The Era of Cognitive Systems: An Inside Look at IBM Watson and How it Works

IBM WATSON

Redguides
for Business Leaders

Rob High

- Learn how cognitive systems, such as IBM Watson, can transform how organizations think, act, and operate
- Understand the natural language processing capabilities and more of IBM Watson
- See how evidence-based responses can drive better outcomes

Redbooks

Executive overview

IBM® Watson™ represents a first step into cognitive systems, a new era of computing. Watson builds on the current era of programmatic computing but differs in significant ways. The combination of the following capabilities makes Watson unique:

- *Natural language processing* by helping to understand the complexities of unstructured data, which makes up as much as 80 percent of the data in the world today
- *Hypothesis generation and evaluation* by applying advanced analytics to weigh and evaluate a panel of responses based on only relevant evidence
- *Dynamic learning* by helping to improve learning based on outcomes to get smarter with each iteration and interaction

Although none of these capabilities alone are unique to Watson, the combination delivers a powerful solution:

- To move beyond the constraints of programmatic computing
- To move from reliance on structured, local data to unlock the world of global, unstructured data
- To move from decision tree-driven, deterministic applications to probabilistic systems that co-evolve with their users
- To move from keyword-based search that provides a list of locations where an answer might (or might not) be located, to an intuitive, conversational means of discovering a set of confidence-ranked responses

Cognitive systems, such as IBM Watson, can transform how organizations think, act, and operate in the future. This IBM Redguide™ publication describes how Watson combines natural language processing, dynamic learning, and hypothesis generation and evaluation to give direct, confidence-based responses.

What language is and why it is hard for computers to understand

Language is the expression of ideas. It is the medium by which we communicate an understanding of things between people. It is how we convey fear, hope, history, and directions for the future. Some say that it is what we use to think, speculate, and imagine. It is at the base of our cognition, our ability to understand the world around us, or at least at the base of our ability to manipulate and exchange that understanding.

And it is incredibly imprecise.

Our language is full of innuendos, idiosyncrasies, idioms, and ambiguity. We have noses that run, and feet that smell. How can a *slim chance* and a *fat chance* be the same, but a *wise man* and a *wise guy* are opposites? How can a house *burn up* as it *burns down*? Why do we *fill in* a form by *filling it out*?

And yet, it can be amazingly accurate.

We convey so much meaning, and accomplish so much collaboration even in the midst of all the difficulties with language. Somehow, we can see through the gaps, the inconsistencies and contradictions, the irregularity, and the lack of clarity and still understand each other with a great deal of accuracy.

This difference between precision and accuracy is important. *Precision* is the mechanical or scientific exactness that can be found in a passage of text. We can determine whether a specific word exists within a passage with a high degree of precision. *Accuracy* is the degree to which one passage infers that another passage might be considered to be true by reasonable people.

What if, when we said "2 + 2," we meant a car configuration, as in two front seats and two back seats?

The answer to "2 + 2" is precisely 4. Mathematics teaches us this fact. It also teaches us that, regardless of how many zeros you place after the decimal to represent greater precision, the answer always derives to 4. But what if, when we say "2 + 2," we did not mean it to be taken literally as a mathematical formula, but rather as an idiom for a car configuration, as in *two front seats and two back seats*. Or what if a psychologist is using "2 + 2" to refer to a family with *two parents and two children*? In those other contexts, the answer *four* might not be an accurate interpretation of what we are trying to convey in the language.

In fact, to accurately answer a question, you must often consider the available context for the question. Without enough evidential information, it is difficult to accurately respond to a question, even though you can precisely answer elements in the question literally.

Shallow natural language processing

Many natural language systems have attempted to emphasize precision within the confines of specific well-formed rules. For example, sentiment analysis often looks for a set of specific words and their synonyms within a social media site. These systems then, without further assessment of the context in which those words are being used, tally the number of times those words are co-located with some brand in the same phrase. For example, it takes the phrase, "… stopped by the IBM Donut Store for a coffee this morning, it was great …" and then asserts that the collocation of the brand name and the term "great" are an indication of a positive sentiment. However, consider if the rest of the phrase is, "…, it was great to hear that a new Fictional Coffee Shop is opening soon, so I am not tempted to eat donuts every morning." Then, the system might miss that the sentiment is not about the IBM Donut Store. We call this concept *shallow natural language processing (NLP)* because, although it might be fairly precise within its more narrow focus, it is not very accurate.

However, it is also important to realize that shallow NLP actually has an important role in many systems. If your intent is to create a statistically relevant assessment of sentiment trends over huge quantities of information, the lack of accuracy for each individual example is likely not an issue. Assuming that there are approximately as many false-positives as there are false-negatives over a sufficiently large sample set, they cancel each other out. And if the pool of canceled tallies remains relatively constant across sample sets over time, the remaining uncanceled data yields statistically relevant trending information. Thus, the additional processing costs that are required for the additional accuracy for any instance might be unwarranted.

> *Shallow natural language processing can be fairly precise within its more narrow focus, but is not very accurate.*

However, when the individual instances matter, the systems that are designed to be precise without focusing on high levels of accuracy tend to be brittle. That is, they perform well within the narrow parameters of their intended design, but they do not perform well when those parameters change. We liken these systems to using brick-laying construction techniques. Bricks are strong and fairly easy to construct with. For decades and centuries, we refined the brick-laying construction technique to be fairly precise. We were able to build relatively large, ornate, and long-lasting structures. However, although brick buildings have great load strength, they have poor tensile strength. They fall down easily in earthquakes and do not support large spans. And after a certain point, their load strength will fail too.

You can observe these same limitations in some consumer products today. For example, you might use your favorite voice-activated personal assistant and say, "Find me pizza." In return, you get a local listing of pizza restaurants, which is exactly what you wanted. Now you say, "Do *not* find me pizza." You still get back a local listing of pizza restaurants, which is not exactly what you asked for. Likewise, if you say "Find me pizza *nearby*" or "Find me pizza *far away*", the same local listings are returned. The point is that these systems are designed according to a specific set of rules and are looking for specific keyword combinations to determine the answer to produce. These systems do not know how to distinguish between things for which there is no rule. They might be precise, but not necessarily very accurate.

Deep natural language processing

To overcome the limitations of brick building, we shifted to using steel and reinforced concrete for larger buildings. Likewise, we are seeing a shift in construction techniques for natural language processing when accuracy is needed over narrow precision. These techniques incorporate much more context into the evaluation of the question. We refer to this concept as *deep natural language processing*, which is sometimes called *Deep Question-Answering (DeepQA)* when the problem is about answering natural language questions.

> *We are seeing a shift in construction techniques for natural language processing when accuracy is needed.*

IBM Watson is a deep NLP system. It achieves accuracy by attempting to assess as much context as possible. It gets that context both within the passage of the question and from the knowledge base (called a *corpus*) that is available to it for finding responses.

When preparing for the quiz show, JEOPARDY!, Watson was asked the following question (clue) from the category Lincoln Blogs:

"Treasury Secy. Chase just submitted this to me for the third time - guess what pal, this time I'm accepting it."

First, notice the abbreviation, "Secy.", which had to be taken to mean *Secretary*. Further notice that *Secretary* is not meant here to be someone who takes dictation and manages an appointment book. The combined terms *Treasury Secretary* is significant here as a noun and a role. Therefore, to answer this question, Watson had to find a passage that involved submitting and accepting something between Treasury Secretary Chase and Lincoln (the category of the clue). However, also notice that the category does not say "President Lincoln" necessarily. The correct answer turned out to be "What is a resignation?".

When describing this example at an elementary school sometime after the broadcast of IBM Watson playing JEOPARDY!, one fifth grade student offered "What is a friend request?" as a possible answer.

Without context, we would be lost.

The answer from this student is interesting in part because it says a lot about the degree to which social media has permeated deeply into the fabric of the next generation of society. However, it is also instructional because it can also be taken as a fairly reasonable answer to the clue. But, we know that this response is inaccurate because we have historical context. We know that Facebook was not available in the late nineteenth century. Notice that context is what enabled us to increase the accuracy of the system in producing this answer. Without that context, we would be lost.

It is worth emphasizing the point that we, as humans, have little difficulty processing our language, even if we do get confused on occassion. But generally we do much better at resolving the meaning of information that we have written than computers typically do.

We have an innate quality about how we disambiguate language that we want to capture and harness in computing systems. This concept has been a key goal of the artificial intelligence community for the past four decades. And to a large extent, we have been able to increase the precision of language processing. But, it is only with Watson that we can finally break through the level of accuracy that is needed for information systems to function well in the real world of broad natural language.

Also a huge driving force seeks to solve this problem. We are experiencing an explosion of data production. Ninety percent of all the data in the world was produced in the last two years. This trend is expected to grow as we interconnect and instrument more of our world. And 80 percent of all the information in the world is unstructured information, which includes text such as literature, reports, articles, research papers, theses, emails, blogs, tweets, forums, chats, and text messages. We need computers that can understand this flood of information so that we can get more out of it.

IBM Watson understands language

Effective navigation through the current flood of unstructured information requires a new era of computing that we call *cognitive systems*. IBM Watson is an example of a cognitive system. It can tease apart the human language to identify inferences between text passages with human-like accuracy, and at speeds and scale that are far faster and far bigger than any person can do on their own. It can manage a high level of accuracy when it comes to understanding the correct answer to a question.

However, Watson does not really understand the individual words in the language. Rather it understands the features of language that are used by people. From those features, it can determine whether one text passage (which we call a *question*) infers another text passage (which we call an *answer*), with a high level of accuracy under changing circumstances.

In the JEOPARDY! quiz show, Watson had to determine whether the question, “Jodie Foster took this home for her role in ‘Silence of the Lambs’” inferred the answer “Jodie Foster won an Oscar for her role in ‘Silence of the Lambs’”. In this case, *taking something home* inferred *winning an Oscar*, but not always. Sometimes *taking something home* infers a cold, groceries, or any number of things. Conversely, you do not always take home the things that you win. For example, you might win a contract for work, but it is not something you take home.

Context matters. Temporal and spatial constraints matter. All of these concepts add to enabling a cognitive system to behave with human-like characteristics. And, to go back to an earlier point, a rules-based approach might need a near infinite number of rules to capture every case that we might encounter in language.

Effective navigation through the current flood of unstructured information requires a new era of computing called cognitive systems*.*

Watson teases apart the question and potential responses in the corpus, and then examines it and the context of the statement in hundreds of ways. Watson then uses the results to gain a degree of confidence in its interpretation of the question and potential answers.

But we must back up a bit. How does Watson derive its responses to questions? It goes through the following process:

1. When a question is first presented to Watson, it parses the question to extract the major features of the question.
2. It generates a set of hypotheses by looking across the corpus for passages that have some potential for containing a valuable response.
3. It performs a deep comparison of the language of the question and the language of each potential response by using various reasoning algorithms.

   This step is challenging. There are hundreds of reasoning algorithms, each of which does a different comparison. For example, some look at the matching of terms and synonyms, some look at the temporal and spatial features, and some look at relevant sources of contextual information.
4. Each reasoning algorithm produces one or more scores, indicating the extent to which the potential response is inferred by the question based on the specific area of focus of that algorithm.
5. Each resulting score is then weighted against a statistical model that captures how well that algorithm did at establishing the inferences between two similar passages for that domain during the “training period” for Watson. That statistical model can then be used to summarize a level of confidence that Watson has about the evidence that the candidate answer is inferred by the question.
6. Watson repeats this process for each of the candidate answers until it can find responses that surface as being stronger candidates than the others.

Figure 1 illustrates how Watson derives a response to a question.

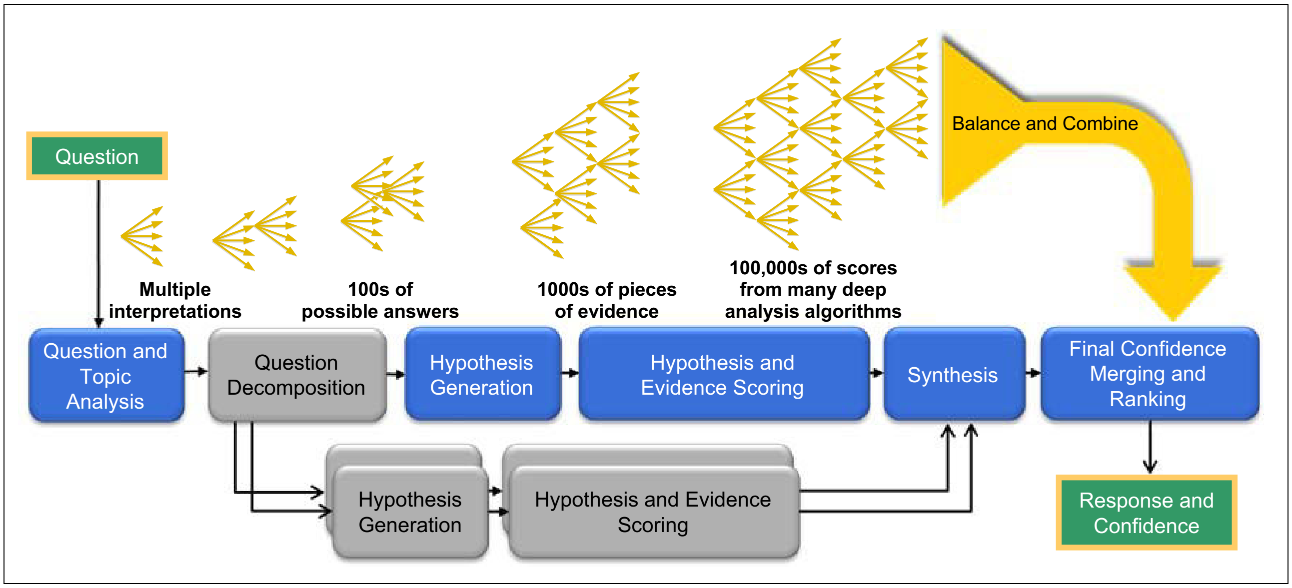

Figure 1 How Watson derives a response to a question

Of paramount importance to the operation of Watson is a *knowledge corpus*. This corpus consists of all kinds of unstructured knowledge, such as text books, guidelines, how-to manuals, FAQs, benefit plans, and news. Watson ingests the corpus, going through the entire body of content to get it into a form that is easier to work with. The ingestion process also curates the content. That is, it focuses on whether the corpus contains appropriate content, sifting out the articles or pages that are out of date, that are irrelevant, or that come from potentially unreliable sources.

Some of the reasoning algorithms focus on spatial and temporal features of the passage, which is critical to disambiguating a tremendous amount of what humans say and write. When we say, “Find me pizza,” it is taken for granted that we mean something nearby. But what is nearby is always relative. In other cases, spatial relationships show up relative to geographic markers, for example, a neighborhood in a city or a state in a country. Likewise, temporal features are also present in the context of much of what we write. When we say, “Get cheese from the store on your way home,” a time frame is inferred. Presumably the writer and the recipient have a shared contextual understanding of when they will be on their way home.

> *Spatial and temporal evaluation has to be performed on both the question and the candidate answer.*

The statement, “In May 1898, Portugal celebrated the 400th anniversary of this explorer’s arrival in India,” demonstrates both spatial and temporal dimensions. The celebration occurred in Portugal, but the event that they were celebrating was the explorer’s arrival in India. Does the statement suggest that the explorer went from Portugal to India? Was he ever in Portugal? Notice that the celebration occurred in 1898, but the event occurred 400 years earlier. So the event actually occurred in 1498. The passage that provided the answer to the question said, “On the 27th of May 1498, Vasco da Gama landed in Kappad Beach.” The spatial and temporal evaluation had to be performed on both the question and the candidate answer passages.

Context is derived from both immediate information and knowledge that is available more broadly. Watson can derive immediate information from the title of a document, other passages in a document, or the source database from where it originated. Context can also come more broadly from a shared history. Remember that we knew that "What is a Friend Request?" was probably an incorrect answer to the clue in the Lincoln Blogs. The reason is because we share a common historical context, which tells us when certain things happened relative to each other. We know that Facebook was created fairly recently, but we know that Abraham Lincoln lived some 150 years ago, well before Facebook became popular. Context and reasoning help us to create a cognitive basis for processing language.

Understanding language is just the beginning

We define cognitive systems as applying human-like characteristics to conveying and manipulating ideas. When combined with the inherent strengths of digital computing, they can help solve problems with higher accuracy, more resilience, and on a massive scale over large bodies of information.

We can decompose a cognitive system as having several key elements (Figure 2). The medium shaded boxes indicate the current capabilities of cognitive systems. The lighter shaded boxes indicate the future capabilities of cognitive systems.

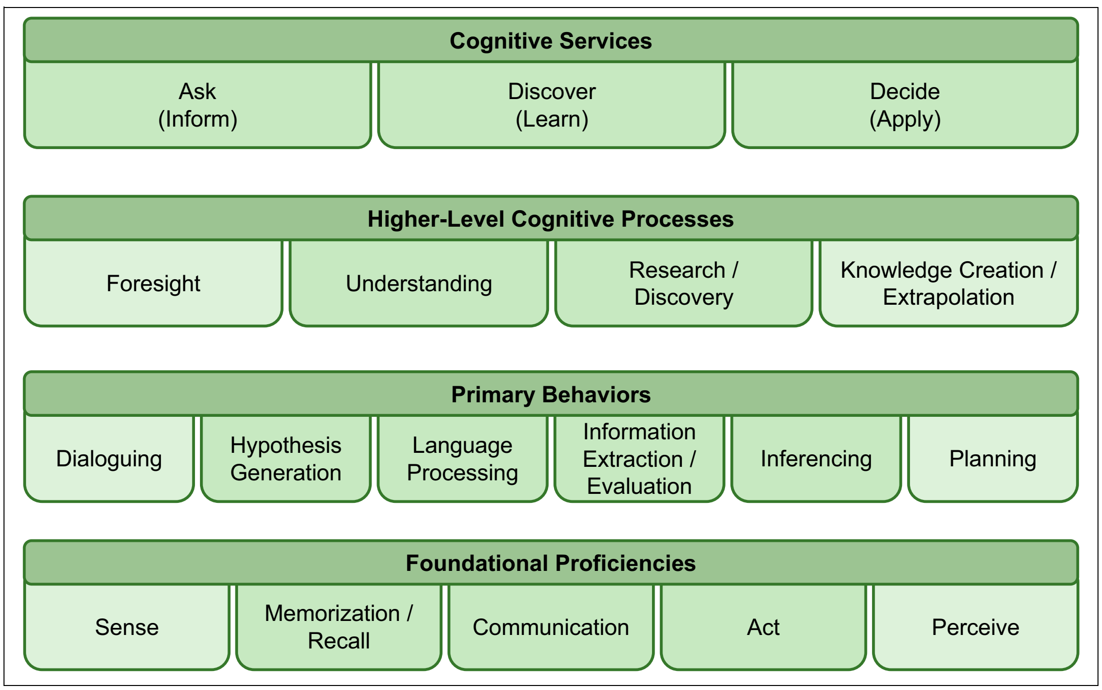

Figure 2 Elements of a cognitive system

Similar to humans, cognitive systems have a way of gathering, memorizing and recalling information, which is the equivalent of human memories. Cognitive systems also have a basic ability to communicate and act. These abilities are organized by certain behavioral constructs such as the following examples:

- The ability to create and test hypotheses
- The ability to tease apart and create inferences about language
- The ability to extract and evaluate useful information (such as dates, places, and values)

These skills are foundational, without which neither computers nor humans can determine the correct correlation between questions and answers.

Higher order cognitive processes can leverage fundamental behaviors to achieve a level of understanding. To understand something requires us to be able to break it apart, down to finer elements that behave in well-ordered ways within a given scale. Just as how things work in physics on human scales is not how things work at cosmic or subatomic scales. Likewise cognitive systems are designed to work at human scales, albeit over enormous collections of humans. As such, understanding language starts with understanding the finer rules of language, not just formal grammar, but the informal grammatical conventions of everyday use.

> *Just as humans do, cognitive systems are driven to understand things by decomposing expressions of an idea and then combining that with context.*

However, as humans do, cognitive systems are driven to understand concepts by decomposing expressions of an idea and then combining the results with context and the probability that certain terms in the passage are being used in a certain way. And, as with humans, our confidence is proportional to the evidence that supports those probabilities and the number of reasoning algorithms that we have available to test our hypotheses.

After we establish a certain level of understanding, decomposing the problem against its probable intent, cognitive systems can recompose the elements in various ways, each of which can be tested to imagine new concepts. These combinations can then be used to drive new discovery and insight, helping us to find answers to questions and to realize the questions that we never thought to ask.

We can then use these capabilities to solve problems that fit certain common patterns. We can ask questions that yield answers. We can use the system to discover new insights and realize concepts that we did not recognize previously. And we can use these systems to support sound decisions or at least to assist people in the decisions that they need to make.

> *As cognitive systems grow richer, they are expected to gain the ability to sense.*

In the future, as cognitive systems grow richer, we expect them to gain the ability to sense. We expect them to do more than to just read text, but to see, hear, and feel so that they have a basic awareness of their environment. And we expect these systems to be able to perceive information, such as to recognize shapes and changing conditions that will further inform their context and ability to infer and reason. We also expect them to adopt higher-order behaviors and cognitive processes, such as to have a dialog, to plan different strategies for solving problems, and to gain foresight and extrapolate it into new knowledge.

In essence, cognitive systems will internalize many of the behaviors that humans find “natural,” and apply them in massive scale to help people solve the problems that today often fall outside their grasp. We are beginning a new era. In this era, computers go beyond just performing routine procedural tasks more efficiently, to employing human-like cognition to make people more intelligent about what they do.

Problems come in different shapes

As we move forward with IBM Watson, we are discovering other uses for it. In the classic “Ask Watson,” a user asks Watson a question (or provides a clue, a patient record, and so on), from which Watson derives a response. Watson has the confidence that the question infers the response and the evidence that supports the response. Watson has found a home in such fields as Oncology Diagnosis, Utilization Management (that is, preapproval of insurance coverage for scheduled medical procedures), Credit Analysis, and basic research.

Watson helps wherever a professional needs assistance in getting the most relevant information to their problem space.

By asking other more important questions, you can begin to think about your business problems in a whole new way.

One of greatest revelations about Watson is that, by using Watson to help answer questions, you might realize that you are fundamentally asking the wrong questions. When Watson responds to your questions, even answering you correctly, you might realize that you need to ask other, better, and more important questions to help consider your business problem in a whole new way. You start to think in ways that help you to understand the competitive threats and opportunities in your marketplace that never occurred to you before.

These discovery type applications are being further improved with work that IBM is doing now in IBM Research and Software Development labs. Recent breakthroughs in inference chaining (determining that *this* infers *that*, which infers something else, and so on) are creating deeper insight. Knowing that diabetes causes high blood sugar is important. However, taking the next step to draw inference between high blood sugar that causes blindness is more critical to caring for the whole patient. These types of multilevel inferences can be captured as an inference graph from which we can observe a broad spectrum of downstream considerations. More importantly, convergence in the graph is a powerful way of deriving more significant inferences, such as answers that can reveal deeper insights and hidden consequences. By coalescing preceding confidence values, we can aggregate and establish higher confidence in an answer as being the preferred answer to the question.

We can produce reverse inferences, in effect, discovering questions to answers that were never asked.

In addition, we can produce reverse inferences, which in effect means that we discover the questions to answers that were never asked. Determining that a patient who has a history of resting tremors and an "unexpressive face" might infer that it has Parkinson's disease. However, determining that the patient also has difficulty walking might further reveal damage to the Substantia Nigra nervous system, which might have been missed without prompting for the previously unasked questions.

IBM is investing in these kinds of major improvements to Watson that we believe are going to lead to further breakthroughs in healthcare, finance, contact centers, government, chemical industries and a more intelligent planet. These types of advances can help propel us into an era of cognitive systems.

In many solutions, Watson is being leveraged with other more traditional forms of computing, such as statistical analytics, rules and business processing, collaboration, and reporting, to solve business problems. For example, consider the idea of IBM combining other statistical analysis with the ability of Watson to answer questions about potential events that can signal a risk for an investment. IBM can help our clients improve their risk and valuation processes for financial institutions. Likewise, insight that we gain about customer responses through deep NLP can suggest changes to buying and usage behavior that otherwise might not be evident in the structured data. In the healthcare industry, Watson is being used to assist insurance companies in their processes to pre-approve treatments as part of their Utilization Management processes.

Accuracy is improved through generalization

As IBM continues to evolve and develop these types of cognitive systems, we must exercise care. We are at a classical juncture that we humans face all the time, which is whether to specialize or generalize. We can specialize NLP technology to one specific domain, concentrating, for example, on just the linguistic features of that domain. This approach is tempting and might even be necessary in the early phases of evolution to ensure the viability of the technology. However, this approach is likely to take us back to the era of building with bricks. If the ability to adapt with human-like dexterity makes cognitive systems special, we must generalize. We need to recognize and draw inferences from a broader set of linguistic variation, under a broader set of circumstances, as our knowledge changes, as the context changes, and as contemporary linguistics change.

By using this approach, we can more readily adapt to new and larger problems. We are already applying Watson to the healthcare and financial services industries, which has the following benefits:

- It brings the advantages of Watson to different domains with high value problems.
- It evolves the language processing algorithms of Watson to handle a broader set of linguistic variation.

This approach enables easier adaptation to other domains and improves the utility of Watson to our existing domain applications.

> *We are at a classic juncture, whether to specialize or generalize.*

Healthcare applications are interesting because they often need precision and accuracy. Accuracy is needed to properly interpret the text in a patient’s health description to infer the patient’s condition. However, the National Comprehensive Cancer Network (NCCN) guideline for breast cancer must be justified by the presence of precise terms in the patient’s health record. Then further accuracy is required to find evidence that supports that treatment.

> *We at the beginning of a new era of computing, one that is less precise, but much more accurate.*

Whenever we encounter a linguistic anomaly (something in the language that we never encountered before), we make a decision about whether the problem is unique to the domain or is characteristic of a broader set of linguistic problems. Wherever possible, we go back to our core algorithms to determine whether we can generalize the algorithm to recognize and evaluate the new situation. As with humans, by using this approach, we can map our understanding to new experiences, and therefore, grow the contextual base for the system.

The expected result is that we increase accuracy, scope, and scale:

- Accuracy of linguistic inferencing (getting the correct answer, for the right reason, in a timely manner)
- Scope of the problem space
- Scaling to massive amounts of data and questions, in many more domains

We expect to see even greater value in “next best action” solutions, social sentiment analysis, petroleum and chemical refinement, and many other applications. We are just at the beginning of a major new era of computing that is less focused on precision, but much more accurate. It is an era of applying human-like behavior to large scale computing problems. It is the era of cognitive systems.

Other resources for more information

For more information about Watson, see the IBM Watson website at:

http://www.ibm.com/innovation/us/watson/index.shtml

The author who wrote this guide

This guide was written by Rob High in collaboration with the International Technical Support Organization (ITSO).

Rob High is an IBM Fellow, Vice President and Chief Technology Officer, IBM Watson Solutions, IBM Software Group. He has the overall responsibility to drive IBM Watson Solutions technical strategy and thought leadership. As a key member of the IBM Watson Solutions Leadership team, Rob works collaboratively with the Watson engineering, research, and development teams. Rob has 37 years of programming experience and has worked with distributed, object-oriented, component-based transaction monitors for the last 21 years, including SOMObject Server, Component Broker, and most recently the IBM WebSphere® Application Server. Rob previously served as Chief Architect for the WebSphere foundation with architectural responsibility for the WebSphere Application Server and the related products that are integrated on that core run time.

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks

- Follow us on Twitter:

  http://twitter.com/ibmredbooks

- Look for us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806

- Explore new IBM Redbooks® publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm

- Stay current on recent Redbooks publications with RSS Feeds:

  http://www.redbooks.ibm.com/rss.html

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 grant 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.

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.

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.

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-4955-00, was created or updated on December 12, 2012.

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 Watson™	Redbooks®	Redbooks (logo) ®
IBM®	Redguide™	WebSphere®

The following terms are trademarks of other companies:

JEOPARDY! is a registered trademark of Jeopardy Productions, Inc. All rights reserved.

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