Artificial Intelligence: Learning Through Interactions and Big Data

An IBM Redbooks® Point-of-View publication by IBM Academy of Technology

Authored by Koichi Takeda, IBM Research Tokyo, and **Tamiya Onodera**, IBM Research Tokyo

Highlights

AI approaches combined with analytics techniques can provide practical approaches to the challenges of creating smarter solutions.

- The current mainstream business analytics research and solution development along with the convergence of AI with machine learning techniques will continue to underpin high-value, human decision support solutions.
- Cognitive systems can combine natural language processing, hypothesis generation and evaluation, and dynamic learning for a powerful, fast, and intelligent problem solving.
- To effectively apply intelligent problem solving using AI solutions, you can extend the capabilities of IBM Watson and its DeepQA technology architecture to other domains.

Artificial intelligence (AI) has been formed as a field of computer science for understanding and building intelligent entities since the 1950s. As with most technologies, the concept has undergone many transformations over time. In particular, the field of AI has experienced several cycles of expectations, commitments, and then disappointments, sometimes referred to as *AI winters*, where interest in the technology has waned only to then peak again at a later time. Some might argue that these cycles of conflict have caused inconsistencies that have made forward progress difficult in the AI field. However, recent practices of data-centric approaches (for example, machine learning and probabilistic reasoning) are beginning to demonstrate the true capability of AI to solve practical decision-making problems.

The concept of AI is often thought of as having two dimensions:

- One dimension emulates the human thought process or simulates intelligent behavior (such as cognitive computing and brain-inspired systems).
- The other dimension achieves human or ideal performance (such as question-answering and natural language understanding capability as demonstrated by IBM Watson and its Deep Question-Answering technology).

This dual nature of AI challenges organizations with how to best approach AI modeling solutions. Conventional AI approaches often tackle knowledge acquisition and reasoning in terms of a set of ad hoc rules and manually crafted weights for computing confidence. As a result, acquiring and maintaining consistency of rules and improving the accuracy of the reasoning can be a challenge when the size of the knowledge base grows to solve real-world problems. This *knowledge bottleneck* for expert systems is often perceived as an inhibitor for building scalable industrial solutions.

However, AI approaches remain central to intelligent problem solving. Unlike precisely defined mathematical problems, where algorithms for computing complete or approximate answers can be applied, many real-world problems are often underspecified, ambiguous, or incompletely defined and understood. AI techniques turn out to be effective for modeling the problem and building practical solutions.

One important aspect of modern AI is that it now embraces machine learning techniques for building statistical models and making statistical reasoning and prediction. Recent advances in statistical approaches along with the availability of a massive amount of data (big data) have been remarkable. These approaches are driving enhancements for high-value business analytics.

According to Bernard Meyerson, Chief Innovation Officer, IBM, the DeepQA technology of IBM Watson "is just a first step into a new era of computing that's going to produce machines that are as distinct from today's computers as those computers are from the mechanical tabulating devices that preceded them."[1] As organizations embrace AI, understanding the methods to approaching intelligent problem solving is central to gaining an advantage with AI solutions.

AI and business analytics drive smarter solutions

Machine learning and statistical or predictive analysis, in many cases, boils down to business analytics, which is a core component of contemporary AI. Thus, business analytics are central to the AI-related activities that build a smarter solution. When you approach methods to use AI technologies, consider the following AI-related activities:

- Adapting AI technologies to other industries
- Applying cognitive computing to mind-like computational models
- Converging AI and business analytics
- Incorporating operational intelligence into expert integrated systems

Adapting AI technologies to other industries

An AI program can use natural language capabilities, hypothesis generation, and evidence-based learning to produce a better outcome. Researchers are now looking at advances in technologies to leverage AI in other industries. Facts strongly suggest that the question and answer architecture behind the AI technology that made IBM Watson possible can be applied to other domains. If quality information sources are available and a reasonable set of sample questions and answers for the domain exists and if the deep linguistic analysis for the source language is available, this type of technology can adapt well to other industries.

For example, the core technology behind AI can help transform the way a doctor interacts with a patient. By feeding data to an AI program, such as IBM Watson, a doctor can better determine possibilities for treatment. In turn, better data might help reduce unnecessary treatments or tests, thus improving patient outcomes and reducing overall costs, for the patient, the practice, and ultimately, the healthcare industry, while also making processes more efficient.

In another example, a cognitive computing assistant, such as IBM Watson Engagement Advisor, can help a business to better serve its consumers. This technology allows brands to crunch big data and transform the way that they engage clients in key functions, such as customer service, marketing, and sales. IBM Watson Engagement Advisor can learn, adapt, and understand a company's data, enabling customer-facing personnel to assist consumers with deeper insights more quickly than previously possible.

Other examples of adapting AI technologies to other industries might be the deployment of speech recognition and voice command interpreters to smartphones, consumer electronics, and automotives and the application of computer vision to robots.

A system that generates hypotheses can recognize different probabilities of various outcomes. It can "learn" from successes and failures and can improve future responses.

Applying cognitive computing to mind-like computational models

Cognitive computers move beyond simple processing of data and begin to learn through experiences. By applying cognitive computing, models can find correlations, can create hypotheses, and then learn from the outcomes. This process is fundamental research for building the coherent and unified model of *mind-like computation*. Potentially, the resulting model can be applied to solve massive sensory information processing and cognitive problems.

Researchers at IBM have been working on a cognitive computing project called Systems of Neuromorphic Adaptive Plastic Scalable Electronics (SyNAPSE). By reproducing the structure and architecture of the brain—the way its elements receive sensory input, connect to each other, adapt these connections, and transmit motor output—the SyNAPSE project models computing systems that emulate the brain's computing efficiency, size, and power usage without being programmed.

[1] http://asmarterplanet.com/blog/2012/12/the-ibm-5-in-5-our-2012-forecast-of-inventions-that-will-change-the-world-within-five-years.html

One of the earlier achievements in the SyNAPSE project is the (cat-scale) cortical simulations (with 109 neurons and 1,013 synapses) that were run by the Blue Gene/P supercomputer with 147,456 CPUs and 144 TB of total main memory at the Lawrence Livermore National Laboratory.

Cognitive systems can transform how organizations think, act, and operate. Learning through interactions, they deliver evidence-based responses that can drive better outcomes.

Converging AI and business analytics

Recent empirical AI approaches deploy analytical methods—called *machine learning*—by assuming a statistical model for capturing uncertainty in the given problem. Machine learning works by clearly defining business goals, exploring data, and validating and deploying models. It relies on domain and data expertise, rather than analytical skills. This method estimates a set of parameters of a model (the *probability*) from an ample amount of observed data that approximates the problem space.

Machine learning techniques are often theoretically sound, and various mathematical properties are known and used to formulate the given problem. Machine learning approaches are inspired by the advances and success of Business Analytics and Mathematical Science (BAMS) research. BAMS conducts active research in several areas of analytical and mathematical sciences such as optimization, algorithms, stochastic analysis, statistics, data mining, predictive modeling, and simulation.

BAMS current projects include a diverse range of industry domains, such as retail and banking service lines, including Supply Chain Management (SCM) and Human Capital Management (HCM) and solution areas, such as Smarter Cities and Smarter Energy. Examples of projects that BAMS has executed include capacity utilization, hiring and training, strategic planning and budgeting, workforce scheduling, contingency planning, dynamic supply chain optimization, customer targeting, smarter water supply management for growth markets, propensity modeling for cross-sell and up-sell, and so forth.

For more information about BAMS, see:

`http://www-07.ibm.com/in/research/bams.html`

Convergence of AI and business analytics approaches can further extend the scope of analytics into structured and unstructured information, human behavioral patterns, and interactive learning of knowledge. The convergence of AI and analytics promises to model intelligent decision processes.

Incorporating operational intelligence into expert integrated systems

Today's applications are complex. They involve a diverse set of hardware and software components that interact with each other horizontally and vertically. Optimizing, deploying, configuring, and managing these components require an enormous amount of manual effort by human experts. This tuning process is extremely laborious and error-prone. Automating this process in a cost-effective manner belongs in the realm of applications for artificial intelligence.

In the mid 2000s, IBM researchers successfully applied control theory to software systems, resulting in two key features in a commercial product:

- Utility Throttling of DB2 v8.2
- Adaptive Self-Tuning Memory of DB2 v9

These features made it easier to configure and maintain databases for all market segments.

In mid 2012, IBM delivered PureSystems, a breed of expert integrated systems that capture and automate what experts do in order to make every part of the IT lifecycle easier. PureSystems uses a semantically rich language, developed by IBM Research, to define its building blocks. The knowledge that formerly sprawled across hundreds of pages of manuals is now distilled into a formal representation.

As the technology integrates further, expert integrated systems will evolve by incorporating more usage data for improving performance and consumability, which will eventually be known as *operational intelligence*.

A grocer who is stocking store shelves can use an instrumented glove designed for a cognitive computing system that monitors sights, smells, texture, and temperature to flag bad or contaminated produce.

Extending the capabilities of IBM Watson and DeepQA

To effectively apply intelligent problem solving using AI solutions, the successful organization must extend the capabilities of today's technologies. IBM Watson represents only the beginning of a new way of thinking about ways that technology can help us live and work better. IBM Watson and its Deep Question-Answering (DeepQA) technology architecture can be applied to other domains or in other languages for QA.

For example, healthcare question-answering for medical diagnosis support is one of the most promising areas for IBM Watson. In October of 2012, IBM announced a collaboration with Cleveland Clinic clinicians, faculty, and medical students to enhance the capabilities of Watson's DeepQA technology for the area of medicine.[2] Watson's ability to analyze the meaning and context of human language and to quickly process information to provide answers can help healthcare decision makers use the huge volumes of information generated in their industry. Cleveland Clinic along with IBM recognized the opportunity for Watson to interact with medical students to help explore a wide variety of learning challenges that the medical industry faces today.

In February 2013, IBM, WellPoint, Inc., and Memorial Sloan-Kettering Cancer Center announced the first commercially developed Watson-based cognitive computing breakthroughs.[3] For more than a year, IBM has partnered separately with WellPoint and Memorial Sloan-Kettering to train Watson in the areas of oncology and utilization management. During this time, clinicians and technology experts have "taught" Watson how to process, analyze, and interpret the meaning of complex clinical information using natural language processing, all with the goal of helping to improve healthcare quality and efficiency.

These innovations represent a breakthrough in how medical professionals can apply advances in analytics and natural language processing to "big data," combined with the clinical knowledge base to create evidence-based decision support systems. These Watson-based systems are designed to assist doctors, researchers, medical centers, and insurance carriers, and ultimately to enhance the quality and speed of care.

From the healthcare industry, the technology behind IBM Watson is expected to branch into other industries that rely on analytic solutions to manage unstructured data, as illustrated in Figure 1.

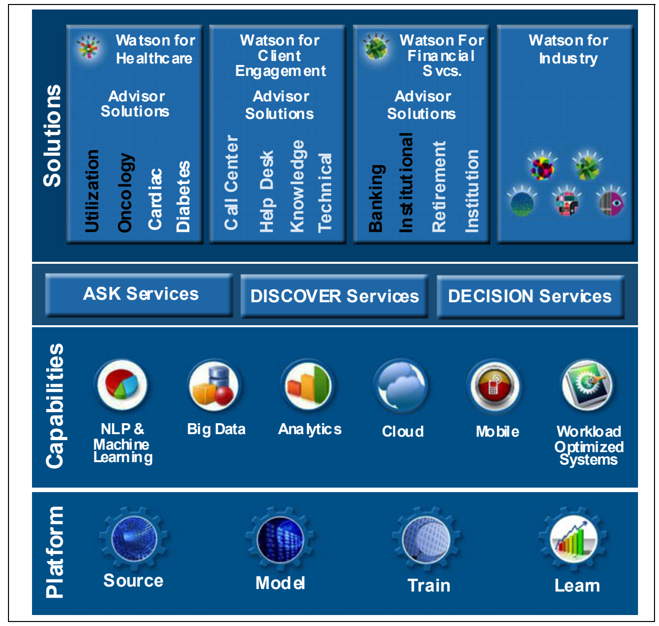

Figure 1 IBM Watson products and infrastructure

Next steps: Leveraging IBM Watson and IBM expert integrated systems

IBM's capability in a diverse range of technical fields such as QA, cognitive computing, and business analytics with workload optimized systems positions IBM as a unique trusted partner to achieve truly intelligent problem solving with optimized platforms.

Your organization can now take advantage of the following IBM solutions, which are powered by Watson:

- IBM Watson Clinical Insights Advisor

  Designed for therapy designers, this solution assists with efficient trials and reduces time to market with new cancer therapies. It can help accelerate research and insights.

- IBM Watson Diagnosis and Treatment Advisor

  Designed for oncologists, this solution assists in identifying individualized treatment options for patients diagnosed with cancer. It can help improve diagnosis and treatments.

[2] http://www-03.ibm.com/press/us/en/pressrelease/39243.wss
[3] http://www-03.ibm.com/press/us/en/pressrelease/40335.wss

- IBM Watson Care Review and Authorization Advisor

  Designed for nurses, this solution streamlines manual review processes between a physician and health plans. It can improve decisions and outcomes.

IBM also offers IBM PureApplication System, a cloud application platform that is designed to simplify and accelerate the deployment of new cloud-based applications. PureApplication System supports a broad range of needs, ranging from traditional enterprise applications to mobile, social, and analytics solutions.

As data grows more and more complex, a simpler approach to managing IT is needed. Expert integrated systems fundamentally change the IT lifecycle by reducing costs, saving time and resources, and speeding innovation for your enterprise. IBM PureSystem combines the flexibility of a general-purpose system, the elasticity of cloud, and the simplicity of an appliance to provide an expert integrated system that has the ability to transform the IT lifecycle.

These IBM smarter solutions show how technology can interact with human experts for learning and enhancing the accuracy, scalability, and speed of decision-making. AI approaches, in addition to analytics techniques, can usher us toward such intelligent systems.

Resources for more information

For more information about the concepts highlighted in the paper, see the following resources:

- *The Era of Cognitive Systems: An Inside Look at IBM Watson and How it Works*, REDP-4955

  http://www.redbooks.ibm.com/redpapers/pdfs/redp4955.pdf

- IBM Smarter Planet solutions

  http://www.ibm.com/smarterplanet/us/en/overview/ideas/

- IBM Analytics for Healthcare

  http://www-01.ibm.com/software/analytics/healthcare/index.html

- *The IBM Next 5 in 5: Our 2012 Forecast of Inventions that Will Change the World Within Five Years*, New York Times, 17 Dec 2012

  http://asmarterplanet.com/blog/2012/12/the-ibm-5-in-5-our-2012-forecast-of-inventions-that-will-change-the-world-within-five-years.html

- *IBM Watson At Your Service: New Watson Breakthrough Transforms How Brands Engage Today's Connected Consumers*

  http://www-03.ibm.com/press/us/en/pressrelease/41122.wss

- *And Now, From I.B.M., Chef Watson*, New York Times, 27 Feb 2013

  http://www.nytimes.com/2013/02/28/technology/ibm-exploring-new-feats-for-watson.html?ref=technology

- IBM Watson Progress and 2013 Roadmap

  http://www.slideshare.net/manojsaxena2/ibm-watson-progress-and-roadmap-saxena

- *IBM Journal of Research and Development*, "This is Watson", Vol.56, No.3/4, May-June 2012

  http://ieeexplore.ieee.org/xpl/tocresult.jsp?isnumber=6177717

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

This document, REDP-4974-00, was created or updated on June 3, 2013.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (or), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

IBM
ILOG
InfoSphere
Rational
Redbooks®
Redguide™
Redpaper
Redbooks (logo)
System z
Tivoli
WebSphere

The following terms are trademarks of other companies:

Other company, product, or service names may be trademarks or service marks of others.