IBM Watson Health
Thought Leadership White Paper

August 2016

Cognitive Computing: The Future of Population Health Management

- *Cognitive computing 101:* The future of population health management (PHM) will be tied to the rise of cognitive computing, which uses massively parallel processing and artificial intelligence to convert unstructured data into structured data, search the medical literature, and find connections among myriad types of data. Clinicians can collaborate with cognitive computing systems, which learn from experience, to improve health care.
- *Natural language processing:* Cognitive computing has improved the capabilities of natural language processing (NLP) to the point where it can extract insights from unstructured data and correlate relevant studies for clinical decision support. This advanced form of NLP understands the context of language.
- *Beyond big data:* Cognitive computing is a next-generation big data approach that integrates knowledge-driven and data-driven decision support. It uses analytics to find valuable connections across many different kinds of data. It can also be used to compare patients who have particular characteristics with cohorts of similar patients to find out which therapies work best.
- *Population health management:* Cognitive computing can help healthcare organizations understand their populations better. It can provide insights into the non-healthcare factors that affect population health, such as demographics, geographical location, behavioral health, transportation, and socioeconomic status.
- *Genomic research:* As researchers make more progress in analyzing genomic data, cognitive computing will be able to tie markers for certain diseases to many environmental and personal factors that affect an individual's health. These analytics will guide physicians as they begin to provide precision medicine that varies with each patient.
- *Other kinds of data:* Cognitive computing will eventually merge advanced image and textual processing with visual reasoning abilities that can identify the relevant information in images. It has already led to innovative ways of using intensive care unit (ICU) data and is starting to integrate mobile monitoring data with electronic health records (EHRs) while giving feedback to patients.
- *Predictive modeling:* Cognitive computing will improve predictive modeling used in health risk stratification and financial risk management.
- *Patient engagement:* Personalized information and social support tailored by cognitive computing can help patients optimize their health. This technology can also use feedback from medical devices to support people in treating their own chronic diseases.
- *Care coordination:* Cognitive computing can help break down information silos, though it is not a cure for the lack of interoperability among systems. The ability to convert unstructured documents into structured data could improve communication across care settings.

To achieve the Triple Aim of an improved patient experience of care, improved population health, and lower per capita cost,[1] a healthcare organization must transform itself into an entity that is capable of population health management (PHM). Health IT is indispensable to achieving that goal.

It is nearly impossible to ensure that most patients receive recommended services and are engaged in their own care without using analytic and automation tools. And, only a data infrastructure grounded in a big data approach can provide the flexibility, scalability, and comprehensiveness that PHM demands.

But the concepts that are now foundational to big data—handling the large volume and accelerating velocity of healthcare data, combining a wide variety of data types, and accounting for the inconsistency of the data, all within a flexible, scalable framework—are only scratching the surface of what today's ever more powerful computers and computer clusters can do with the right software. These new capabilities have given rise to an approach known as *cognitive computing*, which is expected to play a key role in the future of population health management.

Cognitive computing uses massively parallel processing and artificial intelligence to create a humanlike learning ability, combined with the ability to process millions of data points in seconds. Instead of being preprogrammed to provide particular answers to particular questions, a cognitive computing system draws deductions after searching large data sets, and it learns continuously from the feedback it receives about its conclusions. Though the current analytic approach applies one algorithm at a time, a cognitive computing system is a collection of overlapping, reasoning algorithms that can be expanded and updated.

The latest iteration of cognitive computing offers healthcare organizations a wide array of choices in areas ranging from clinical decision support and clinical quality improvement to medical research tools and analytics for managing population health. Here are some of the new directions explored in this paper:

- *Natural language processing:* About 80 percent of the data in electronic health records (EHRs) is unstructured and therefore unavailable to analytic applications. Up to now, natural language processing has not been successful in recognizing concepts in unstructured text. Cognitive computing improves the accuracy of NLP by understanding language in context. As a result, it can convert unstructured data into structured data and extract insights from it. It can also use NLP to search the medical literature for information relevant to clinical decision-making at the point of care.
- *Genomic data:* Precision medicine is starting to become a real possibility as the cost of genomic sequencing drops. Genomic research generates massive amounts of data that requires big data solutions. Building on those solutions, cognitive computing will vastly increase the amount and types of data that can be combined with genomic information to support precision medicine. Eventually, it will enable physicians to use the results in everyday care.
- *Clinical quality improvement:* In some academic medical centers, cognitive computing systems are being used to improve clinical pathways. Quality improvement teams are starting to feed clinical data back into cognitive computing systems to upgrade those systems' algorithms. The next step is to combine these techniques with insights that cognitive computing gleans from the medical literature. This is already happening in oncology and will spread to other medical fields.
- *Understanding populations:* Big data techniques are being used to aggregate, normalize, and analyze multiple kinds of data and generate actionable reports for providers and care managers. However, many pieces are still missing from our understanding of population health. Cognitive computing can ingest and analyze data from many more sources than existing systems, including remote monitoring data and information on the social and behavioral determinants of health.

The health IT that will support future population health management (PHM) can be compared to a three-legged stool:

- **One leg is the kind of data infrastructure described earlier, which can aggregate and normalize different kinds of data to provide timely reports and answers to clinical questions.** Such an infrastructure must be flexible enough to answer a wide range of queries.

- **The second leg consists of analytic and automation tools that improve patient engagement and coordination across the continuum of care.** These applications can be used both to manage the care of high-risk patients and to engage all kinds of patients in their health care. One aim of these solutions is to ensure that all individuals receive recommended preventive and chronic care. This reduces the likelihood that they will be hospitalized or readmitted to the hospital.
- **Cognitive computing, the third leg, can expand the analytic abilities of current big data techniques while adding entirely new categories of data, including information from relevant studies.**

In addition, cognitive computing can identify previously unknown connections among these different kinds of data. The learning ability of cognitive computing systems and their unparalleled speed and power can support personalized medicine and supply new insights into population health.

Cognitive Computing 101

The big data approach is gaining traction in larger healthcare organizations because of its superior ability to deal with the volume, velocity, and variety of healthcare data. A recent review article noted:

> The increasing scale and availability of large quantities of health data require strategies for data management, data linkage, and data integration beyond the limits of many existing information systems, and substantial effort is underway to meet those needs. As our ability to make sense of that data improves, the value of the data will continue to increase. Health systems, genetics and genomics, population and public health: all areas of biomedicine stand to benefit from big data and the associated technologies.[2]

Following the publication of the first human genome sequence in 2003, the article pointed out, the rapidly expanding genomics field was the initial driver of big data techniques in health care. Since then, however, data scientists have pushed into other fields, and “the use of big data has now reached all areas of health care, biomedical research, and population health,” the paper stated.

Cognitive computing, a branch of artificial intelligence, uses machine learning to expand big data techniques. For example, predictive analytics are typically based on algorithms that are trained through a “supervised learning” approach, in which the outcome is known ahead of time. That approach is fine when only certain kinds of data are being used. But when the algorithm encounters unknown data types and situations, as is frequently the case in health care, it may be unable to make a reliable prediction. That is where cognitive computing has the advantage—its algorithms can make predictions based on the data they have and test them in the real world. As the results are fed back into an algorithm, it is modified to reflect what it has learned. The algorithm becomes more and more accurate as it gains experience.

Another advantage of cognitive computing is its ability to link together many different kinds of data and analyze it in a particular context. This feature can help solve many IT challenges in health care and other industries. As a group of IBM researchers pointed out in a recent paper:

> Conventional programmatic computing, where you program a computer to process a known data set, is not adequate for managing the volume and inconsistency in big data…It requires cognitive computing, using data-centric, probabilistic approaches to data, where, after a fashion, the computer “thinks.”[3]

COGNITIVE COMPUTING FOR POPULATION HEALTH MANAGEMENT

Cognitive computing systems can process massive amounts of data to understand, reason and learn from it, helping providers develop enhanced care plans.

Types of Structured and Unstructured Data

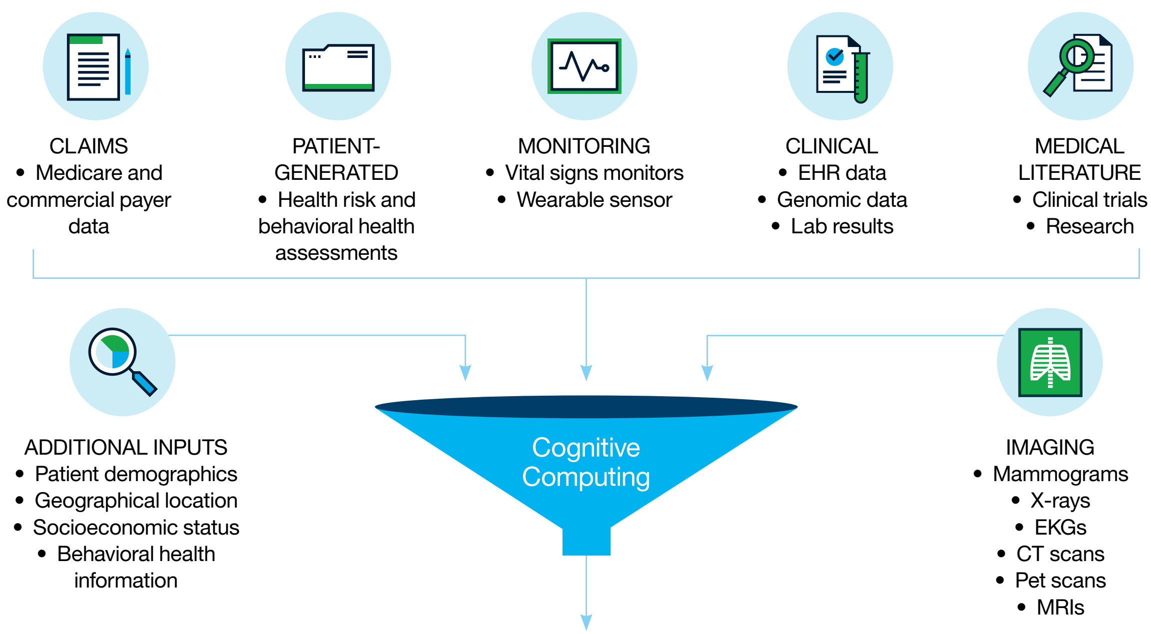

Cognitive Insights and Probablistic Recommendations

- Provides timely insights on populations and individuals
- Proactively identifies at-risk patient groups
- Predicts patient health needs and cost of care
- Uses natural language processing to convert unstructured data
- Supports personalized medicine and clinical decision-making
- Recommends interventions based on probability of success
- Improves patient engagement and communication across care settings (through the use of technologies like mobile health apps and wearable sensors)

Jeopardy!

Cognitive computing arrived with a bang in 2011 when IBM Watson, a learning machine based on a new form of massively parallel processing, won the *Jeopardy!* game on national television.[4] Watson's software, which harnessed the power of a large computer cluster, used probabilistic algorithms to perform natural language processing that was able to understand language in the context in which it was spoken. Though Watson was trained in many different fields of knowledge, it learned far more than it had been taught as it prepared for the *Jeopardy!* game.

Immediately after that victory, which was the culmination of 10 years of research, IBM set out to apply the capabilities of Watson in a number of fields, including health care. Though other cognitive computing systems exist, Watson has had the most practical impact so far, and it is on this system that most of this paper is based.

Recognizing that few companies could afford to purchase Watson outright, IBM made two critical decisions:

- It moved Watson to the cloud, where it could be accessed at any time, from anywhere.
- It "deconstructed" Watson into a group of services that customers could implement, along with application programming interfaces (APIs) that developers could use to write applications for Watson. Some of these developers work for IBM, but an increasing number of entrepreneurial firms are building the Watson ecosystem.

To create a cognitive computing application, a developer must first look at the kinds of questions that target users are likely to ask. The cognitive system must be seeded with a sufficient number of questions and answers to start the machine learning process. Typically, between 1,000 and 2,000 question/answer pairs seem to be the right number to start the process. After that, machine learning takes over as the application seeks data to answer questions and analyzes new data in light of its experience.[5]

A key component of cognitive computing in health care is natural language processing, which it uses to research topics in the medical literature and to search for relevant information locked up in unstructured data. To help clinicians make better decisions at the point of care, cognitive computing integrates knowledge-driven decision support, which draws data from existing knowledge such as journal articles and practice guidelines, with data-driven decision support, which looks for patterns in structured and unstructured healthcare data.

Among the unique features of cognitive computing is its ability to:

- Comb through millions of documents in near-real-time to find particular facts or concepts.
- Use analytics to find useful or illuminating connections among many different kinds of data.
- Compare data sets to find matches or mismatches, which might include patients whose characteristics match those of particular cohorts or missing data in structured EHR fields, respectively.

Cognitive computing can help provider organizations manage population health. By searching text documents rapidly and converting unstructured data into structured concepts, it can help these organizations understand their populations and individual patients much better than they do now. Many non-healthcare factors contribute to health and disease, and cognitive computing can also help providers comprehend the roles that those factors play. After those factors are better understood, it will be easier to engage patients in their own care and improve their health behaviors.

Natural Language Processing

Natural language processing (NLP, for short) has been around for 50 years, but it has not been widely adopted in health care. One reason is the uncertainty about whether its accuracy is sufficient for clinical applications. No matter how accurate NLP becomes, it will not be 100 percent accurate. For some purposes, however, the prevailing accuracy rate—which ranges from 70 percent to around 95 percent—is enough to make NLP highly useful.[6]

There are three types of traditional NLP approaches:

- A *rules-based* system, which searches for patterns in documents, works well for extracting simple measures, such as a blood pressure score or an ejection fraction, but is easy to confuse the system because of variations in terminology.
- A *grammar-based* NLP system considers how terms are used syntactically to map them to dictionaries of concepts, but it doesn't always identify matches.
- A *machine learning* system can learn patterns and excels at finding the most likely matches, but requires extensive training.[7]

Developers of speech recognition applications have tried to use rules-based systems to translate certain medical concepts into discrete data. But these systems have repeatedly stumbled over the issue of context.

For example, if a doctor orders a test to rule out a diagnosis, such as cancer, the word cancer refers not to a diagnosis but to the reason for the test. Yet NLP might not notice the difference. If a visit note says that a patient stopped smoking three days ago, that doesn't mean the person is a nonsmoker, though the program might interpret it that way.[8] Rules-based NLP programs also have difficulties with the "negation" problem. In this type of scenario, a physician may say that the patient does not have a symptom associated with a particular health condition, yet the application interprets that statement as a positive indication.[9]

A cognitive computing system combines the grammar-based and machine learning approaches into a flexible, comprehensive tool for searching unstructured text. It has the ability to understand what language means, including its nuances. It also has the computational power to search large amounts of text quickly to find matches of terms used in a query and rank them by confidence level. According to IBM, Watson's NLP accuracy rate is in the mid-90 percent range.

Unstructured EHR data

The structured data in an EHR can provide the basic facts about a patient's health status and medical history, if the correct data has been entered in the appropriate fields. But the narrative of the patient's complaint, the care plan, and observations that don't fit into EHR templates are usually documented in free text, which is unstructured. In fact, some physicians simply dictate their findings, documenting nearly all their encounters as free text that is then transcribed by a person or a speech recognition program.

Unstructured documents contain many of the most important facts about a medical case, noted an IDC Health Insights paper:

> In many cases, unstructured text remains the best option for providers to capture the depth of detail required, for example, in a clinical summary, or to preserve productivity by incorporating dictation and transcription into the workflow. Unstructured text records contain valuable narratives about a patient's health and about the reasoning behind healthcare decisions.[10]

That being the case, it is essential to make this unstructured data accessible to analytics. An advanced NLP system can mine unstructured data so that it can be analyzed along with structured EHR and claims data. "This powerful set of capabilities will serve as the basis for a new generation of information access and analysis systems that improve population health, discover factors for readmission, predict infection outbreaks, or improve quality measures," the IDC study pointed out.

Seton Healthcare Family, a Texas healthcare provider that operates 38 facilities serving 1.9 million people, has used a solution that combines NLP with predictive analytics to collect and analyze structured and unstructured data. In a study it conducted with IBM, Seton Healthcare Family developed a method to identify which patients with congestive heart failure were likely to be readmitted. Among the 18 top predictors were factors that are not usually documented in EHRs as structured data elements, including drug and alcohol use and living arrangement (whether the patient lives alone or not, for example).[11]

One large EHR vendor has incorporated IBM's NLP and content analysis software into its system. The EHR is able to convert unstructured notes into structured data for clinical decision support. It can also code some of this data for billing purposes.[12]

Carilion Roanoke Community Hospital in Roanoke, Virginia, has this kind of EHR, and the hospital used the EHR's NLP capability to "read" millions of notes. In the course of 18 months, the insights that Carilion extracted from the unstructured text enabled its clinicians to increase their ability to identify patients who had heart failure and had been previously undiagnosed.[13]

By analyzing unstructured data, cognitive computing can also find omissions in the structured record, such as diagnoses that have not been entered into a patient's problem list. Cleveland Clinic has leveraged Watson for this purpose, comparing physician and hospital admission notes with diagnoses listed in the EHR. "When the unstructured data is compared to the problem list from the EMR [electronic medical record], all too often omissions are identified," noted a recent book on cognitive computing.[14]

Medical literature

Advanced forms of NLP can also be used to search the medical literature for evidence and insights that support clinical decision-making. Moreover, cognitive computing can integrate such knowledge-based analytics with data-driven analytics that look for patterns in real-world data.

This kind of approach reflects an ongoing change in how researchers are using EHR data in clinical decision support. Rather than simply trigger alerts about a person's allergies or medications, EHR data can now be used — in combination with insights from the literature — to provide support for how a physician thinks about diagnosis and treatment in individual cases.

This method starts with a statistical analysis that matches a particular patient with a large cohort of patients who have similar characteristics and were treated in various ways. By combining an analysis of those patients' outcomes with treatment guidelines based on studies that show the value of certain interventions for patients in clinical trials, the cognitive computing system can give physicians more information than they could find through a literature search alone. This combination of case-based and knowledge-based information "can be used for a range of decision support applications, including guideline adherence monitoring and personalized prognostic predictions."[15]

As in any approach based on machine learning, Watson requires extensive training within a particular medical specialty to reach this point. Oncologists at institutions such as Memorial Sloan Kettering and MD Anderson Cancer Center are already taking advantage of Watson's decision support capabilities.[16,17] In primary care, it will take longer to train Watson because of the large number of different health problems that primary care physicians encounter every day. But the cognitive computing system recently took a big leap forward by "attending" Cleveland Clinic's medical school and "learning" everything that the medical students were taught.[18]

Data Types

Many diverse factors are converging to propel health care — including population health management — into the future. Genomic research is exploding; mobile health apps and devices are proliferating; imaging data is growing exponentially and becoming more liquid; social media is being scanned for clues to health behavior; and research on the social and behavioral determinants of health is revealing new opportunities for patient engagement.

Cognitive computing has the potential to harness all these factors in population health management. The research to support the integration of some kinds of data, such as genomic data and social media, is still in an early phase. In other areas, such as mobile health data, providers are not yet ready to use the information. But cognitive computing is already starting to analyze imaging data and will be ready to accommodate other kinds of data when it is available.

Genomic data

Researchers have made impressive strides in analyzing genomic data, which requires a big data approach. A lot of number crunching is required in order to compare genotypes with the phenotypes of individual patients. Scientists do this to match particular patient characteristics and disease states with variations in the patient's DNA that make him or her different from other people.[19]

The significant drop in genomic sequencing costs, coupled with the widespread adoption of EHRs, has made this approach feasible, but it's still in an early phase. A more direct approach to harnessing genomic discoveries in clinical care has been embraced by a number of cancer centers. These institutions sequence cancer tumors and use cognitive computing to analyze the patient factors involved. They also employ the technology to search the literature for information about treatments for that particular type of cancer. Based on its analysis of all that data, IBM Watson has been asked to provide an additional source of input for expert discussions about how to treat particular patients.[20]

In the long run, as genomics begins to yield its secrets in fields other than cancer, cognitive computing will provide the key for translating those advances into approaches that doctors can use in everyday medicine. At the same time, genomics will become integrated with case-based and knowledge-based analytics to provide more personalized care and a deeper understanding of patient populations.

Imaging data

NLP can extract meaning from unstructured text found in journal articles and EHR notes, but other kinds of unstructured data can also contribute to better medical decisions, including imaging data. A great deal of important information is available in X-rays, EKGs, mammograms, CT scans, PET scans, and MRI images.

Cognitive computing can use pattern recognition, coupled with clinical training, to "read" these images. Watson, for example, is being trained on a library of 30 million X-ray images. After the system is trained, it will be able to compare a patient's X-rays with clinically normal images to detect disparities. The radiologist can agree or disagree with its findings, but the computer can perform the readings much faster and with fewer errors than a radiologist who reads thousands of images a day.[21]

Moreover, cognitive computing can combine image readings with other data to monitor the results of specific interventions. For example, the system can tell how a cancer patient is doing by tracking the patient's tumor regression and can correlate that outcome with the patient's drug treatment, lifestyle, and, someday, genetic makeup. In addition, the system can compare the images with those of millions of similar patients to see how they differ and offer explanations for those variations.

Specifically, this is how IBM researchers described their plans for imaging data:

> We are developing a cognitive system with advanced textual, image processing, and visual reasoning abilities that is able to gather information from across an enterprise and identify important elements in the data to help clinicians in their decision-making processes. It is designed as a new layer above PACS [picture archiving and communication systems] and EHR management systems, which combines imaging data along with textual and structural patient data.[22]

Monitoring data

Monitoring data comes from vital-signs monitors in hospitals and other healthcare facilities, as well as from home monitoring and mobile monitoring—including wearable sensors.

In hospitals, the automatic documentation of vital signs has begun to spread from ICUs to medical wards, allowing both real-time analysis and analysis of data over time. With training and feedback from clinical teams, cognitive computing can be used to provide more reliable alerts regarding certain conditions than traditional methods can. By combining the monitoring data with structured and unstructured clinical data, the system can also produce insights that improve the quality of care.

For example, a neonatal department in a Toronto hospital used cognitive computing to predict 24 hours in advance which babies might develop late-onset neonatal sepsis. The key to their success was the analysis of monitoring data streams over time.[23]

ICU data is normally thrown away after a day because there is so much of it. A 50-bed ICU, for instance, may generate a quarter of a terabyte of data per month. By storing this data in the cloud and analyzing it with the help of cognitive computing, clinicians could someday prevent complications such as adult sepsis, hospital-acquired infections, and respiratory, cardiac, and neurological events.[24]

Remote patient monitoring provides the other major opportunity here. Home monitoring is increasingly being used to track the vital signs of high-risk patients, such as those with congestive heart failure. The widespread use of smartphones, meanwhile, has created a universe of mobile health apps. Though most of these apps focus on fitness and wellness, some are used in tandem with specialized devices to measure health conditions such as diabetes (glucometers), hypertension (blood pressure monitors), and heart disease (electrocardiogram [ECG] strips and heart rhythm monitors).[25]

Most physicians are not reviewing this data yet, nor are they prescribing mobile health (mHealth) apps to patients. They would argue that they are not being paid for this activity (except in the Medicare Chronic Care Management program), that most mHealth apps haven't been validated, and that the amount of data that remote monitoring could generate would be overwhelming in any case. A reliable method of screening this data must be developed before mHealth will become a routine element of clinical care.[26]

Nevertheless, some companies are already delivering applications that can help patients understand their own monitoring data and improve their health. For example, Apple's HealthKit allows patients to collect their mobile monitoring data in one place, and some EHR vendors have integrated their platforms with HealthKit, allowing providers to give feedback to patients.[27]

Non-healthcare data

As explained earlier, health care is just one of many factors that affects and determines a person's health. According to the World Health Organization, healthcare accounts for 10 percent to 25 percent of the variance in health over time. The remainder is related to genetic factors (up to 30 percent), health behaviors (30 percent to 40 percent), social and economic factors (15 percent to 40 percent), and environmental factors (5 percent to 10 percent).[28]

People with multiple health and social needs are high consumers of healthcare services. High-risk patients generate roughly 80 percent of healthcare costs. And social determinants such as homelessness, substance abuse, physical disability, and economic factors add massive costs to the healthcare system.[29]

There are many available sources of data on the social and behavioral determinants of health. What all of this means and how it relates to clinical and genetic data has yet to be worked out in detail. But cognitive computing offers methods to integrate and analyze this information to help clinicians engage patients in their own health care. In addition, this wide-ranging assessment of health determinants can aid healthcare organizations in understanding their patient populations.

Population Health Management

Cognitive computing brings a new dimension to population health management. Though the data lake approach is essential for crunching administrative, clinical, claims, and patient-generated data quickly and efficiently, cognitive computing can add insights from many other sources, including the medical literature, genomic data, and social and behavioral information. Unifying all these disparate areas of knowledge can help healthcare organizations understand and manage population health.

For starters, cognitive computing brings a new power of statistical analysis to the task of describing and analyzing a specific patient population. Just as a cognitive computing system can be used to compare a patient's characteristics and the treatments the patient has received to those of many similar patients with the same health condition, such an approach can also be harnessed to compare a population or subpopulation to similar patient populations. Moreover, matching is not limited to how the population is constituted in demographic and disease categories. It can include myriad other factors that have been shown to be important in health.

These kinds of comparisons may indicate that a particular subpopulation — diabetic patients who are on Medicaid, live alone, and have poor diets because of limited access to fresh food, for example — may benefit from a particular kind of intervention that has worked in other places. Or cognitive computing may reveal that patients who have hypertension and who receive texted prompts to avoid salty foods have better outcomes than those who aren't prompted.

Some innovative healthcare systems are already starting to meld their clinical and administrative data with insights from geographical information systems. For example, Carolinas HealthCare System, based in Charlotte, North Carolina, began developing such a matrix a decade ago. Using census-based demographic data and geocoding software in conjunction with data on emergency room visits, the organization devised a method for "using information about patients' patterns of healthcare utilization and community-level data to identify areas where access to primary care services was limited and could be improved."[30]

This kind of data could be integrated with other attributes of patients who recently visited a healthcare system's emergency department to alert care managers that certain high-risk patients, for example, had no access to primary care or had disabilities that made it difficult to get to a clinic. These are the kinds of insights that a cognitive computing system could provide — in a more comprehensive, scalable way than any healthcare organization could do on its own — to support population health management.

Predictive modeling

Cognitive computing can mine vast amounts of data and process that data in near-real-time, using algorithms that learn from experience. The sweet spot of cognitive computing is its ability to do this with a large number of different data streams and to correlate its insights from that data.

Predictive modeling is an essential part of the data infrastructure for PHM. Predictive modeling, in theory, should benefit from having more kinds of data available for analysis. So, cognitive computing should be able to improve the risk stratification of a population and the ability to predict the overall disease burden and health costs of that population. In addition, it could help care managers prioritize the patients who need help the most at any given time.

Patient engagement

PHM cannot succeed without patient engagement. Although human contact is a necessary part of this strategy, especially for high-risk individuals, automation tools are required in order to scale patient engagement to entire populations. Cognitive computing takes automation to the next level while preserving the human touch.

For example, the Denver firm Welltok uses a Watson-based application, CafeWell Concierge, to provide personalized information and social support that helps individuals optimize their health. Designed for use by population health managers, health plans, and employers, CafeWell Concierge integrates large amounts of internal and external data. It uses natural language processing, machine learning, and analytics to provide personalized recommendations and answers to questions posed by health consumers.

Each CafeWell Concierge user receives an Intelligent Health Itinerary based on the user's health insurance benefits, health status, preferences, interests, demographics, and other factors. This document is a personalized action plan that includes resources, activities, health content, and condition-management programs. Users can also ask questions related to health and wellness. The cognitive application is trained to answer open-ended questions in seconds. It can also learn which rewards have the most effect in motivating individuals to change their health behaviors.[31]

Cognitive computing can also use feedback from medical devices to support people in treating their own chronic diseases. For example, IBM Watson Health has done some research with device manufacturer Medtronic on how patients with diabetes control their blood sugar levels. The initial investigation showed that, by running data from home glucometers through analytic models, it was possible to forecast that certain patients would likely become hypoglycemic a couple of hours before that occurred. If these patients could be alerted in real time and could do something about their risk factors, it would be more effective than having them show their doctor three months of data and try to remember what was going on when they last crashed.[32]

Mobile health apps and wearable sensors are the new frontier for patient engagement. By using Apple HealthKit, for example, health consumers can aggregate their mHealth data from multiple devices, and patients and care managers can receive insights about health problems from the cognitive analysis of the integrated data streams. The results can also be fed back into the Watson cloud so that it can learn which interventions produced positive results.[33]

One further step is required to make this mobile health data truly useful in population health management: A care manager who receives the HealthKit data analysis must be able to translate it into action. To do so, the care manager can use complementary applications that identify high-risk patients and help engage them in improving their health behavior. In addition, mobile data analyses can be used to inform automated online campaigns that educate and engage patients who have moderate chronic diseases or are healthy and need only preventive care. In fact, these campaigns could be conducted in a feedback loop with smartphone users.

Care coordination

Cognitive computing does not magically make disparate healthcare applications interoperable. But it can break down information silos by aggregating and normalizing data from healthcare providers across care settings.

For example, a cognitive computing system that has access to multiple clinical information systems may discover that a patient who sees a primary care physician during flu season has already had a flu shot at a CVS pharmacy. Though claims data could supply the same connection, that information is much less timely than data that comes directly from provider information systems.

There are still many obstacles to gaining access to this data from all care settings. For example, post–acute care facilities such as nursing homes and home care agencies use information systems that are far less capable of delivering discrete clinical data than are hospital or ambulatory care EHRs or pharmacy systems.34 But cognitive computing's use of NLP to scan unstructured documents, coupled with its normalization of aggregated data, could supply the missing links.

Workflow integration

If doctors and nurses have to leave their EHR workflow to obtain additional patient information or clinical decision support from a cognitive computing system, they are less likely to use it. So, the insights from cognitive computing must be available within the clinicians' workflow; the user interfaces must be intuitive and seamless; and clinicians should be able to get answers to their questions at the point of care.

Fortunately, there are already population health management solutions that can provide this kind of support within EHR workflow, along with automation tools that make the information actionable. When combined with the cognitive computing insights from structured and unstructured data culled from many different sources, these PHM solutions can present providers and care managers with reliable, timely analyses to support improved care delivery and patient engagement.

As mentioned earlier in this paper, Watson's NLP capabilities and some of its analytics have been integrated with one of the leading EHRs. Besides enlarging the corpus of data available to clinicians, the integrated system is also being used to accelerate clinical trial recruitment. With the help of EHR data, Watson's clinical trial matching application allows physicians to identify patients during office visits who fit the criteria for particular studies. Since only 5 percent of cancer studies in the United States recruit sufficient participants, this solution could help drive medical research forward at a faster pace.[35]

Other EHR-integrated applications of cognitive computing are expected to play a role in quality improvement and disease management. Again, it is already possible to tailor treatments to particular patients by comparing the outcomes of many other similar patients who were treated in various ways. When genomic science progresses further, cognitive computing will also be able to present physicians with best practice guidelines that apply to patients with a particular genetic makeup, along with other specific characteristics.

Conclusion

Population health management has become increasingly dependent on big data because of the volume, variety, and velocity of information that healthcare organizations must manage to ensure that everybody receives appropriate care. But though many organizations are still struggling in this stage, it is only the first step on the road to a vastly more capable way to apply big data in health care.

Cognitive computing can use NLP to extract insights from the unstructured data that holds most of the relevant information about patients. It can rapidly search the medical literature and correlate what it finds with clinical and other kinds of data on a patient. It can also aggregate and analyze an extensive assortment of data, including imaging, monitoring, and genomic information as well as data on the behavioral and social determinants of health.

Cognitive computing marries data-driven analytics — using clinical, administrative, claims, and other data — with knowledge-driven analytics, which are applied to medical studies and clinical guidelines. In this dynamic process, the artificial intelligence of cognitive computing constantly learns and improves from the feedback it receives in the real world.

By identifying health problems that are missing from clinical data, providing clinical decision support, and predicting the health status of individual patients, among other tasks, cognitive computing can become a dependable and essential assistant to physicians and to organizations engaged in population health management. But it will never replace clinicians. As one paper stated:

The collaboration between human and machine that is inherent in a cognitive system supports a best practices approach that enables healthcare organizations to gain more value from data and solve complex problems.[36]

Returning to the analogy of the three-legged stool, the population health management of the future will include two legs that some organizations are already using:

- Data aggregation, normalization, and analysis based on a flexible, scalable data lake approach
- PHM tools that automate care coordination, care management, and patient engagement

The third leg of the stool is cognitive computing, which can provide much more advanced and timely insights, using a far wider range of data than the analytic tools that are now prevalent in PHM.

Together, these three legs can support vastly improved health care at a lower cost and a better patient experience, thus meeting the criteria of the Triple Aim.

Endnotes

1. Donald M. Berwick, Thomas W. Nolan, and John Whittington, "The Triple Aim: Care, Health and Cost," *Health Affairs*, May/June 2008, 759-769.

2. F. and K. Verspoor, "Big Data in Medicine Is Driving Big Changes," *Yearbook of Medical Informatics*, 2014;9(1):14–20.

3. Martin S. Kohn, Jimeng Sun, Sarah Knoop, Amnon Shabo, Boaz Carmeli, and Daby Sow, IBM Research, "Health Analytics and the Transformation of Health Care," unpublished working paper.

4. John Markoff, "Computer Wins on 'Jeopardy!': Trivial, It's Not," *The New York Times*, Feb. 16, 2011, accessed at www.nytimes.com/2011/02/17/science/17jeopardy-watson.html?pagewanted=all&_r=0

5. Judith Hurwitz, Marcia Kaufman, and Adrian Bowles, *Cognitive Computing and Big Data Analytics*, Chapter 11: "Building a Cognitive Healthcare Application," 184. Indianapolis: John Wiley & Sons, Inc., 2015.

6. Leonard D'Avolio, "6 Questions to Guide Natural Language Processing Strategy," *InformationWeek*, Feb. 12, 2013, accessed at www.informationweek.com/big-data/big-data-analytics/6-questions-to-guide-natural-language-processingstrategy/d/d-id/1108629.

7. Ibid.

8. Michael Weiner, IBM Watson, personal communication.

9. Eduardo Blanco and Dan Moldovan, "Some Issues on Detecting Negation from Text," Proceedings of the 24th International Florida Artificial Intelligence Research Society Conference, accessed at www.aaai.org/ocs/index.php/FLAIRS/FLAIRS11/paper/download/2629/3031.

10. Susan Feldman and Judy Hanover, "Unlocking the Power of Unstructured Data," IDC Health Insights white paper, June 2012, accessed at public.dhe.**ibm.com**/software/data/sw-library/ecm-programs/IDC_UnlockingThePower.pdf.

11. Ibid.

12. IBM press release, "Doctors' Notes Turn into EMR Insights with Natural Language Processing Software," Feb. 19, 2014, accessed at www-03.**ibm.com**/press/us/en/pressrelease/43232.wss.

13. Weiner and IBM press release, "IBM Predictive Analytics to Detect Patients at Risk for Heart Failure," Feb. 19, 2014, accessed at www-03.**ibm.com**/press/us/en/pressrelease/43231.wss

14. "Building a Cognitive Healthcare Application," 193…

15. "Health Analytics and the Transformation of Health Care"…

16. Bruce Upbin, "IBM Watson Hits Daily Double Fighting Cancer with Memorial Sloan Kettering," *Forbes*, March 22, 2012, accessed at www.forbes.com/sites/bruceupbin/2012/03/22/ibm-watson-hits-daily-double-fighting-cancer-with-memorial-sloan-kettering/#466f22fb582c.

17. MD Anderson Cancer Center press release, "MD Anderson Taps IBM Watson to Power 'Moon Shots' Mission," Oct. 18, 2013, accessed at www.mdanderson.org/newsroom/news-releases/2013/ibm-watson-to-power-moon-shots-.html.

18. Ken Terry, "IBM's Watson Hits Medical School," *InformationWeek Healthcare*, Nov. 2, 2012, accessed at www.informationweek.com/healthcare/clinical-information-systems/ibmswatson-hits-medical-school/d/d-id/1107199.

19. Terry, "Big data makes a difference at Penn Medicine," Oct. 8, 2015, CIO.com, accessed at www.cio.com/article/2990428/healthcare/big-data-makes-a-difference-at-pennmedicine.html.

20. Kathy McGroddy, IBM Watson, personal communication.

21. "Health Analytics and the Transformation of Health Care," and Weiner, personal communication…

22. "Health Analytics and the Transformation of Health Care"…

23. "Building a Cognitive Healthcare Application," 177…

24. "Health Analytics and the Transformation of Health Care"…

25. Terry, "Remote patient monitoring: fulfilling its promise," *Medical Economics*, Sept. 3, 2015, accessed at medicaleconomics.modernmedicine.com/medical-economics/news/remote-patient-monitoring-fulfilling-its-promise.

26. Ibid.

27. Ibid.

28. Edward Blatt, Eloise O'Riordan, Ljubisav Matejevic, and Martin Duggan, IBM Cúram Research Institute, "Addressing Social Determinants and Their Impact on Health Care," February 2013, accessed at www.longwoods.com/articles/images/SocialHealth_IBM.pdf.

29. Ibid.

30. Greg Goth, "GIS, Boots on Ground Prove Powerful Healthcare Combination," Health Data Management, July 21, 2015, accessed at www.healthdatamanagement.com/news/GIS-Boots-on-Ground-Prove-Powerful-Healthcare-Combination-50918-1.html?portal=community-health.

31. "Building a Cognitive Healthcare Application," 187–188…

32. McGroddy, personal communication.

33. IBM press release, "IBM Watson Health Cloud Capabilities Expand," Sept. 10, 2015, accessed at www-03.**ibm.com**/press/us/en/pressrelease/47624.wss.

34. Terry, "EHR Interoperability With Long-Term Care Providers Wanted, but Who Will Pay?" *iHealthBeat*, July 30, 2015, accessed at hpi.umich.edu/news/ehr-interoperability-long-term-care-providers-wanted-who-will-pay.

35. Jane E. Brody, "The Road to Cancer Treatment Through Clinical Trials," *The New York Times*, March 23, 2015, accessed at well.blogs.nytimes.com/2015/03/23/the-road-to-cancer-treatment-through-clinical-trials.

36. "Building a Cognitive Healthcare Application," 179…

IBM Corporation
Route 100
Somers, NY 10589

Produced in the United States of America
August 2016

HLW03028-USEN-01