

Webcast – Speech Recognition, Knowledge Graphs, and AI for Intelligent Customer Operations – April 3, 2019

Presenters – Burt Smith, N3 Results and Jans Aasman, Franz Inc.

In the typical sales organization the contents of the actual chat or voice conversation between agent and customer is a black hole. In the modern Intelligent Customer Operations center (e.g. N3 Results – www.n3results.com) the interactions between agent and customer are a source of rich information that helps agents to improve the quality of the interaction in real time, creates more sales, and provides far better analytics for management.

Join us for this Webinar where we describe a real world Intelligent Customer Operations center that uses graph based technology for taxonomy driven entity extraction, speech recognition, machine learning and predictive analytics to improve quality of conversations, increase sales and improve business visibility.

View the recorded webinar.

ГРАФОВЫЕ БАЗЫ: ПРИНЦИП РАБОТЫ

И ПРИМЕНЕНИЕ — GRAPH BASES: PRINCIPLE OF OPERATION AND APPLICATION

Всеволод Дёмкин удаленно работает во Franz Inc. над графовой базой AllegroGraph. Преподает в Projector курс «Natural Language Processing». В свободное время делает open-cors для обработки природных текстов на Lisp'e.

Мы рассмотрим создание программы для агрегации текстов из разных источников, таких как twitter, блоги, reddit и т.д., — их автоматической, а затем ручной обработки для формирования дайджеста новостей по определенной теме. На этом примере мы проанализируем, какие преимущества дает использование графовых баз данных, обсудим их возможности и ограничения.

В качестве конкретной БД будет использована система Franz AllegroGraph и мы ознакомимся с ее экосистемой, включающей возможности построение API и веб-приложений, а также со средой Allegro Common Lisp, на которой она построена. Особое внимание будет уделено использованию машинного обучения и NLP при решении задач работы с текстом, в частности, внутри AllegroGraph.

Обсудим:

- В чем особенности, как работают, преимущества/недостатки графовых БД;
- Как решать базовые задачи обработки текстов с использованием инструментария ML/NLP;
- Как построить полноценное приложение с ядром обработки текста на основе графовой БД и ML/NLP технологий;
- Как устроена экосистема Common Lisp и как можно задействовать ее для создания серверных приложений.

Лекция будет полезна: разработчикам, которые интересуются темой графовых баз данных и/или ML/NLP.

Semantic Web and Semantic Technology Trends in 2019

Dataversity – January 2019

What to expect of Semantic Web and other Semantic Technologies in 2019? Quite a bit. DATAVERSITY engaged with leaders in the space to get their thoughts on how Semantic Technologies will have an impact on multiple areas.

Dr. Jans Aasman, CEO of Franz Inc. was quoted several times in the article:

Among the semantic-driven AI ventures next year will be those that relate to the healthcare space, says Dr. Jans Aasman, CEO of Semantic Web technology company Franz, Inc:

“In the last two years some of the technologies were starting to get used in production,” he says. “In 2019 we will see a ramp-up of the number of AI applications that will help save lives by providing early warning signs for impending diseases. Some diseases will be predicted years in advance by using genetic patient data to understand future biological issues, like the likelihood of cancerous mutations – and start preventive therapies before the disease takes hold.”

If that's not enough, how about digital immortality via AI Knowledge Graphs, where an interactive voice system will bring public figures in contact with anyone in the real world?

“We’ll see the first examples of Digital Immortality in 2019 in the form of AI Digital Personas for public figures,” says Aasman, whose company is a partner in the Noam Chomsky Knowledge Graph:

“The combination of Artificial Intelligence and Semantic Knowledge Graphs will be used to transform the works of scientists, technologists, politicians, and scholars like Noam Chomsky into an interactive response system that uses the person’s actual voice to answer questions,” he comments.

“AI Digital Personas will dynamically link information from various sources – such as books, research papers, notes and media interviews – and turn the disparate information into a knowledge system that people can interact with digitally.” These AI Digital Personas could also be used while the person is still alive to broaden the accessibility of their expertise.

On the point of the future of graph visualization apps, Aasman notes that:

“Most graph visualization applications show network diagrams in only two dimensions, but it is unnatural to manipulate graphs on a flat computer screen in 2D. Modern R virtual reality will add at least two dimensions to graph visualization, which will create a more natural way to manipulate complex graphs by incorporating more depth and temporal unfolding to understand information within a time perspective.”

Read the full article at [Dataversity](#).

Solving Knowledge Graph Data Prep with Standards

Dataversity – December 2018

There's a general consensus throughout the data ecosystem that Data Preparation is the most substantial barrier to capitalizing on data-driven processes. Whether organizations are embarking on Data Science initiatives or simply feeding any assortment of enterprise applications, the cleansing, classifying, mapping, modeling, transforming, and integrating of data is the most time honored (and time consuming) aspect of this process.

Approximately 80 percent of the work of data scientists is mired in Data Preparation, leaving roughly 20 percent of their jobs to actually exploiting data. Moreover, the contemporary focus on external sources, Big Data, social and mobile technologies has exploded the presence of semi-structured and unstructured data, which accounts for nearly 80 percent of today's data and further slows the preparation processes.

Read the full article at [Dataversity](#).

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](#)

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.

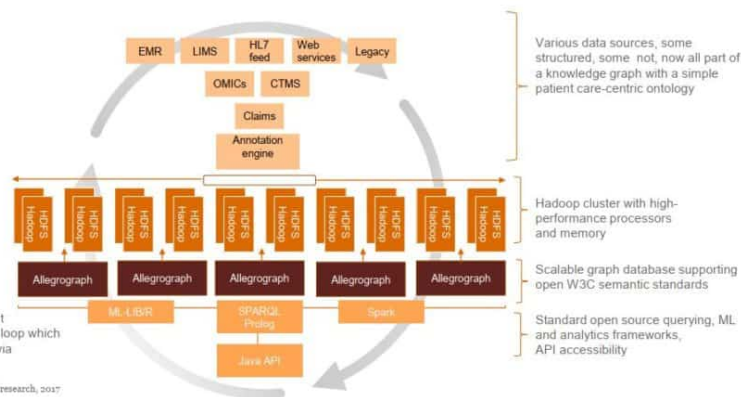
*From my notes on a talk Parsa Mirhaji of Montefiore Health System gave in 2017. **Montefiore uses Franz AllegroGraph, a distributed RDF graph database.** He modeled a core patient-centric hospital knowledge need using a simple standard ontology with a 1,000 or so concepts total.*

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.



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

.

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.