Inside Analysis article – How Many Trips to the ER Does It Take?

In the days following my daughter's first birthday, she had an out-of-the-blue severe asthma attack. After three nights in the hospital, we were sent home with a nebulizer, a hefty prescription of Albuterol and instructions to give her breathing treatments every 4-6 hours.

One month later, she was back in the ER, having had, despite a rigorous routine of medicine, another severe attack.

I questioned her pediatrician whether these issues could be caused by allergies, since neither instance was prefaced with high levels of physical activity. The doctor was firm in her assertion that one-year-olds can't be diagnosed with asthma or allergies. Why not? Her answer: because that's what the American Academy of Pediatrics recommends.

Medical associations naturally make these recommendations based on data. The medical community arguably has collected more data than any other sciences division. Every association, from cancer research to genomics, has a ton of information. The data is truly BIG.

Consider that data attached to a single patient can include personal information (name, address, phone number, gender, age), employer, demographics, medical history, imaging files, medications, billing, existing conditions, primary care provider, insurance information, etc. And then there is data about disease, treatments, pharmaceuticals, diagnoses, drug interactions, medical devices and so on.

But every domain has its own ontologies and taxonomies, and hence, the data used by the governing body for pediatrics is

probably not the same as the data used by, say, The American Cancer Society. If these data stores could be shared and combined, it could open up the possibility of new discoveries and reduce the number of missed healthcare opportunities.

The Precision Medicine Initiative (PMI), announced last year, seeks to leverage big data to better understand the relationships between people, disease and other factors and harness that information to make individualized diagnoses and treatments. Although it's a top-down initiative, it's up to the health networks themselves to develop and implement the technology.

No relational database could manage a spread of structured data, unstructured data, images, clinical narratives, geolocation data, etc. Such enormous volumes of complex data seem destined for Hadoop or a data lake. However, the most feasible way to understand patterns and relationships buried in this data is by using semantic graph analytics.

One of the collaborations poised to capitalize on the PMI is Montefiore Medical Center and Franz Inc. They have teamed up (along with Intel, Cloudera, and Cisco) to create a Semantic Data Lake (SDL) for healthcare, helping to make precision medicine accurate and reliable. At the core of the SDL is Franz's flagship product, AllegroGraph, a semantic graph analytics product integrated with Hadoop.

The SDL for Healthcare platform combines the scalability of Cloudera's Hadoop distribution with the semantic query capabilities of AllegroGraph. The solution, which has some trillion edges, links domain specific ontologies, terminologies and taxonomies with electronic medical records, ERP data and device data to form a fully integrated repository. The fusion of these data sources enables the type of semantic consistency that is much needed but missing from the average data lake.

Combined with meticulous data governance and access controls, this repository of graph data provides users the ability to run complex queries, in real time, and visually discover and extract sophisticated insights and predictive analytics.

Precision medicine and accountable care rely on the ability to detect patterns, find relationships and make connections. The SDL for Healthcare links billions of entities and objects, truly providing an environment that empowers researchers and practitioners in the medical community to uncover important insights that could never ben found using a traditional database.

This will no doubt make a profound impact on the medical community and hospital networks. While technology like the SDL for Healthcare can also be leveraged for insights into fraud detection, operational analytics or risk assessments, it ultimately leads to improved patient outcomes and better quality care. And when healthcare systems and recommendation agencies have access to a governed, accurate semantic data lake of virtually all known data, they can make more informed decisions and turn precision medicine into a real possibility.

Returning to the opening narrative, my daughter went through her second year of life without any more asthma attacks. But right around her second birthday, we were back in the ER. I ignored the pediatrician's opinion and took her to an allergist.

The allergist asked a ton of questions. Do you have animals in the house? Yes. Do you often have cut flowers in the house? Yes. Do you spend a lot of time outside? Um, yes. *Why didn't the pediatrician ask any of these questions?*Turns out, a blood test revealed she's highly allergic to dogs and everything that grows during that time of year. Problem solved.

So how many trips to the emergency room would it have taken for the pediatrician to suggest allergies as the culprit? I

know a lot of healthcare is guesswork. But putting precision medicine into practice with a comprehensive, knowledge-based analytics solution like the SDL for Healthcare is bound to fuel a revolution in how and why decisions are made.

About the Authors:

Rebecca Jozwiak is Senior Analyst and Editor for The Bloor Group. She does research and writes publications on enterprise information management software, primarily on the topics of database technologies and analytics solutions. Having previously worked in the public sector, she has experience with computer networking, project management, technology training and teaching.

Robin Bloor is co-founder and Chief Analyst of The Bloor Group. He has more than 30 years of experience in the world of data and information management. He is the creator of the Information-Oriented Architecture, which is to data what the SOA is to services. He is the author of several books including, The Electronic B@zaar, From the Silk Road to the eRoad; a book on e-commerce and three IT books in the Dummies series on SOA, Service Management and The Cloud. He is an international speaker on information management topics. Robin has written scores of white papers, research reports and columns on a wide range of topics from database evaluation to networking options and comparisons to the enterprise in transition.