

AllegroGraph named to 2019 Trend-Setting Products

Database Trends and Applications – December 2018

You can call it the new oil, or even the new electricity, but however it is described, it's clear that data is now recognized as an essential fuel flowing through organizations and enabling never before seen opportunities. However, data cannot simply be collected; it must be handled with care in order to fulfill the promise of faster, smarter decision making.

More than ever, it is critical to have the right tools for the job. Leading IT vendors are coming forward to help customers address the data-driven possibilities by improving self-service access, real-time insights, governance and security, collaboration, high availability, and more.

To help showcase these innovative products and services each year, Database Trends and Applications magazine looks for offerings that promise to help organizations derive greater benefit from their data, make decisions faster, and work smarter and more securely.

This year our list includes newer approaches leveraging artificial intelligence, machine learning, and automation as well as products in more established categories such as relational and NoSQL database management, MultiValue, performance management, analytics, and data governance.

[Read the AllegroGraph Spotlight](#)

Knowledge Graphs – The path to true AI

Published in SD Times – December, 2018



Knowledge is the foundation of intelligence– whether artificial intelligence or conventional human intellect. The understanding implicit in intelligence, its application towards business problems or personal ones, requires knowledge of these problems (and potential solutions) to effectively overcome them.

The knowledge underpinning AI has traditionally come from two distinct methods: statistical reasoning, or machine learning, and symbolic reasoning based on rules and logic. The former approach learns by correlating inputs with outputs for increasingly progressive pattern identification; the latter approach uses expert, human-crafted rules to apply to particular real-world domains.

Read the full article at SD Times.

What is the most interesting use of a graph database you ever seen? PWC responds.

From a Quora post by Alan Morrison – Sr. Research Fellow at PricewaterhouseCoopers – November 2018

The most interesting use is the most powerful: standard RDF graphs for large-scale knowledge graph integration.


*From my notes on a talk Parsa Mirhaji of Montefiore Health System gave in 2017. **Montefiore uses Franz AllegroGraph, a distributed RDF graph database.** He modeled a core patient-centric hospital knowledge need using a simple standard ontology with a 1,000 or so concepts total.*

That model integrated data from lots of different kinds of heterogeneous sources so that doctors could query the knowledge graph from tablets or phones at a patient's bedside and get contextualized, patient-specific answers to questions for diagnostic purposes.

Fast forward to 2018, and nine out of ten of the most value-creating companies in the world are using standard knowledge graphs in a comparable fashion, either as a base for multi-domain intelligent assistants a la Siri or Alibot or Alexa, or to integrate and contextualize business domains cross-enterprise, or both. The method is preparatory to what John Launchbury of DARPA described as the Third Wave of AI.....

Montefiore's semantic data lake

Doctors can query the graph or harness ML + analytics and receive answers from the system at the point of care via their handhelds.



Montefiore Health, Franz, Intel and PwC research, 2017

PwC | Collapsing the IT stack

29

Read the full article over at Quora

2019 Trends in Data Governance: The Model Governance Question

From an AI Business Article by Jelani Harper – November 2018

The propagation of the enterprise's ability to capitalize on data-driven processes—to effectively reap data's yield as an organizational asset, much like any other—hinges on data governance, which arguably underpins the foundation of data management itself.

There are numerous trends impacting that foundation, many of which have always had, and will continue to have, relevance as 2019 looms. Questions of regulatory compliance, data

lineage, metadata management, and even data governance will all play crucial roles.

Franz's CEO, Dr. Jans Aasman was quoted:

Still, as Aasman denoted, "It's extremely complicated to make fair [machine learning] models with all the context around them." Both rules and human supervision of models can furnish a fair amount of context for them, serving as starting points for their consistent governance.

Read the full article at [AI Business](#).

AI Requires More Than Machine Learning

From Forbes Technology Council – October 2018

This article discusses the facets of machine learning and AI:

Lauded primarily for its automation and decision support, machine learning is undoubtedly a vital component of artificial intelligence. However, a small but growing number of thought leaders throughout the industry are acknowledging that the breadth of AI's upper cognitive capabilities involves more than just machine learning.

Machine learning is all about sophisticated pattern recognition. It's virtually unsurpassable at determining relevant, predictive outputs from a series of data-driven inputs. Nevertheless, there is a plethora of everyday, practical business problems that cannot be solved with input/output reasoning alone. The problems also require the

multistep, symbolic reasoning of rules-based systems.

Whereas machine learning is rooted in a statistical approach, symbolic reasoning is predicated on the symbolic representation of a problem usually rooted in a knowledge base. Most rules-based systems involve multistep reasoning, including those powered by coding languages such as Prolog.

Read the full article over at Forbes

.

Gartner – Knowledge Graphs Emerge in the HypeCycle

From Gartner – August 2018

Gartner Identifies Five Emerging Technology Trends That Will Blur the Lines Between Human and Machine

Gartner's HypeCycle report is now acknowledging Knowledge Graphs, a market area that Franz has been leading with AllegroGraph.

Read Jans Aasman's IEEE paper on the Enterprise Knowledge Graph for more insight.

From the Gartner Press release:

Digitalized Ecosystems

Emerging technologies require revolutionizing the enabling foundations that provide the volume of data needed, advanced compute power and ubiquity-enabling ecosystems. The shift from compartmentalized technical infrastructure to ecosystem-enabling platforms is laying the foundations for entirely new business models that are forming the bridge between humans and technology.

*This trend is enabled by the following technologies: Blockchain, Blockchain for Data Security, Digital Twin, IoT Platform and **Knowledge Graphs**.*

“Digitalized ecosystem technologies are making their way to the Hype Cycle fast,” said Walker. “Blockchain and IoT platforms have crossed the peak by now, and we believe that they will reach maturity in the next five to 10 years, with digital twins and knowledge graphs on their heels.”

Read the full article over at Gartner.

Venture Beat Features Montefiore's Healthcare project with AllegroGraph

From VentureBeat August 2018

This article discusses Montefiore's PALM project that uses AllegroGraph:

Montefiore is one of the largest employers in New York State. It's also one of the busiest health care complexes – hundreds of thousands of patients pass through its sprawling campuses, which include Montefiore Medical Center, the Albert Einstein College of Medicine, and Montefiore Medical Park.

Those logistical challenges catalyzed the development of Montefiore's Patient-centered Analytical Learning Machine (PALM), a machine learning platform built from the ground up to predict and prevent life-threatening medical conditions and minimize wait times.

PALM juggles lots of datasets – electronic medical records, insurance billing codes, drug databases, and clinical trial results, to name a few. And its analytical models recently expanded to handle voice, images, and sensor inputs from internet of things devices.

Core to the semantic graph model are triplestores, which are a type of database optimized for filing away and retrieving triples. They're an entity composed of subject-predicate-object – "John has tuberculosis," for example – which PALM builds dynamically, as needed. Along the way, the system uses a frame data language, or FDL, to resolve ambiguities, like when some electronic records refer to medication by its brand instead of by its scientific name (e.g., "Advil" or "Motrin" instead of ibuprofen).

Read the full article over at Venture Beat.

Transmuting Machine Learning into Verifiable Knowledge

From AI Business – August 2018

This article covers machine learning and AI:

According to Franz CEO Jans Aasman, these machine learning deployments not only maximize organizational investments in them by driving business value, but also optimize the most prominent aspects of the data systems supporting them.

“You start with the raw data...do analytics on it, get interesting results, then you put the results of the machine learning back in the database, and suddenly you have a far more powerful database,” Aasman said.

Dr. Aasman is further quoted:

For internal applications, organizations can use machine learning concepts (such as co-occurrence—how often defined concepts occur together) alongside other analytics to monitor employee behavior, efficiency, and success with customers or certain types of customers. Aasman mentioned a project management use case for a consultancy company in which these analytics were used to “compute for every person, or every combination of persons, whether or not the project was successful: meaning, done on time to the satisfaction of the customer.”

Organizations can use whichever metrics are relevant for their businesses to qualify success. This approach is useful for determining a numerical rating for employees “and you could put that rating back in the database,” Aasman said. “Now you can do a follow up query where you say how much money did I make on the top 10 successful people; how much money did I lose on the top 10 people I don’t make a profit

on.”

Read the full article over at AI Business.

Using AI and Semantic Data Lakes in Healthcare – FeibusTech Research Report



Artificial intelligence has the potential to make huge improvements in just about every aspect of healthcare. Learn how Montefiore Health Systems is using semantic data lakes, architectures, and triplestores to power AI patient-centered learning. With origins in post-9/11 municipal emergency projects, Montefiore Health Systems platform – called PALM, short for patient-centered Analytical Learning Machine – is beginning to prove itself out in the Intensive Care Unit, helping doctors save lives by flagging patients headed toward respiratory failure.

Intel and Montefiore in collaboration with FeibusTech have released a Research Brief covering Montefiore’s PALM Platform (aka – The Semantic Data Lake) powered by AllegroGraph.

“Just atop all the databases is what’s known as a triplestore, or triple, construct. That’s a key piece of any semantic data architecture. A triple is a three-part data series with a common grammar structure: that is, subject-predicate-object. Like, for example, John Smith has hives. Or Jill Martin takes ibuprofen.”

“Triples are the heart and soul of graph databases, or graphs, a powerful, labor-saving approach that associates John and Jill to records of humans, hives to definitions of maladies and Ibuprofen to catalogues of drugs. And then it builds databases on the fly for the task at hand based on those associations.”

Read the full article on Intel’s website to learn more about healthcare solutions based on AllegroGraph.

The Cornerstone of Data Science: Progressive Data Modeling

From AI Business June 27, 2018

This article covers Single Schema, Universal Taxonomies, Time Series Analysis, Accelerating Data Science and features some thought leadership from Franz Inc.’s CEO, Jans Aasman:

‘Contemporary data science and artificial

intelligence requirements simply can't wait for this ongoing, dilatory process. According to Jans Aasman, CEO of Franz, they no longer have to. By deploying what Aasman called an "events-based approach to schema", companies can model datasets with any number of differences alongside one another for expedited enterprise value.'

'The resulting schema is simplified, uniform, and useful in multiple ways. "You achieve two goals," Aasman noted. "One is you define what data you trust to be in the main repository to have all the truth. The second thing is you make your data management a little more uniform. By doing those two things your AI and your data science will become better, because the data that goes into them is better."'

Dr. Aasman goes on to note:

'The events-based schema methodology only works with enterprise taxonomies—or at least with taxonomies spanning the different sources included in a specific repository, such as a Master Data Management hub. Taxonomies are necessary so that "the type of event can be specified," Aasman said.'

'Moreover, taxonomies are indispensable for clarifying terms and their meaning across different data formats, which may represent similar concepts in distinct ways. Therefore, practically all objects in a database should be "taxonomy based" Aasman said, because these hierarchical classifications enable organizations to query their repositories via this uniform schema.'

Read the full article over at AI Business.

AI IN FOCUS

The Cornerstone of Data Science: Progressive Data Modeling