

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.

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](#).

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.

Optimizing Fraud Management with AI Knowledge Graphs

From Global Banking and Finance Review – July 12, 2018

This article discusses Knowledge Graphs for Anti-Money Laundering (AML), Suspicious Activity Reports (SAR), counterfeiting and social engineering falsities, as well as synthetic, first-party, and card-not-present fraud.

By compiling fraud-related data into an AI knowledge graph, risk management personnel can also triage those alerts for the right action at the right time. They also get the additive benefit of reusing this graph to decrease other risks for security, loans, or additional financial purposes.

Dr. Aasman goes on to note:

By incorporating AI, these threat maps yields a plethora of information for actually preventing fraud. Supervised learning methods can readily identify what events constitute fraud and which don't; many of these involve classic machine learning. Unsupervised learning capabilities are influential in determining normal user behavior then pinpointing anomalies contributing to fraud. Perhaps the most effective way AI underpins risk management knowledge graphs is in predicting the likelihood—and when—a specific fraud instance will take place. Once organizations have data for customers, events, and fraud types over a length of time (which could be in as little as a month in the rapidly evolving financial crimes space), they can compute the co-occurrence between events and fraud types.

Read the full article over at [Global Banking and Finance Review](#).



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](#).



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.

New York Times Article – Is There a Smarter Path to Artificial Intelligence?

From the New York Times – June 20, 2018

This article caught our attention because they featured a startup that was using Prolog for AI. We have been strong proponents of Prolog for Semantic Graph solutions for many years.

For the past five years, the hottest thing in artificial intelligence has been a branch known as deep learning. The grandly named statistical technique, put simply, gives computers a way to learn by processing vast amounts of data. Thanks to deep learning, computers can easily identify faces and recognize spoken words, making other forms of humanlike intelligence suddenly seem within reach.

Companies like Google, Facebook and Microsoft have poured

money into deep learning. Start-ups pursuing everything from cancer cures to back-office automation trumpet their deep learning expertise. And the technology's perception and pattern-matching abilities are being applied to improve progress in fields such as drug discovery and self-driving cars.

But now some scientists are asking whether deep learning is really so deep after all.....

.....Those other, non-deep learning tools are often old techniques employed in new ways. At Kyndi, a Silicon Valley start-up, computer scientists are writing code in Prolog, a programming language that dates to the 1970s. It was designed for the reasoning and knowledge representation side of A.I., which processes facts and concepts, and tries to complete tasks that are not always well defined. Deep learning comes from the statistical side of A.I. known as machine learning.

Our Tweet with links to AllegroGraph Prolog documenation and the full article:

nytimestech “computer scientists are writing code in **#Prolog**... It was designed for the reasoning and knowledge representation side of **#AI** ...” <https://buff.ly/2lmYwkv> – **#AllegroGraph** is the only **#GraphDatabase** to include **#Prolog** for your AI apps. <https://buff.ly/2yv0IzF>

How AI Boosts Human Expertise at Wolters Kluwer



Wolters Kluwer, a long time AllegroGraph customer, recently spoke with Alex Woodie at Datanami to describe how they are using AI tools such as AllegroGraph:

Thousands of companies around the world rely on Wolters Kluwer's practice management software to automate core aspects of their businesses. That includes doctor's offices that use its software make healthcare decisions in a clinical setting, corporate law offices that use its software to understand M&A activities, and accounting firms that use its software to craft tax strategies for high net-worth clients.

Wolters Kluwer is embedding a range of AI capabilities – including deep learning and graph analytics – across multiple product lines. For example, its Legalview Bill Analyzer software helps to identify errors in legal bills sent from outside law firms to the corporate counsels of large companies. The typical recovery rate for people reviewing bills manually is 1% to 2%. By adding machine learning technology to the product the recovery rate jumps to 7% to 8%, which can translate into tens of millions of dollars.

*Wolters Kluwer is using graph analytic techniques to accelerate the knowledge discovery process for its clients across various professions. **The company has tapped Franz's AllegroGraph software** to help it drive new navigational tools for helping customers find answers to their questions.*

*By arranging known facts and concepts as **triples in the AllegroGraph database** and then exposing those structures to users through a traditional search engine dialog box, Wolters Kluwer is able to surface related insights in a much more interactive manner.*

"We're providing this live feedback. As you're typing, we're providing question and suggestions for you live," Tatham

said. "AllegroGraph gives us a performant way to be able to just work our way through the whole knowledge model and come up with suggestion to the user in real time."

Read the full article over at Datanami.

Why Smart Cities Need AI Knowledge Graphs

A linked data framework can empower smart cities to realize social, political, and financial goals.



Smart cities are projected to become one of the most prominent manifestations of the Internet of Things (IoT). Current estimates for the emerging smart city market exceed \$40 trillion, and San Jose, Barcelona, Singapore, and many other major metropolises are adopting smart technologies.

The appeal of smart cities is binary. On the one hand, the automated connectivity of the IoT is instrumental in reducing costs associated with public expenditures for infrastructure such as street lighting and transportation. With smart lighting, municipalities only pay for street light expenses when people are present. Additionally, by leveraging options for dynamic pricing with smart parking, for example, the technology can provide new revenue opportunities.

Despite these advantages, smart cities demand extensive data management. Consistent data integration from multiple locations and departments is necessary to enable

interoperability between new and legacy systems. Smart cities need granular data governance for long-term sustainability. Finally, they necessitate open standards to future-proof their perpetual utility.

Knowledge graphs—enterprise-wide graphs which link all data assets for internal or external use—offer all these benefits and more. They deliver a uniform, linked framework for sharing data in accordance with governance protocols, are based on open standards, and exploit relationships between data for business and operational optimization. They supply everything smart cities need to realize their social, political, and financial goals. Knowledge graphs can use machine learning to reinsert the output of contextualized analytics into the technology stack, transforming the IoT's copious data into foundational knowledge to spur improved civic applications.

Read the full article at [Trajectory Magazine](#)

