Smarter Healthcare
Built on Informatics and Cybernetics

Fred A. Maymir-Ducharme[1], Lee Angelelli[2] and Prithvi Rao[3]
[1]*IBM US Federal CTO Office; University of Maryland University College (UMUC), Maryland, U.S.A.*
[2]*Analytics Solution Center Architect, IBM US Federal CTO Office, Maryland, U.S.A.*
[3]*Federal Healthcare Executive Architect, IBM US Federal CTO Office, Maryland, U.S.A.*

Keywords: Healthcare, Knowledge Management, Advanced Analytics, Contextual & Cognitive Computing, Watson.

Abstract: Applying advanced analytics to big data in healthcare offers insights that can improve quality of care. This paper focuses on the application of health informatics in care coordination, payment, wellness, and healthcare decision management. Cognitive computing and analytics can be used to capture and extract information from large volumes of disparate medical data. Applying natural language processing, probabilistic computing, and dynamic learning can achieve intelligent healthcare systems that users can interact with to drive business and medical insights across patient populations and result in greater patient safety, care quality, wellness, and improvements in payer programs. As is the case for most organizations with large and disparate data sets, the ability to manage information across the enterprise becomes extremely challenging as the size and complexity of the knowledge management infrastructure grows. Interconnecting healthcare systems and applying advanced cognitive analytics and health informatics would provide medical organizations, clinicians, and payers with the information they need to make the best treatment decision at the point of care. The authors address related research in the following areas: image analysis for anomaly detection; electronic healthcare advisors for clinical trial matching and oncology treatments options; advanced models and tools that can accelerate geo-spatial disease outbreak detection and reporting.

1 INTRODUCTION

The healthcare industry over the past century has evolved into a series of independent providers and processes, focused on intervention rather than prevention, while devaluing primary care and population health. The problem is that many "healthcare systems" are not run as "systems" at all. There is little coordination of data, care, or services. The industry is so fragmented, in fact, that in 2010 economists ranked healthcare the least efficient industry in the world, with more than $2.5 trillion wasted annually (IBM, 2012). Physicians have been on information overload for decades, contributing to the estimated 15 percent of diagnoses that are inaccurate or incomplete (Harvard Business Review, April, 2010). Population growth, increased life expectancy, an increase in chronic diseases across an aging population, health laws, and a lack of trained clinicians exacerbate the present healthcare situation. An over demand of patients, under supply of healthcare workers, and financial incentives that reward volume over value are increasing the complexity and inefficiencies of healthcare, which results in higher costs per patient and degraded quality. Among healthcare executives interviewed for the 2010 Global CEO study, 90 percent expected a high or very high level of complexity over the next five years, but more than 40 percent said they were unprepared to deal with it (IBM, 2010).

This paper addresses leading-edge technology advancements in the field of Health Informatics, Analytics and Cybernetics. The advancements in Health Informatics technologies to analyze Big Data, detect patient medical risks, share reliable health and operational information, and extract tacit knowledge out of biomedical and healthcare data in combination with explicit knowledge to intervene with specific treatments across the care continuum is a fundamental game changer for healthcare organizations, which are challenged to increase their quality of care, improve clinical outcomes, and at the same time reduce costs.

Today's Healthcare Knowledge Management

Maymir-Ducharme F., Angelelli L. and Rao P..
Smarter Healthcare - Built on Informatics and Cybernetics.
DOI: 10.5220/0005261705250533
In *Proceedings of the International Conference on Health Informatics* (HEALTHINF-2015), pages 525-533
ISBN: 978-989-758-068-0

Systems (HCKMS) require a multi-disciplinary approach to meet an organization's need to innovate and thrive. Gotz, Stavropoulos, Sun and Wang describe the use of a predictive risk-assessment analytics system that analyzes large volume of heterogeneous medical data to provide various forms of risk analytics across the patient population. By applying text and data mining techniques, pattern recognition, and cognitive computing techniques on large volumes of clinical and genomic data, smarter health informatics and analytics can drive knowledge discovery that have a direct impact on patients at the point of care (Gotz et al, 2012).

Polanyi described tacit knowledge as implicit information that is difficult to capture linguistically. (Polanyi, 1966) This challenge is further exacerbated when attempting to automate the capture and process of healthcare tacit knowledge. Nonaka and Takeuchi created a variety of models as a means of capturing and communicating (Nonaka and Takeuchi, 1995).

Healthcare organizations are under constant pressure to cut costs while improving patient care, clinical outcomes, and reducing hospital re-admission rates. IBM's Patient Care and Insights solution offers the ability to analyze a wealth of structured and unstructured data from medical records and by using similarity analytics can auto-select a cohort of clinically similar patients for further analysis (Gourlay, 2002). The solution leverages healthcare annotators, natural language processing, and analytics to track trends, patterns, anomalies, and deviations; it produces a personalized proactive care solution that can reduce hospital readmission rates, predict the onset of ailments based on a patient's social and behavioral patterns, ensure drug treatment effectiveness, and help in aligning physician and care teams to patients.

With the explosion of data in terms of volume, velocity, variety, and veracity, it is imperative and time critical to have the ability to analyze a high number of healthcare data points and correlate the data in real time. Clinicians need timely and accurate information so that they can make real time decisions and proactively intervene to save human lives, shorten hospital stays, reduce hospital readmission rates, and improve the overall quality of care. Medical devices that monitor patients' health conditions can generate streams of data related to heart rate, blood oxygen saturation, and respiratory rates. The ability to fuse real-time streaming data from various sources and apply analytics can help address and prevent many life threatening conditions. Health Informatics that leverages and processes data streams in real time with an advanced analytic engine can provide the ability to quickly ingest, analyze and correlate information from thousands of real-time resources to predict and save lives in a timely fashion across intensive care units (Dollard 2013). Big Data and analytics will play a major role in the future of healthcare and will have a direct impact on the medical practitioners and their care of patients. With the onset of aging population, managing chronic illness and conditions is becoming a major factor in healthcare. Providers are under pressure to deliver better and safer care at lower cost with increased transparencies in the quality of care and outcomes.

At the same time, payers including health insurance companies and the US Government are under increased pressure to monitor the accuracy of claims and quickly detect a fraudulent claim. Historically, a pure 'pay and chase' model has been used to combat healthcare fraud. This model is expensive to execute and thus not sustainable. A more reliable solution would be to identify fraudulent claims on entry by parsing the claim text, looking for trends and anomalies that depict fraudulent patterns, and flagging these claims early in the payment workflow cycle for review and analysis. Here cognitive computing in the form of Watson Policy Advisor can play a significant role in ingesting claim policies and guidelines. Going through domain adaptation and training to detect fraudulent patterns can result in significant savings in cost, time, and resources for the payer, provider, and the patient.

Healthcare providers are expected to be current on the latest medical research and advancements. Physicians are expected to be aware of the 360 view of the patient's medical history, allergies, drug treatments, and clinical and lab work when they walk into the consulting room. This puts tremendous pressure on the practitioners' time as they need to not only discern their patient's history and analyze test results, but also maintain a current medical knowledge base. (David Dugdale et.al., 1999).

Search is a fundamental element of a HCKMS. With the growth of multimedia over the past decade, new search and analytics systems have evolved that provide the ability to search through pre-processed metadata and automatically analyze files for the search criteria (e.g., image recognition, speech-to-text). These advanced search techniques enable the capture of tacit knowledge in non-linguistic forms. Gourlay's knowledge management (KM) framework stresses the value of non-verbal modes of information, such as behaviors and processes, to convey a variety of perspectives (Gourlay, 2002.

There is no single Knowledge Management System (KMS) that integrates all of the evolving KM capabilities into a single store, and provides a single search and analytics user interface; hence, the most effective KMS' today act as a federation of KM systems.

Search and analytic capabilities have evolved considerably over the past few decades. Search was the initial focus, since one must logically "find" the data before one can analyze it. The early work focused on a variety of "relevancy ranking techniques" (e.g., key word proximity, cardinality of key words in documents, thesaurus, stemming, etc.) and federated search, which provides the ability to search across a variety of DBMS, XML-stores, file stores, and web pages from a single search interface. In parallel, but at a slower pace, analytics capabilities have followed suit in their evolution by extending traditional search engine reverse-indices to store and exploit search criteria and other metadata captured during previous searches or data crawls.

A variety of analytics have been developed that extend traditional business intelligence and data warehousing techniques to provide predictive analysis, and even stochastic analysis. Predictive models depend on a large corpus of data and can be calibrated by modifying the model attributes and the parameter ranges. Stochastic modeling applies a variety of techniques to discover new aggregations of data and other data affinities that can later be used to augment the organization's predictive models. These techniques rely on capabilities such as object-based computing, contextual computing, and cognitive computing.

The remainder of this paper elaborates on the concepts relevant for HCKMS, including Predictive Analytics, Stochastic Analytics, and technology like Cognitive Computing. In addition the paper documents the use of these technologies in exploitation of tacit knowledge in HCKMS. IBM's WATSON (Ferrucci, 2012) is built on these technologies, and the company's research continues to exploit these capabilities for the healthcare industry.

2 HEALTH CARE INFORMATICS

There are numerous analytic concepts and technologies one can apply to HCKMS. Traditional Data Warehousing and Business Intelligence (BI) technologies, such as Master Data Management and crawling or indexing data sources, provide the ability to integrate multiple data sources into a single repository and lexicon to better understand the data an organization has in their KM repositories and thus better understand and manage the associated assets, products, and services.

The next level of maturity is embodied by Predictive Analytics, which can leverage data models, such as entity and relationship maps, to predict current or future events, meet objectives, and identify risks such as fraud and market changes. Predictive analytics models are typically represented as aggregations of data in models that generally represent "observable insights." When sufficient data is available and mapped to a model, leading a level of assurance, an organization is able to make predictions. Organizations in mature industries with a large corpus of empirical data can effectively calibrate their predictive models by changing the models' attributes and ranges – and then assessing the accuracy and fidelity of the model (e.g., "false-positives" and "false negatives") against the corpus of data.

Stochastic modeling goes beyond the deterministic aspects of predictive modeling by introducing non-determinism to create or discover new models. Stochastic approaches generally include affinity link analysis (e.g., petri-net or semantic modeling) that extends predictive models with new data clusters (e.g., new associations that co-occur, are temporally or geodetically associated, or are statistically related by an additional data element or parametric such as data attributes and ranges).

2.1 Cognitive Analytics

Advanced analytics has been a primary focus of the authors for the last decade. Having previously worked on traditional DBMS On-Line Analytical Processing and a variety of Data Warehousing projects prior to 2000, the new millennium challenge has been unstructured data and multimedia. (Maymir-Ducharme and Angelelli, 2014).

This section describes cognitive analytics, which the authors define to include contextual analytics. Contextual analytics is used to assess information within a confined set of data sources, e.g., a set of 200 documents resulting from a federated search query and ingest. Contextual analytics applies techniques such as relevancy ranking, entity extraction, entity-relationship modeling, parts of speech tagging, etc., to analyze the data from a federated search. Cognitive analytics extends the scope of the analyses to include "implicit knowledge" and perspectives that may be represented by lexicons, taxonomies, models, or rules-based computations tailored to the areas of interest represented by the data in the KMS. These frames of reference are then used to build understanding and insights about the explicit knowledge in the KMS. Cognitive Analytics builds on four key capabilities: collection of data, contextual computing, cognitive computing, and exploitation of results (cybernetics).

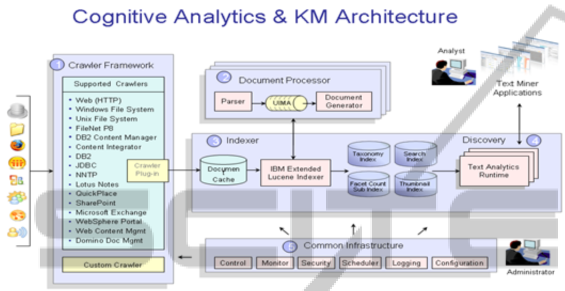

Figure 1: Cognitive Analytics Architecture.

Figure 1 illustrates the functional architecture of IBM's Content Analytics (ICA), a Cognitive Analytics solution. The architecture above does not reflect networks, physical or logical views, or storage, which is traditionally a combination of DBMS for the structured data and a content management system for the unstructured text and multimedia. The Crawler Framework provides the ability to identify and ingest data from a variety of sources and in various formats. ICA uses an IBM search engine that extends the open source Lucene indexer for efficient storage of the information content, metadata, and other analytical results.

The first phase of cognitive computing involves ingesting and indexing a data set. General indexing may include metadata such as URL, URI, document name, document type, date stamp, and date indexed. The original data and documents may or may not become persistent in the KMS depending on organizational policy, storage costs, and legal data rights.

The "Document Processor components provide the contextual analytics capabilities. The second phase of cognitive computing involves applying various analytics to individual documents, and ultimately across all documents. This includes the ability to implicitly or explicitly extract features and relationships from diverse data and data sources, creating metadata to continuously build and update context around the data elements. The figure above illustrates the "stand-alone" ability to apply a unique set of analytics on a select document or group of documents, which is at times is a valuable capability outside of the broader analytical framework. The Search and Index API (SIAPI) is used to analyze one or more documents, resulting in the "annotations" illustrated as output from the SIAPI.

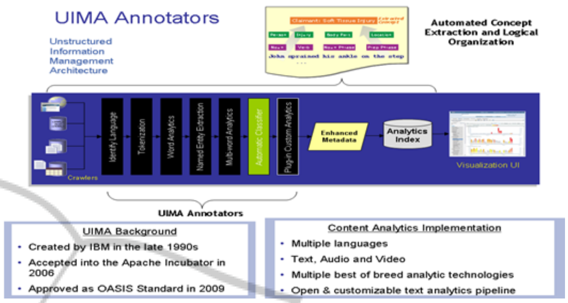

Figure 2: UIMA Framework.

The UIMA framework illustrated in Figure 2 provides the core contextual analytics through a variety of annotators. Annotators support an array of analytical processing capabilities, such as language identification, entity extraction, entity type extraction, parts of speech tagging, tokenization, machine translation, speaker identification, and tagging. The annotations are used for both contextual analytics and cognitive analytics.

The UIMA framework can be used for structured and unstructured data – and for multimedia. IBM has created numerous translingual annotators that enhance search, e.g., supporting transliterations, foreign character sets in UTF-16, multi-lingual search, and relevancy ranking algorithms such as stemming and polymorphic analysis.

Figure 3 illustrates some of the results of text analytics annotators. The keywords are labeled and color coded to facilitate user's finding the terms of interest within a document. The "Relations in this text" also identified relationships between the entities extracted. Behind the scenes tokenizers recognize multiple spellings of the same name (refining associative analytics), and parts of speech tagging recognizes pronouns and is therefore able to include name-relationships associated with an individual. Applying these analytics resulted mapping "she" and "her" to Robin Cook.

IBM has extended these text analytics to support multiple foreign languages, including English, Russian, Chinese, and Arabic. The company's translingual technology includes two types of machine translation (MT): Rules Based MT, and Statistical MT. The search engine supports searching in English, in a foreign language, or in multiple languages, and it can search foreign language transliterations in the native language vernacular and foreign character sets.

To apply text analytics to multimedia, one uses multimedia annotators such as:

- Language Identifier to identify one of 128 languages within the first three phonemes;
- Speech to text translations to convert speech into text and then apply the appropriate text analytics;
- Speaker Identifier to identify and distinguish multiple speakers in an audio or video clip and tag conversations appropriately;
- Speaker Authentication to authenticate a speaker if their voice has been enrolled into the system as an identity's voice.

An important observation is that this system now enables users to search multimedia natively – as opposed to the traditional limitations of static tagging techniques – and apply the same analytics, e.g., implicit knowledge and discovery, to any combination of structured text, unstructured text, images, audio, and video – concomitantly.

The Analytics Server in Figure 1 illustrates the ability to extend the text and multimedia analytics previously discussed, to support lexicons, taxonomies, and other analytical models. The "text miner" in the analytics server uses a multi-element mapping structure to create analytic models. An element can be an entity, an entity type, a relationship, or an attribute, e.g., tagged or derived metadata. Various analytic tools provide an entity-relationship (E-R) model that one can then visualize a number of individuals and their direct and transitive relationships graphically. One can then create a variety of models (e.g., entity social network models, hierarchical, organizational, risk/threat, etc.) to represent an organization's predictive models or categorization schemes. Valuable attributes such as time, geo-location, and demographics can be used to provide unique analytic models and visualizations.

Experience shows that the quality of the lexicon and taxonomy created for the KMS directly impacts the soundness of the resulting models and analyses. These multi-element mappings are processed by the text miner and result in an XML tabularized list that can be used for analysis or as input to a variety of visualizations.

There are numerous types of visualizations one can derive from the multi-element mapping results (models). ICA includes Entity - Relationship Model Views, Entity-type – Relationship Model Views, Hybrid Entity-type – Relationship Model Views, Categorization Model Views, and traditional Measurement and Metrics Views, e.g., Pie charts. Another view is Facets, which represent different views a user can define to visualize either search results (e.g., views based on a taxonomy) or analytics results (e.g., multi-element mappings viewed as models. ICA Facets present the results of a search query that is mapped to the user's categorization scheme. It provides a quick glance at the taxonomy nodes where the majority of the search results fall, and other related keywords from the lexicon that were not explicitly in the search query. Facets provide the ability to model various clusters of data elements and help intuitively guide the prioritization, evaluation, and focus areas of search results and analytics.

Recognizing the speed with which analytics technology is emerging and maturing, a KMS Analytics solution needs to be extensible and provide interfaces for new inputs – as well as output to other systems. ICA has the ability to input and crawl a variety of data sources. The system also provides the ability to export data and results to other analytic components, either in an XML document or directly into a relational DBMS. There are three different export opportunities in the ICA architecture:

1. Crawled data can be exported before indexing or performing any other text analytics;
2. Indexed data can be exported before applying multi-element mappings and other analytics; and
3. Analyzed data can be exported to be used by different visualization technologies, or additional analytic engines.

Section 1 discusses stochastic analysis, which is not supported by this Cognitive Analytics solution. One approach to stochastic analytics could be to export the predictive models from ICA into another modeling engine such as SPSS Modeler, which includes petri-net and semantic models that could be used to discover new aggregations, thereby extending the exported multi-element mappings and relevant metadata.

2.2 Stream Processing

There are many other data attributes to consider, and one more in particular is worth noting. A new paradigm has evolved over the last decade. Stream computing differs from traditional systems in that it provides the ability to analyze **data in motion (on the network)**, rather than **data at rest (in storage)**. The parallel design of stream processing and the "pipe and filter" architecture allows a variety of analytics to be applied to data in motion (Maymir-Ducharme, 2013). Stream processing dynamically supports analytics on data in motion. In traditional computing, you access relatively static information to answer evolving and dynamic analytic questions. With stream processing, one can deploy an application that continuously applies analysis to an ever-changing stream of data before it ever lands on disk – providing real-time analytics capabilities not possible before.

Stream computing is meant to augment current data at rest analytic systems. The best stream processing systems have been built with a data centric model that works with traditional structured data and unstructured data, including video, image, and digital signal processing. Stream processing is especially suitable for applications that exhibit three characteristics: compute intensity (high ratio of operations to I/O), data parallelism allowing for parallel processing, and ability to apply data pipelining where data is continuously fed from producers to downstream consumers. As the number of intelligent devices gathering and generating data has grown rapidly alongside numerous social platforms in the last decade, the volume and velocity of data that organizations can exploit have mushroomed. By leveraging Stream Processing practitioners can make decisions in real-time based on a complete analysis of information as it arrives from monitors and equipment (measurements and events) as well as text, voice transmissions, and video feeds.

2.3 IBM Watson System

Some organizations are using KMS based on Cognitive Analytics to capture, structure, manage, and disseminate knowledge throughout their organization and thus enabling employees to work faster, reuse best practices, and reduce costly rework from project to project. Society has now entered an era where the complexity of our world and the risks thereof demand a capacity for reasoning and learning that is far beyond individual human capability.

Watson takes the Cognitive Analytics technology described in Section 2.1 to a new level of innovation, adding advanced natural language processing, automated reasoning, and machine learning to the Cognitive Analytics components, e.g., information retrieval, knowledge representation, and analytics. Watson used databases, taxonomies, and ontologies to structure its knowledge, enabling the processing of 200 million pages of unstructured and structured text (stored on four terabytes of disk storage.) Watson used the UIMA framework as well as the Apache Hadoop framework to support the required parallelism of the distributed system, which consisted of ninety IBM Power 750 processors and sixteen terabytes of RAM.

Figure 3 illustrates the various components of the Watson architecture. More than 100 different techniques and technology were used to provide natural language analytics, source identification, hypothesis generation and discovery, evidence discovery and scoring, and hypotheses merging and ranking.

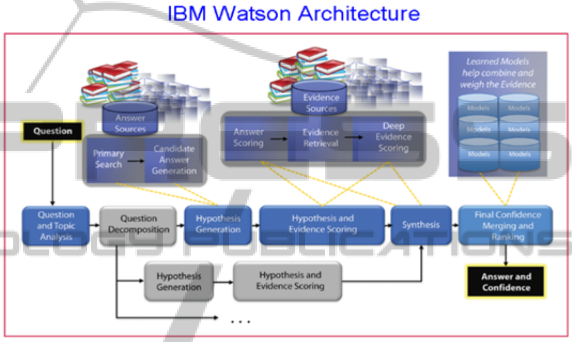

Figure 3: IBM Watson Architecture.

IBM is leveraging its Watson technology to create the concept of Cogs computing, which is designed to follow and interact with people and other cogs and services inside and across cognitive environments. A "cog" represents a specific frame of reference and the associated data. IBM is using Cogs computing to create **"Industry of Knowledge"** expert advisors.

3 HEALTH CARE CYBERNETICS

Watson can be used to analyze structure and unstructured Healthcare data, such as medical journals, patient medical records, physician notes, lab results, and produce confidence weighted and evidence based responses to queries. In the process, Watson will become progressively smarter and help the healthcare industry reduce cost and deliver quality healthcare.

3.1 Watson Oncology Case Study

Memorial Sloan-Kettering Cancer Center (MSKCC) is using Watson to help Oncologists battle cancer.

Traditionally, Oncologists diagnose cancer using a patient's chart, x-rays, laboratory data, a few medical books; and they might then recommend either the general radiation therapy or three types of chemotherapy. Today, Oncologists face a perpetually growing sea of data in their efforts to effectively deal with every aspect of their patients' care. The associated medical information doubles every five years. An Oncologist must maintain awareness and understanding of the ever expanding medical books and articles; electronic patient and family medical records; more than 800 different cancer therapies sequencing 340 cancer tumors (each with multiple mutations), 20,000 genes, correspondence with more than 1,000 physicians. Traditional processes for cancer prognosis and the recommendation of therapies are no longer able to effectively harness all of the available data. Hence physicians are turning to Watson to develop precision based medicine in cancer.

MSKCC and IBM are training Watson to compare a patient's medical information against a vast array of treatment guidelines, published research, and other insights to provide individualized, condensed, scored recommendations to physicians. Watson's Natural Language Processing capabilities enable the system to leverage this sea of unstructured data, including journal articles, multiple physicians' notes, as well as the guidelines and best practices from the National Comprehensive Cancer Network.

The evolving IBM Watson Oncology Diagnosis and Treatment Advisor includes supporting evidence, and with every suggestion, it provides transparency and assists physicians with decision-making and patient discussions. Watson interactively points out when more information is needed and updates its suggestions as new data is added (IBM, February 2013).

For example, MSKCC had a cancer case involving a 37 year-old Japanese non-smoking patient who was diagnosed with lung Adenocarcinoma cancer. The physician asked Watson for a recommend therapy. Watson's initial case analysis recommended Chemo-Erlotinib treatment at a 28% confidence interval. Watson needed more information and recommended the physician perform a molecular pathology test to detect if there were any EGFR mutations because Watson's knowledge base indicated that 57% of all EGFR mutation in women with Adenocarcinoma cancer are missed. The Lab results identified the presence of an EGFR exon 20 mutation. Watson referenced a medical paper citing an exception, where the EGFR exon 20 mutation did not respond to Erlotinib treatment. Analyzing the new lab information along with the medical article, Watson then recommended Cisplatin / Pemetrexed treatment with 90% confidence. "There are only about two or three physicians in the world who would know this information" said Dr. Jose Baselga (Physician-in-Chief at MSKCC) at the IBM Watson Group Launch in New York Event (9 January, 2014).

3.2 Watson Medical Insurance Case Study

WellPoint Inc., is the largest for profit managed health care company in the Blue Cross and Blue Shield Association. WellPoint identified that a significant amount of time, and monetary resources were not being used optimally to deliver quality care to their constituents. For example, WellPoint needed to examine and pre-approve a high number of medical claim forms before a patient could be treated. The pre-approvals required WellPoint claim processors to process a high volume of claims containing both structured and unstructured data, assimilate and coordinate across many external systems to track and validate claims across clinical and patient data, which was a labor intensive and time consuming process. Long wait times frustrated both the provider and patients, especially when time sensitive information and approvals were requested. (IBM, June 2014). WellPoint explored rigid code based procedures and claim processing guidelines and discovered that they too had limitations and inefficiencies.

WellPoint turned to Watson to leverage its power to easily handle natural language processing to streamline healthcare insurance pre-approvals for patients and providers. WellPoint trained Watson with 25,000 historical cases to help their claim processors with tools to streamline their pre-approval claims processing and assist them with an automated decision making process. Watson on its part ingests the plethora of treatment guidelines, policies, medical notes, journals and documents and using its natural language processing capability processes the unstructured data such as text based treatment requests, leverage its hypothesis and evaluation processes to provide confidence-scored and references backed evidences to help WellPoint make better and more efficient decisions about the treatment authorization requests. This iterative process enables Watson to improve iteratively as payers and providers use it in their claims processing and help in the proper decision making outcome to either approve or deny a pre-authorization claim for medical procedures. The new system provides responses to all requests in seconds, as opposed to 72 hours for urgent pre-authorization and three to five days for elective procedures with the previous UM process. Elizabeth Bingham, Vice President of Health IT Strategy at WellPoint Inc., stated, “We know Watson makes us more efficient and is helping us turn around requests faster: It also ensures we are consistent in our application of medical policies and guidelines.” (IBM, June 2014)

3.3 Watson Neonatal Care Case Study

IBM is applying Streaming computing and advanced analytics such as machine learning to develop state-of-the- art multimodal medical informatics that combine analyses of multiple information sources, e.g., text, image, molecular to help intensive care units (ICU) improve prediction, prevention and treatment of diseases. IBM Stream computing and advance analytics are making great advancements in the development of an instrumented, interconnected, and intelligent ICU platform that integrates human knowledge with medical device data and enables real-time automated holistic analysis of patient’s vital signs to detect early subtle patient’s warning signs before symptoms become critical and life threatening. This is another example of leveraging advanced analytics and cybernetics to achieve smarter healthcare for humans.

ICU clinicians review outputs from multiple medical devices to continuously monitor patient vital signs such as blood pressure, heart rate, and temperature. The medical devices are designed to issue a visual and/or audio alert when any vital sign goes out of the normal range, which notifies clinicians to take immediate action. The medical devices produce massive amounts of data: Up to 1,000 readings per second are summarized into one reading every 30 to 60 minutes. Some data is stored for up to 72 hours before being discarded (Gotz et al, 2012). Even the most skilled nurses and physicians have difficulty reading all the medical devices monitoring a patient’s vital signs to detect subtle warning signs that something is wrong before the situation becomes serious. “The challenge we face is that there’s too much data,” says Dr. Andrew James, staff neonatologist at The Hospital for Sick Children (Sick Kids) in Toronto. “In the hectic environment of the neonatal intensive care unit, the ability to absorb and reflect upon everything presented is beyond human capacity, so the significance of trends is often lost.” (Gotz et al, 2012)

It is difficult to detect early signs of hospital-acquired infections. IBM T.J. Watson Research Center’s Industry Solutions Lab (ISL) is working with hospitals and medical research labs including the University of Ontario Institute in Project Artemis to develop a first-of-its-kind, stream-computing platform that ingests a variety of unstructured monitoring-device data *(instrumented)*; integrates clinical/medical knowledge from multiple sources *(interconnected)* for automated analysis using advance machine learning algorithms to detect subtle patient vital signs *(intelligent)*; and alert hospital staff to potential health problems before patients manifest clinical signs of infection or other issues. The Project Artemis brings human knowledge and expertise together with medical device-generated data captured in machine learning algorithms to produce a holistic early detection patient analysis. The Project Artemis stream computing platform capabilities, which enable nurses and physicians to detect medical significant events even before a patient exhibit symptoms, will enable proactive treatment that can increase the success rate and potentially saving lives. Physicians at the University of Virginia can detect nosocomial infections 12 to 24 hours in advance before “preemies” exhibit symptoms. (IBM, December 2010)

4 CONCLUSIONS

This paper provides an overview of advanced analytics technology such as cognitive computing that can be applied in support of Healthcare Informatics. The interactive Watson system demonstrates the use of cognitive analytics and illustrates various Healthcare Cybernetic automations as medical and insurance decision support systems

Although the current Watson emphasis is on "validated" data & sources from mature industries, there is further research focused on broadening the types of data and applying cognitive computing in other mature industries as well as in less stable industries such as entertainment.

ACKNOWLEDGEMENTS

The authors wish to thank Joe Angelelli, Ph.D., Director, Health Services Administration at Robert Morris University and Susan Carlson, MBA, PMP for providing technical consulting.

REFERENCES

"Redefining Value and Success in Healthcare" IBM Healthcare and Life Sciences. February 2012. http://www.ibm.com/

Capitalizing on Complexity: Insights from the 2010 IBM Global CEO Study" IBM Institute for Business Value. May 2010. http://www.ibm.com/ceostudy.

David Gotz et al "ICDA: A Platform for Intelligent Care Delivery Analytics", American Medical Informatics Association Annual Symposium (AMIA), Nov 3, 2012.

Wes Rishel et. al 'IBM's Patient Care and Insights: A Promising Approach to Healthcare Big Data'. Gartner,Research, 27 November 2012.

David C. Dugdale, Ronald Epstein, Steven Z. Pantilat 'Time and the Patient-Physician Relationship', Journal General Internal Medicine, Volume 14, January 1999.

Polanyi, M. (1966) *The Tacit Dimension*, London: Routledge & Kegan Paul.

Nonaka, I. and Takeuchi, H. (1995) *The Knowledge-Creating Company*, Oxford: Oxforbd University Press.

Gourlay, Stephen (2002) "Tacit Knowledge, Tacit Knowing or Behaving," *3rd European Organizational Knowledge, Learning and Capabilities Conference;* 5-6 April 2002, Athens, Greece.

Gal, Y. et al. "A Model of Tacit Knowledge and Action," IEEE Computational Science and Engineering, 2009 International Conference p. 463-468.

Maymir-Ducharme, Fred and Angelelli, Lee "Cognitive Analytics: A Step Towards Tacit Knowledge?" proceedings of the 8th International Multi-Conference on Society, Cybernetics and Informatics, July 2014.

Maymir-Ducharme, F. and Ernst, R. "Optimizing Distributed and Parallel TCPED Systems," US Geospatial Intelligence Foundation (USGIF) Technical Workshop, Denver CO, July 17-19, 2013.

Ferrucci, D. "Introduction to 'This is Watson'," IBM Journal of Research and Development, Vol 56 May/July 2012.

Memorial Sloan-Kettering Cancer Center IBM Watson Helps Fight Cancer With Evidence-Based Diagnosis and Treatment Suggestions, Feb. 8, 20132/8/1313, http://www.ibm.com/services/multimedia/MSK_Case_Study_IMC14794.pdf.

IBM Watson – WellPoint Inc, IBM Watson Platform. June 2014 http://www.ibm.com/smarterplanet/us/en/ibmwatson/assets/pdfs/WellPoint_Case_Study_IMC14792.pdf.

Vince Dollard, Emory University Hospital Explores 'Intensive Care Unit of the Future', Emory University Health Sciences, Woodruff Health Science Center, Nov 11, 2013.

Shirley S. Savage, 'How Watson and Cognitive Computing are changing healthcare and customer service', IBM Systems Magazine, October, 2013.

IBM Corporation - University of Ontario Institute of Technology. December 2010 http://www.ibmbigdatahub.com/sites/default/files/document/ODC03157USEN.PDF.