



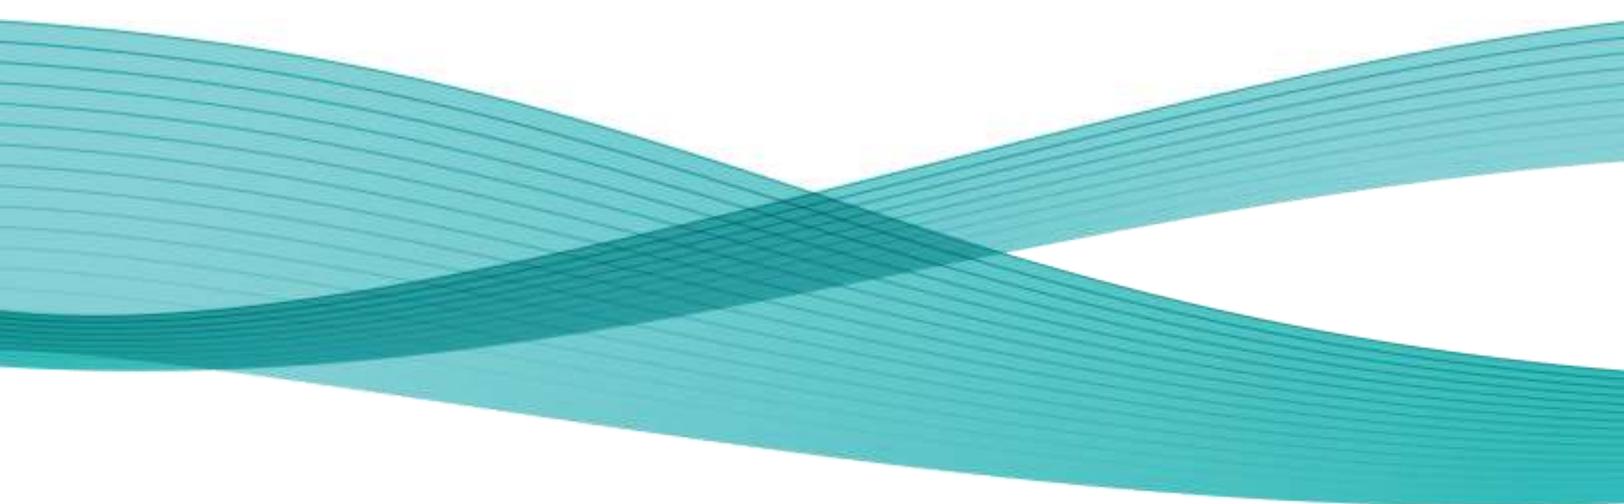
Big Data Analytics

Federal Business Analytics

April 2013

James Bessin, Vice President

Anand Das, Vice President and Chief Technology Officer



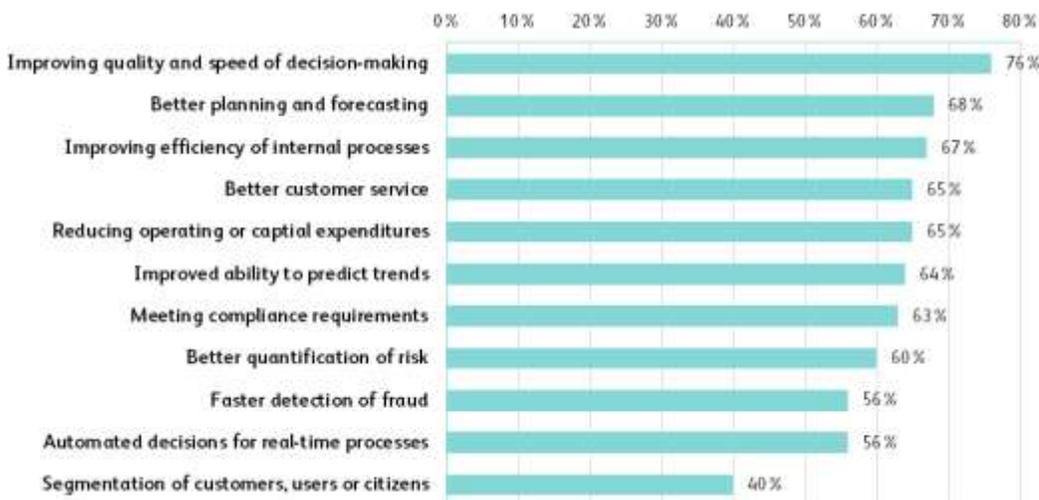
Federal Agency Data Challenges

Federal data is growing at an astounding rate, expected to double every two years. However, data collected from a wide variety of sources such as blogs, emails, videos, social media, photos and other types of sensors often are unused. What makes their analysis difficult is their volume, the velocity with which they arrive, their variety, and the validity of their pedigree over the entire data life cycle. Xerox understands the challenges and the opportunities posed by the explosion in data, and inherent value data analytics brings to produce new insights. By using natural language processing, graph analytics, Hadoop, machine learning, and predictive analytics, agencies can realize the latent value potential sitting in large document stores and other operational datasets.

To benefit from the value of data in their possession, agencies must extract relevant information and aggregate or translate it into information relevant to increasing automation, optimizing business processes, improving efficiency, or boosting productivity. As a leader in operational excellence providing business process outsourcing to thousands of commercial and government clients, Xerox is well poised to help agencies harness the potential value in their data.

Federal Data Analytics Opportunity

Petabytes of information are accumulating across government. Military and civilian personnel records, veterans benefit and health data, genomics data, economic reports, centuries of climate records, decades of stock trades and financial reports, health records for hundreds of millions of individuals, and even the results of NASA experiments in space sit in various federal agencies. The era of “Big Data” has arrived in government just as it has across numerous other sectors of business. Digital documents, transactions, intelligence, photos, video, Web content, server logs, and electronic correspondence are filling storage systems to the brim. At the same time, IT budgets are flat, agencies are being pressed to consolidate data centers, and agencies want to move from managing IT infrastructure to securing and managing data to support agency missions. The President’s Council of Advisors on Science and Technology has decreed that every Federal agency needs to have a ‘big data’ strategy’ because of the exploding volume and variety of data.



On average, 31% of federal agency data is unstructured; this figure is rising. Agencies want to tap this potential. Big Data Market by Federal Buyer Segment, FY 2012-2017 (Deltek, 2012) ranks the reasons agencies are pursuing data analytic projects.

Sample Breadth of Federal Data

Corporate Functions	Operations	Policy Development	Security and Reuse
Corporate Performance – responsiveness to changing requirements / capabilities	Business Process – operational service delivery: cost, performance, value	Data Driven Decisions – market data, customer data, behavioral trends, international benchmarking	Legislative Scrutiny – asking the right questions, performance data, etc.
Workforce Analytics – absenteeism, recruitment, performance, training, staff unit cost and skill mgmt.	Customers – behavior, requirements, supply and demand across product groups.	Regulatory Impact – cost and administrative burden, economic competitiveness, societal benefits	Transparency – improved awareness of policy effectiveness.
Financial Performance – asset mgmt., resource allocation, credit, cash, resource accounts	Technology – performance, cost, utility, interoperability, and technical problems	Service Demand – supply and demand projections, aggregation across silos (e.g., social care)	Third Party Reuse – key data to foster innovation and new ways to apply data
Asset Analysis – equipment, real estate, location, and utilization	Fraud, Error, and Waste – frequency, cost, problem areas, benchmarking	Market Failure – threats and risks, tracking outcomes, market abuse.	Social Media Channels – Organize data to be accessed by social media.
Procurement Analysis – risk based maintenance and regulation	Systems Performance – system benchmarking across customer channels	Externalities – Forecast primary and secondary impact as third parties respond	Commercial Value – new revenue streams by licensing government data

Source – Improving Public Services through Data Analytics, Deloitte LLP, 2011.

Data analytics is a hot topic because, for the first time in thirty years that widespread attention is being given to databases and data management. Tremendous amounts of innovation are taking place around NoSQL, streaming databases, in-memory databases, column stores, graph databases, and document stores. Innovation spans broad disciplines including managing data volume, data structures, data hardware, and query optimization.

Many federal agencies have already launched Big Data pilots and/or operational initiatives. One 2012 study¹ identified 155 federal data analytic projects including 46 in the healthcare arena, 31 back office projects, 30 related to science, space or technology, 21 national defense related, seven pertaining to financial oversight and fraud reduction, and four energy projects. The table above illustrates the breadth of the potential for data analytics.

Realizing this potential will require overcoming many challenges such as:

- Data is trapped in legacy systems or has not been digitized.
- Up to 80% of data is unstructured, un-cleansed, and/or duplicated.
- A long data lifespan (for retention) stresses storage while the value of data diminishes over time.
- The volume, variety, and velocity of data require new data governance mechanisms, systematic thinking about data inventories, data stewardship, and master data management.
- Planning is needed to leverage the technology, skills, equipment, and standards needed to realize value.
- Finding skilled data scientists – those who devise theories, design experiments, and test hypothesis to extract relevant information from data – is challenging as relatively few exist and they are in high demand.

Approach and Methodology

Our Approach

We regularly handle a variety of very interesting data sets including those from retail purchase, click-stream, email, hospital admissions, health care claims, system logs, fare and toll collection logs, credit and debit card transactions, social media, and student loan servicing. Xerox Research Centers bring together experts across multiple disciplines, such as computer science, machine learning, and natural language processing, to work with subject matter experts in business units to create innovative business process and document services.

Beneath unstructured, seemingly non-relational data lie hidden treasures of new insights and opportunities. To capture them, we need to collect, organize, and analyze data which requires highly sophisticated processing, modeling, and analytics capabilities.

For years, we have focused our research and development efforts on finding new and better ways to work with information. Our approach has been to concentrate our efforts on documents which we see as inextricably linked to information-related business processes. Documents are the containers for unstructured data, but on their own documents are not that smart. Our goal has been to bring intelligence to the document to make it easier to extract the relevant information. [Smarter Document ManagementSM](#) technologies, including the capabilities below, are uniquely differentiated components coming out of the world-class research. Developed within our state-of-the-art R&D research labs, these natural language processing technologies can be customized and combined to meet the most advanced content analytic needs.

Our solutions are developed within our state-of-the-art R&D research labs and applied to real client problems by subject matter experts who understand the terminologies (semantics), rules of behavior, and decision rules in a specific domain so that the data extracted is relevant to a business problem and organized for use in the field. We find that this approach results in clients successfully deploying and continuing to use analytic tools such as Hadoop, graph analytics and contextual intelligence to make optimal business decisions.

Methodology

To support data analytics projects, we have adapted our [SPARK-ITS[®]](#) program management methodology to provide managers with an integrated set of best practices for deploying data analytic systems, integrating data sources, refining data analysis, transferring knowledge, and receiving feedback. SPARK-ITS[®], depicted below, is a critical component of our approach to project management, quality management, and data governance – allowing us to manage and execute high-quality technical solutions both on time and within budget.



In addition, SPARK-ITS® is continually being refined based upon factors such as industry trends, feedback from data analytic projects, client surveys and evaluations, and internal improvement efforts. Our formal change, configuration, and release management approach ensures projects are provided with the most current documents, templates, tools, and standards to effectively execute project management and data analytic processes.

We tune our SPARK-ITS® method to meet client-specific needs and requirements, whether their project is large, midsize, or small. Review and approval processes help ensure projects maintain a consistent approach, minimize the learning curve, leverage proven practices, and maintain alignment with important industry standards such as Institute of Electrical and Electronic Engineers (IEEE), CMMI, ITIL, PMBOK® Guide, and the Standard for Program Management. As a result, each SPARK-ITS® solution is standardized, consistent, and optimized, yet specifically addresses the needs of each Xerox client.

Data Analytic Capabilities

Text Categorization

Our text categorization uses patented linguistic analysis technologies and machine learning algorithms involving two essential, tightly integrated data mining techniques: clustering and classification. Our expertise in natural language processing contributed to the success of companies such as PowerSet (acquired by Microsoft), Scansoft (Nuance), Microlytics, and Inxight.

Clustering technology classifies documents by splitting them into smaller groups of similar objects whose characteristics of similarity may not be known in advance. An example of cluster analysis is looking at fraudulent documents and finding other documents with similar linguistic features. Similarly, we use [ClusteriX](#) in the digital mailroom to route similar documents to the same work queue.

Classification identifies the categories in which new data points belong based on prior experience. We use advanced software such as [CategoriX](#) to analyze and filter all types of documents, such as letters, forms and invoices, to determine what type (category) of document they are, or to whom they belong, in a business process.

Image Categorization

Accessing the textual context of a document is not always possible or practical. Image categorization is a general-purpose image classification engine that is capable of categorization multiple everyday image content types, including buildings, airplanes, books, faces, and documents. Using a combination of expertise in image processing, computer vision, and machine learning, the computer is “trained” to map the key features, including the subject, texture, structure, layout, and colors of an object. Simply put, image categorization can recognize documents based on key visual features.

This system has proven state-of-the-art performance in a large number of scenarios such as consumer photography annotation, vehicle identification, and document object recognition. The technology is accurate, fast, and can cope effectively with hundreds of categories and the variations usually encountered in imaging scenarios. Image categorization often is used in classifying documents, such as recognizing pre-defined document types and routing accordingly or recognizing typewritten characters from handwritten ones and guiding OCR to the zones where the relevant information is located.

FactSpotter

Fact extraction aims at determining the meaning of statements and of the relations that link them together. Most commercial fact extraction systems cover a limited number of objects of named entities (such as names, locations, or dates) and identify simple links between them (such as the verb associated with a noun or synonyms). Semantics can offer a lot more as they enable users to find exactly what they are looking for without having to plough through entire documents. They can extract knowledge from large collections of unstructured documents including news media, scientific literature, financial reports, health records and claims, and enterprise document stores.



Clustering splits data into clusters that contain similar data items whose characteristics of similarity may not be known in advance.



Factspotter returns relevant answers instead of thousands of unrelated responses.

We have developed **FactSpotter**, a semantic text mining tool that goes beyond conventional keyword search, enabling users to extract quickly the one or two golden nuggets of critical information buried in mountains of documents. Factspotter has an easy-to-use interface so that anyone can query the system in everyday language. Unlike traditional enterprise search tools, it looks not only for the keywords contained in a query but also at the context in which those words appear. Because it “understands” the context it highlights only the relevant answers instead of returning thousands of unrelated responses. As a result, Factspotter promises a significant boost in productivity for data-intensive environments including legal discovery, risk management, pharmaceutical research, security intelligence, and fraud detection by significantly reducing search times and improving the relevance of results.

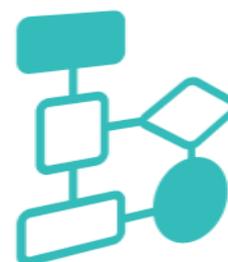
Advanced Text Analysis

Recognizing simple index fields, such as telephone, customer ID, social security number, IP address or other simple patterns is the easy stuff. Advanced text analysis uses powerful linguistic engines to recognize more advanced elements or concepts such as dates, addresses, company and person names, locations or organizations. It becomes extremely powerful analyzing documents to detect facts and other more advanced relationships, such as temporal relationships (as in before or after a date or event), and other types of relations (as in Company Y will earn Z dollars).

Advanced text analysis uses an impressive range of semantic and parsing technologies including:

- **Syntactic Analysis** – splits documents into sentences and words to tokenize text (standardize and normalize words), perform syntactic tagging (identify parts of speech in context), and chunk (group words together such as noun phrases or prepositional phrases) to produce a syntactic representation of the text. Such low level applications are the bricks of many information retrieval systems.
- **Sentence Normalization** – builds on syntactic analysis to resolve grammatical problems that the general grammar cannot easily solve. The module can detect when two sentences with different appearance have the same meaning and, conversely, when the same word has different meanings based on the context.
- **Entity Recognition** – is another module that extracts and recognizes specific terms and expressions and associates them with a topic. Examples of different types of entities that can be recognized include percentages (10% or 10 percent), dates (March 4, 1991), monetary expressions (\$26 billion), locations (San Francisco, Wall Street), names of persons (Barack Obama, Angela Merkel), organizations (Virgin America, JP Morgan), events (company merger, cancer treatment), and legal references (Dodd-Frank).
- **Concept Matching** – can recognize abstract concepts like “people” or “building” when expressed in completely different ways and retrieve all of the words that fit within that category. For example, it can detect that the following sentences refer to business acquisitions by a Japanese company: “Nekoosa will end up being owned by a Japanese company” and “A family-owned Japanese company called Kamori Kanko Co Ltd. became the new owner of Nekoosa.”
- **Semantic Disambiguation** – uses context to automatically select the correct meaning of a word in a sentence from the number of possible meanings that the word might have. Deep semantics is needed to distinguish between, for example, the different meanings of “treated with” in the following sentence: “The doctor treated his patients with respect” and “the doctor treated his patients with antibiotics.”
- **Co-reference** – conveys how individual expressions within or across documents are connected. At its simplest, for example, co-reference can tell that in the sentence “Bill said he would buy” that “Bill” and “he” probably refer to the same person (and thus are co-referent). These algorithms are at the very core of natural language processing and provide another powerful tool to optimize fact extraction from documents.
- **Temporal Expressions** – contained in expressions such as “next year, three years ago, or during”, provide key input for many applications such as information extraction, question answering, and summarization. By identifying such expressions, robust Xerox text analyzers are able to establish chronological timelines for a document enabling documents to be “anchored” in time.

These technologies can be deployed in unique combinations to support complex analytics such as for risk management and fraud prevention and detection.



Advanced text analytics provides for customizable data flows and works best for documents with a fairly rigid writing style such as financial reports.

Hadoop in the Cloud

To a significant extent, the current availability of big data solutions can be attributed to the development of Hadoop. In its current implementation, Hadoop assumes a homogeneous cluster of compute nodes. This assumption manifests in Hadoop's scheduling algorithms, but is also crucial to existing approaches for diagnosing performance issues, which rely on the peer similarity between nodes.

As more aspects of the data center are virtualized, the problems of diagnosing issues and optimizing data center performance becomes harder and harder. We have extended Hadoop and other cloud computing platforms by applying model-based diagnosis, machine learning, and artificial intelligence planning and scheduling to 1) manage heterogeneous clusters of compute nodes (violating the standard Hadoop peer-similarity assumption), 2) optimize the scheduling of jobs onto available resources, and 3) identify abnormally performing cluster nodes, and 4) to diagnose the type of fault occurring on the node (such as CPU contention or disk I/O contention). Applying these techniques improves the throughput of Hadoop and reduces the manual overhead involved with diagnosing faulty nodes.

Graph Analytics

While Hadoop is the right tool for massively parallel tasks, like many text processing tasks, it is rather inappropriate for various other kinds of problems. Graph analysis is one of them. Applications where the data takes the form of graphs and the analysis needs to take this graph structure into account are widespread, such as for social network analysis and PageRank. To accommodate these needs, and to enable real-time analysis of graph data, we are developing a high-performance, in-memory graph analysis engine that exploits parallelized search and compact graph representations that are amendable for traversal queries. We can then compute relevant graph properties and queries several orders-of-magnitude faster than existing solutions. For instance, we can compute the single-source shortest path for a set of 40 million Twitter users with 1.5 billion connections in less than 3 seconds, using only 6GB of RAM.

Predictive Modeling

Our predictive analytics use available information to generate predictions while modeling the entire population. They are not a single technology, but rather a data-driven approach we use to develop and deploy customized solutions to meet difficult business challenges. Our predictive analytics are domain agnostic and are used in financial services, transportation, and healthcare for a diverse set of business processes such as fraud identification, risk analysis, business activity monitoring, service forecasting, sentiment analysis, human resources and congestion.

We currently use predictive modeling in a diverse set of verticals including developing an early warning system for borrowers, dynamic queuing of (intelligent routing for) inbound telephone calls, analyzing web traffic, optimizing call center capacity, predicting the availability of parking spaces, and hiring workers more likely to stay longer.

Case Study – Predictive Analytics in Finance

Loan Risk Early Warning System

Loan risk modeling is a hard problem in quantitative analysis. Loan originators, servicing agencies and analytics providers participate in some aspect of loan modeling and/or subsequent use. Predictive analytics in this field have focused on scoring loan portfolios based on current and future risk. Scoring approaches take either known or proprietary metrics and calculate the risk associated with an individual or group.

We have invented a new approach called scalable multi-application predictive analytics for loan risk modeling. This approach incrementally accepts multi-application big data intuitively and incrementally. For example, it can accept input from payment records as well as call center and portal logs. With more information our approach can assess the effect of various parameters in an inclusive fashion. Scalability comes from (a) decentralized operations on top of industry-standard Hadoop/Map-Reduce framework without canned, external software and (b) easy application to other risk or loan analytics.

Approach and Data

Our approach begins with regression using time series data from multiple aforementioned data sources – supplemented by data from external sources such as the Department of Labor or public records.

Assuming we have data for a total of N months. From that we pick n (where $n < N$) months to act as our independent variables. To predict for k months in advance, we train the regression model with the (n+k)th variable as the response variable. We then pick another dataset in the same fashion to run the prediction.

Student loan data aggregated to form time-series include: (a) personal information such as address and age, (b) payment data such as principal balance outstanding, days since last payment and delinquency days, (c) web portal activity and inbound and outbound calls – often an indicator of the borrower’s interest in the loan, and (d) local economic data such as unemployment. As credit scores can be expensive, we utilize them only in a selective fashion.

Predicting Using the Model

The regression model predicts risk on a 0-100 scale (with 100 indicating high propensity of default and 70s indicating high risk for example) and acts as an early warning system. This means that the model gives an early warning (up to six months ahead) to the servicing organization to act in a manner that the business deems fit. To accommodate the diverse portfolio, the model predicts risk differently for private loans than for federal loans.

The model result can be validated by considering that the prediction should be close (say within ± 10 , ± 20 points) to the ground-truth severity of the same month. Generally, this validation can be done only when the current month’s data becomes available. So while the model tries to make a prediction several months in advance it is validated only subsequently.

Glimpses of Prediction Results

We experimented with a portfolio of about 5M loan records and has been able to achieve an accuracy of around 85%. We have assumed that a bracket of $[-10, 10]$ is the tolerable limit for the prediction accuracy. While prediction accuracy improves with more data and training, a conservative prediction is useful for operational reasons.

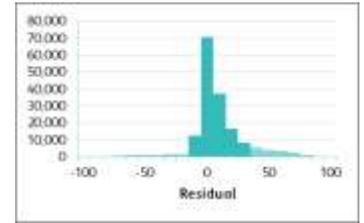
Making the model conservative means the prediction may be slightly higher than the actual risk. This assumes that it is safer to consider some borrowers to be more risky than they are and then be proven wrong rather than the alternate situation. To achieve this, we considered relative statistics and time series for variables such as falling outstanding balance and a stable volume of web communication – both desirable trends in loan servicing. Using several indicator variables like these, we proportionally boost the risk of those borrowers whose values vary from the normal. While some of the risks associated with these borrowers may be larger than actual, it provides the comfort of not missing them.

Response Strategy

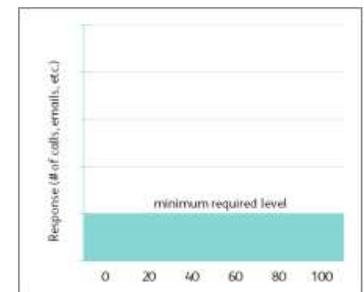
We respond to the risk using available machinery such as collections, call center, communications, and the automated dialer. We use three response strategies: a standard response or control group and two enhances response strategies. Each strategy reacts differently to borrower risk and use different outreach procedures with the goal of helping borrowers return to active repayment.

Conclusion

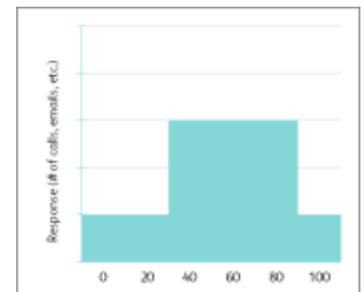
Our investigation targets the area of loan risk prediction and a proactive data-driven approach to loan processing. Currently we have a model that is more than 85% accurate i.e. we have an eight in ten chance of accurately predicting a high risk borrower. Efforts are underway to increase the extensibility, accuracy and the conservative application of the model in a real servicing scenario. We are currently piloting the solution with a few portfolios and proactively engaging borrowers on the basis of the predictions. We are collecting data to validate the efficacy of our response strategies. These response strategies on the predicted risk groups are also expected to decrease the proportion of customers in “past due” status.



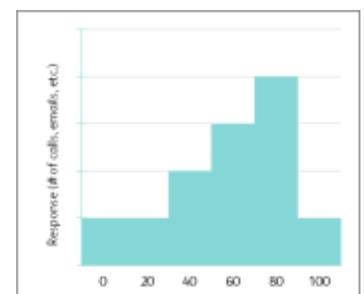
The model was calibrated so that the observed prediction accuracy fit in a $[-10,30]$ bracket.



The baseline model provided the minimum required level of interaction with borrowers.



Scenario A illustratively represents a heightened response strategy for some borrowers based on risk.



Scenario B illustratively represents a different heightened response strategy for other borrowers.

Case Study – Predictive Analytics in Health

Processing information available in computerized patient records is complex to manage as it requires many disciplines each with their own specifics and constraints. At the same time, the different sources of information available in a computerized patient record provide rich information that can be used to epidemiological surveillance and can alert health professionals of risk to the health of patients while maintaining high standards for patient confidentiality.

ALADIN – Predicting Hospital Acquired Infection

Each year, hospital -acquired infections (HAIs) affect millions of patients around the world, killing hundreds of thousands. While doctors and nurses have stepped-up hand washing and other methods to curtail infection, linguists at Xerox have teamed up with medical researchers in France to explore how “language technology” can help. During a three-year project, researchers used an advanced text analytics developed by us to analyze medical records, automatically identifying patients who could be at risk of contracting an HAI.

Hospitals are experimenting different mechanisms to detect and report these events. The most common process currently in involves medical staff reporting suspicious cases and experts from risk assessment departments analyzing reports to identify HAI. Efforts to find HAI more quickly and reduce subsequent infections are complex as one must link medical reports and subsequent infection events using complex tools across multiple disciplines such as medical terminology, linguistics, and conceptual modeling.

The ALADIN project used Factspotter to review medical records and identify specific terms and sequences of facts that indicate a patient may have contracted an HAI. The software not only pinpoints meaningful pieces of information, such as patient symptoms, drugs and names of bacteria but also how they are linked to each other. When these links identify potential risk of an HAI, the system automatically alerts the staff, so preventative measures can be taken.

According to the Centers for Disease Control, hospital-acquired infections in the U.S. result in an estimated 1.7 million infections and as many as 99,000 deaths each year. The annual cost is pegged at \$45 billion. In France alone it is estimated that 4,000 HAI related deaths occur each year and that a third of these could have been prevented

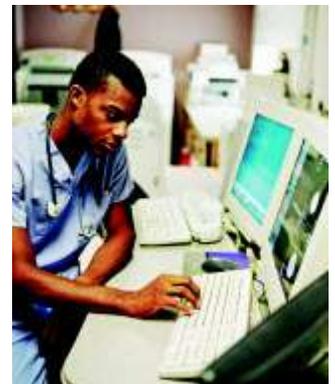
Approach and Data

The project brought together a range of unique competencies in the fields of natural language processing, terminology, knowledge representation, epidemiological surveillance, medicine and care associated infections. Technology experts come from Xerox Research Centre Europe, CISMef: the Catalog and Index of French-language Health Internet resources; Vidal -the French equivalent of the Physician's Desk Reference; and the research team in public health and epidemiology from Université Claude Bernard – Centre National de Recherche Scientifique UMR 5558 in Lyons.

“HAI infections are complex and can have many different causes. They may be the direct result of the type of care or completely independent of it but linked to a patient’s illness or condition,” said Dr. Marie-Hélène Metzger, medical lead on the project. “Linguistic technology plays a vital role in extracting the information required to correctly judge the situation and make the right decision. That’s what this project is all about.”

“Every patient is different, which makes it impossible to capture every piece of relevant information in a checklist or form,” said Frédérique Segond, principal scientist at the Xerox Research Centre Europe and coordinator of the project. “Using our advanced text mining technology to analyze entire patient records, we can extract information specific to each case to help doctors evaluate the patient risk and quickly take the right action.”

Another difficulty in detecting HAI is that most of the time there is no explicit mention of it in patient records. Rather, medical records only contain a sequence of events in a given period of time. Further, events typically are spread across multiple records over a period of time and may appear in different reports from different specialists.



We used our advanced text analytic tools to identify the presence of Hospital Acquired Infection when not explicitly stated in medical records.

The consortium selected the official terminologies required to cover the coding of symptoms, diagnosis, bacteriology, types of microorganisms, biological tests, antibiotics, and types of surgery. The team developed an indexing tool to make these terminologies accessible and an annotation tool for formal encoding of events in medical records.

We used its advanced text analysis tools to detect specific semantic relations between the events using specific linguistic rules customized for the medical language and a dedicated dictionary of 4000 medical terms. Characterizing HAI events also required addressing issues such as negation (no fever is semantically different from Temperature rise) and the detection of a specific chronology of events (time stamp tagging). Finally, Our linguists worked with medical experts to design decision rules. The consortium found that the decision rules needed to be less strict than official CDC definitions as practical experience with actual medical records suggested that all of the elements required by official CDC definitions often were not present. (For example, common knowledge typically does not appear in clinical records. Also, in practice health care providers simply don't record everything; that is, clinical records typically do not meet the standards of scientific research.)

Temporality – A Key Element

Taking into account the chronology of events turned out to be a key element in the decision mechanism for identifying HAI. Creating a chronology required analyzing each event date and categorizing the event (such as presence of micro-organisms, presence of diagnostic symptoms, or treatment provided). Finding a set of events is not enough to identifying HAI. Rather, one must establish a chronology of events to determine whether events occurred before or after a reference date for monitoring for HAI. For example, a knee surgery event indicates that the patient is hospitalized and should be monitored for HAI. Only the events in the chronology can be taken into account by the parser to find indications of infection (such as unexpected fever or delivery of antibiotics) in a follow-up period defined in the context of standard monitoring protocols for knee surgery.

To meet this challenge, we developed rules for identifying temporal phases of the events described and added a second language parser to index the time of each detected event in the patient record. This was possible because of the fairly rigid writing style of patient records.

Project Results

The project successfully removed personally identifying information, such as names of patients, doctors and hospitals. Further, the project successfully indexed the patient records using a specific terminology for HAI to support decision rules for reporting. A formal evaluation showed that the detection system found 93.3% of actual HAI surpassing objectives set at the beginning of the project.

About Our Services

All federal government agencies these days share a common goal to cut costs and increase efficiencies and we have the answer with our BPO – Business Process Outsource approach to business. We have a solid understanding of government operations and a vision of emerging federal government service trends. Among the services we can provide cost-saving solutions are:

- Customer Care Solutions
- Call Center and Help Desk
- Document Management Solutions
- Healthcare Solutions
- IT Solutions – VDI, MDM & Cloud Services
- Loan Servicing Solutions
- Transaction Processing Solutions

In closing, in a climate of shrinking IT Federal budgets our BPO message has the attention of the market by cutting costs and increasing efficiencies. We can help you enhance your federal services, under fixed-price, performance-based contracts. Anand and I look forward to working with you.

