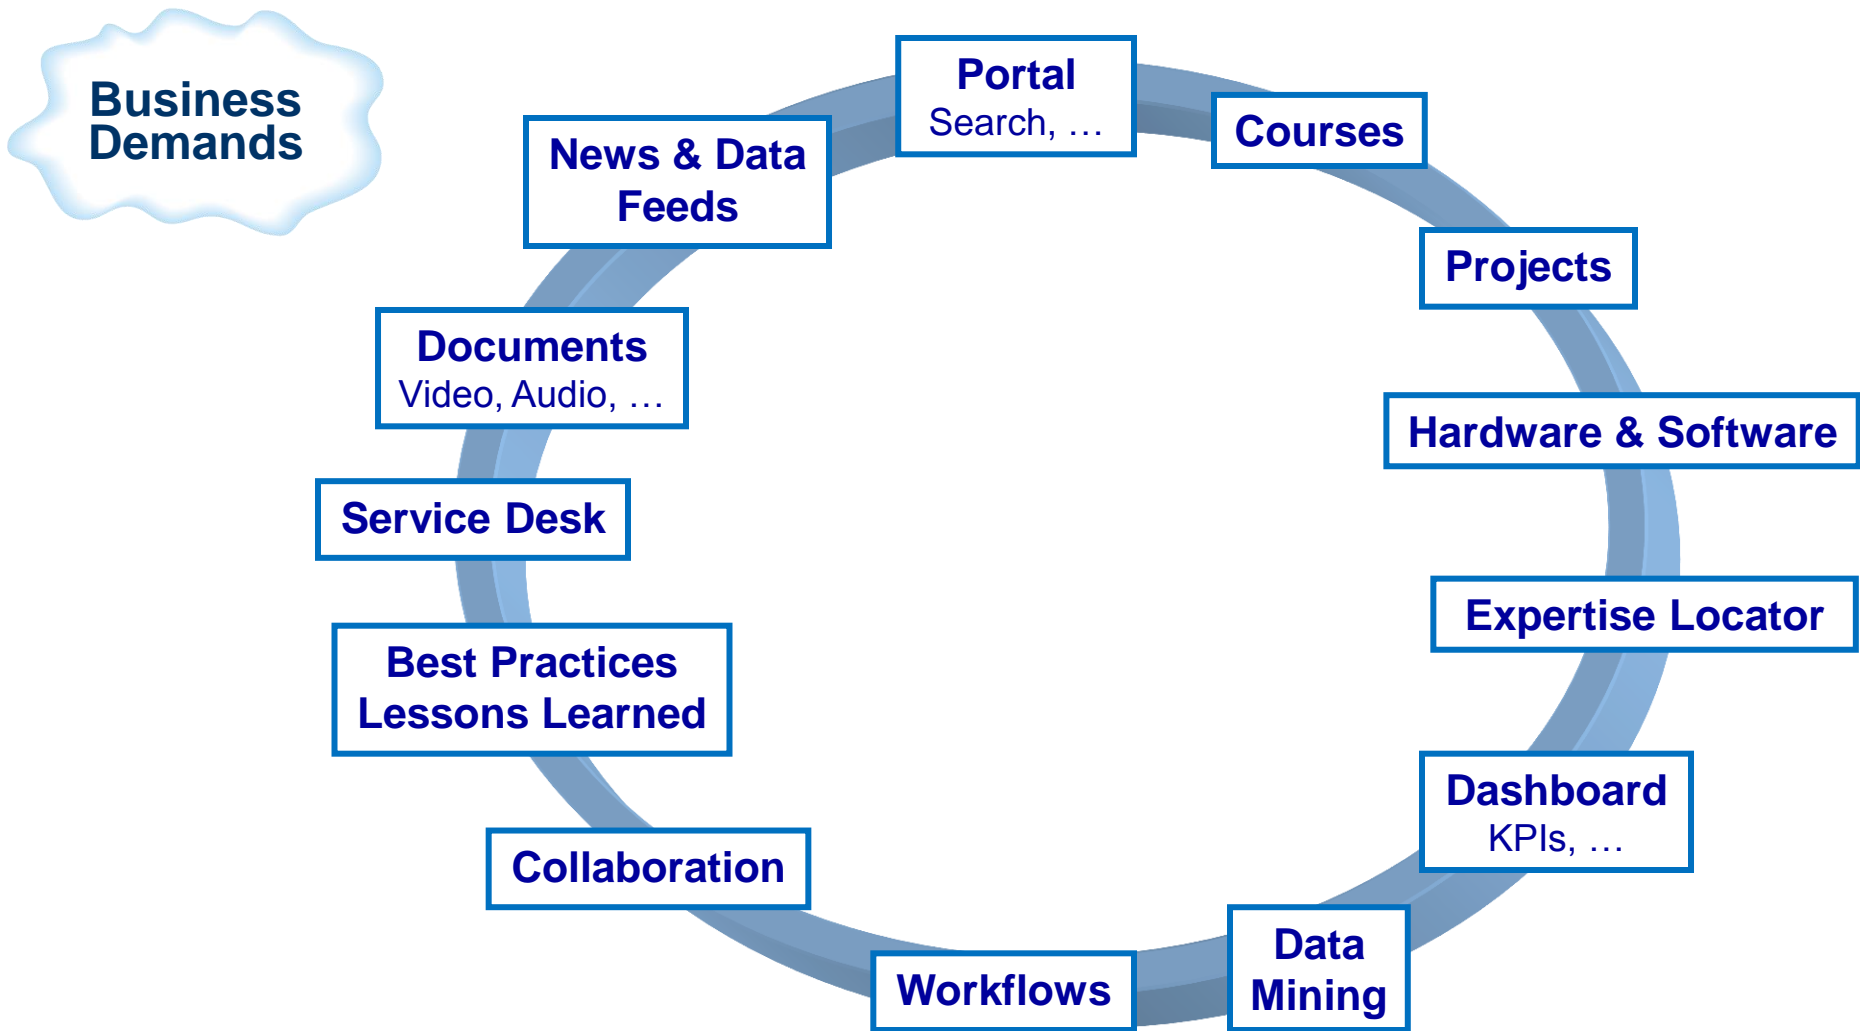
Knowledge Management

- Systematic approaches to help information and knowledge grow and flow to the right people at the right time

Generally Accepted Ideas

- **People** – The “Killer App”
 - **Community of Practice** – Fundamental organizational unit
- **Processes** – Focus on knowledge flow
- **Connect, Collect, Collaborate** – Standard approaches
- **Technology** – Merely an enabler
 - Low level, domain-independent

The Larger Context: Knowledge Management



Exploit AI technology to change the game of meeting the educational responsibilities

- Automate finding and organizing information on the research, the people and the applications of Artificial Intelligence.**
- Leverage efforts currently underway: EAAI, AI Genealogy, ...**
- Build teaching units on specific topics**
- Analyze data: Classics, IAAI, ...**
- Define AAI seal of approval process for content**
- Support student projects and competitions**