NLP: Unlock the Hidden Business Value in Voice Communications

By Dr. Jans Aasman, CEO, Franz Inc.

Today organizations capture an enormous amount of information in spoken conversations, from routine customer service calls to sophisticated claims processing interactions in finance and healthcare. But most of this information remains hidden and unused due to the difficulty of turning these conversations into meaningful data that can be effectively analyzed through Natural Language Processing (NLP).

Simply applying speech recognition software to voice conversations often results in unreliable data. State-of-the-art speech recognition systems still have trouble distinguishing between homophones (words with the same pronunciation, but different meanings), as well as the difference between proper names (i.e. people, products) and separate words. In addition, there is also the challenge of identifying domain-specific words accurately. Thus, in most cases, using speech recognition software alone doesn’t produce accurate enough data for reliable NLP.

Domain-specific taxonomies are key to understanding conversations via speech recognition systems. With them, we can feed conversations to knowledge graphs that understand the conversation and make connections in the data. Knowledge graphs provide the ability to extract the correct meaning of

text from conversations and connect concepts in order to add business value.

Knowledge graphs fed with NLP provide two prime opportunities for monetization. First, organizations can better understand their customers to improve products and services more to their liking, which in turn boosts marketing, sales and customer retention rates. Secondly, this analysis gives contact center agents real-time support for optimizing customer interactions to produce faster resolutions, better conversion rates, and cross-selling and up-selling opportunities. These approaches enable companies to capitalize on speech recognition knowledge graphs, accelerate their ROI, and expand their bottom lines.

Taxonomy Driven Speech Recognition

The story of taxonomy-driven speech recognition closely relates to knowledge graphs. The first wave of knowledge graphs was built from taking structured data and turning it into semantic graphs that support the linked open data movement. The next wave is all about unstructured data. People started doing Natural Language Processing on documents and textual conversations like emails and chats. Doing so accurately for a given domain requires a taxonomy to understand the words and concepts. Otherwise, downstream processes like entity extraction and event detection won’t work.

Read the full article at DZone.

The Future of AI: Machine Learning and Knowledge Graphs

Bringing knowledge graph and machine learning technology together can improve the accuracy of the outcomes and augment the potential of machine learning approaches. With knowledge graphs, AI language models are able to represent the relationships and accurate meaning of data instead of simply generating words based on patterns.

Read this special report to dive into key uses cases, best practices for getting started, and technology solutions every organization should know about.

The Future of AI: Machine Learning and Knowledge Graphs

Gartner Case Study: Entity-Event Knowledge Graph for Powering AI Solutions (Montefiore)

Gartner featured Franz's customer, Montefiore Medical Center, in a research report on Montefiore's Entity-Event Knowledge Graph:

"AI solutions are often hindered by fragmented data and siloed point solutions," according to Gartner's Chief Data and Analytics Officer Research Team. "Montefiore's data and analytics leader used semantic knowledge graphs to power its AI solutions and achieved considerable cost savings as well as

improvements in timeliness and the prediction accuracy of AI models.” Source: Gartner Case Study: Entity-Event Knowledge Graph for Powering AI Solutions (Montefiore) – Subscription required.

Copy Available from Montefiore/Einstein.

Understanding What Matters With Text Analytics and NLP

Dr. Jans Aasman was quoted extensively in this KMWorld Article:

Whether employing traditional rulesbased approaches to text analytics or leveraging more modern machine learning strategies, users must initially train the systems on relevant business domains. One way to do so is with comprehensive taxonomies of terms, their synonyms, and their meanings—which are traditionally associated with rules-based models. According to Franz CEO Jans Aasman, “There’s a part of NLP where people create taxonomies and ontologies. That is just a very acceptable way of doing NLP.” Historically, such defined hierarchies of vocabularies were paired with rules to find patterns in text and create actions such as classifications or entity extraction.

The trade-off between this approach and the taxonomic one is clear: Organizations can forsake the extensive time required

to build taxonomies by simply using annotated training data. The objective is to “just throw statistics and machine learning at the problem so it will all automatically work,” Aasman said. Although reduced time-to-value is an advantage of the deep learning approach, there are issues to consider, including the following:

♦ Training data: Machine learning models require immense amounts of training data, which organizations might not have for their domains. Transfer learning solves this problem by enabling subject matter experts to upload a couple of hundred examples (instead of thousands), highlight them, and teach dynamic models “the representative entities, key-value pairs, and classes they’re trying to derive from these documents,” Wilde noted.

♦ Controlled vocabularies: Transformers and techniques such as Bidirectional Encoder Representations from Transformers (BERT) reduce the training data quantities for machine learning models, broaden the array of training data that’s relevant, and implement a controlled vocabulary that otherwise isn’t as defined as taxonomic ones. Thus, organizations can take a phrase and “generate a similar phrase that means the same, but can be used in multiple reports in a controlled way,” Mishra said. Additionally, it’s possible to simply purchase libraries of terms and definitions. “Many companies end up buying those things to be able to incorporate those capabilities,” Shankar added.

♦ Practical knowledge: Exclusively using machine learning models to train text analytics decreases the real-world understanding and applicability of text. “People that do machine learning don’t want to spend the effort to create a vocabulary or the pragmatics or the semantics,” Aasman noted. “Machine learning has a place in all of this, but it misses part of the whole future solution where we have systems that understand what people are talking about.”

Read the full article at KMWorld.

KMWorld 100 Companies that Matter Most – Franz Inc.

Franz Inc., is proud to announce that it has been named to The 100 Companies That Matter in Knowledge Management by KMWorld. The annual list reflects the urgency felt among many organizations to provide a timely flow of targeted information. Among the more prominent initiatives is the use of AI and cognitive computing, as well as related capabilities such as machine learning, natural language processing, and text analytics.

“Flexibility, agility, and the ability to pivot are attributes that have become critical to forward-thinking companies—and that is particularly the case now. Successful organizations don’t want to merely survive; they want to dominate their market sectors. But to do that, they need the right tools and products,” said Tom Hogan, Group Publisher at KMWorld. “Amidst the dramatic changes taking place today, innovative organizations are seeking new approaches to improve their processes. The 2021 KMWorld 100 is a list of leading-edge knowledge management companies that are helping their customers to expand access to information, leverage new opportunities, and accelerate growth.”

Read More about Franz Inc.

Maximizing Your Data Fabric's ROI via Entity Data Modeling

Data fabrics are emerging as the most effective means of integrating data throughout the enterprise. They deliver a single access point for all data regardless of location – whether it's at rest or in motion. Experts agree that data fabrics are the future of data analytics and management. Gartner recommends:

"Data and analytics leaders must upgrade to a data fabric design that enables dynamic and augmented data integration in support of their Data Management strategy."

Forrester states that "Enterprise Architecture (EA) pros should use data fabric to democratize data across the enterprise for various use cases."

However, the adoption rate of data fabrics hinges on the ROI of their use cases. One such use case is to make it easier to do advanced Data Science on available data sources. Currently, extracting machine learning features is an exacting, time-consuming process because relevant data is trapped in silos. Data fabrics and knowledge graphs have a unique, symbiotic relationship because they substantially streamline the processes to extract data from the myriad sources that populate these platforms. Knowledge graphs are key to providing fundamental capabilities enabling data fabrics to accomplish this objective.

Read the full article at Dataversity.

Data-Centric Architecture Forum – DCAF 2021

Data and the subsequent knowledge derived from information are the most valuable strategic asset an organization possesses. Despite the abundance of sophisticated technology developments, most organizations don't have disciplines or a plan to enable data-centric principles.

DCAF 2021 will help provide clarity.

Our overarching theme for this conference is to **make it REAL**. Real in the sense that others are becoming data-centric, it is achievable, and you are not alone in your efforts.

Join us in understanding how data as an open, centralized resource outlives any application. Once globally integrated by sharing a common meaning, internal and external data can be readily integrated, unlike the traditional "application-centric" mindset predominantly used in systems development.

The compounding problem is these application systems each have their own completely idiosyncratic data models. The net result is that after a few decades, hundreds or thousands of applications implemented have given origin to a segregated family of disparate data silos. Integration debt rises and unsustainable architectural complexity abounds with every application bought, developed, or rented (SaaS).

Becoming data-centric will improve data characteristics of findability, accessibility, interoperability, and re-usability (FAIR principles), thereby allowing data to be exported into any needed format with virtually free integration.\

Dr. Jans Aasman to present – Franz's approach to Entity Event Data Modeling for Enterprise Knowledge Fabrics

---

AllegroGraph 7 Named a Trend Setting Product for 2021

AllegroGraph receives numerous industry awards and recognition for 2020

Franz Inc., an early innovator in Artificial Intelligence (AI) and leading supplier of Graph Database technology for Knowledge Graph Solutions, today announced that AllegroGraph has been named a Trend Setting Product in Data and Information Management by Database Trends and Applications (DBTA). In 2020, AllegroGraph has received numerous industry awards and independent analyst firms have positioned AllegroGraph 7 as a Champion and Strong Performer.

"Data management and integration demands continue to increase as organizations are faced with more data flowing in from a greater variety sources than ever before," said Tom Hogan, Group Publisher at DBTA. "To help make the process of identifying useful products and services easier, each year Database Trends and Applications magazine presents a list of Trend-Setting Products. These products, platforms, and

services range from widely accepted offerings that continue to evolve to meet the needs of their loyal constituents to breakthrough technologies that are in the early stages of adoption. However, the common denominator for all is that they represent a commitment to innovation and seek to provide organizations with tools to address changing market requirements."

"We are honored to be recognized again by DBTA as a trend setter in data management," said Dr. Jans Aasman, CEO of Franz Inc. "Organizations across a range of industries are realizing the critical role that Knowledge Graphs play in creating rich, yet flexible enterprise data fabrics and AI-driven applications. In just the past couple of years, we have helped our customers create large-scale, multi-model and innovative knowledge graph solutions for diverse use cases, such as healthcare real-time AI decision support, NLP 360 customer intelligence with real-time agent support, and social network data privacy compliance."

During 2020, AllegroGraph and Franz were recognized by the following industry analysts and technology media.

* AllegroGraph named a DBTA 2021 Trend Setting Product.

* Franz was positioned as a Strong Performer in the first Forrester Wave™: Graph Database Platforms 2020, Q4.

* Bloor Research positioned AllegroGraph as a Champion in the 2020 Bloor Research Graph Database report, which recognized AllegroGraph as a multi-model RDF database.

* KMWorld named Franz an AI 50: The Companies Empowering Intelligent Knowledge Management.

* Database Trends and Applications (DBTA) named Franz a Big Data 50—Company Driving Innovation in 2020.

* AllegroGraph was a KMWorld 2020 Trend Setting Product, as

a noteworthy solution transforming information into insight.

* Franz CEO Dr. Jans Aasman was featured as an expert in The Knowledge Graph Cookbook, which was released April 22, 2020 and explains why and how to build Knowledge Graphs that help enterprises use data to innovate, create value and increase revenue.

* Franz was recognized as one of the 100 Companies That Matter Most in Data by Database Trends and Applications (DBTA).

* KMWorld's 100 Companies that matter in Knowledge Management named Franz Inc. to this exclusive list.

AllegroGraph 7 is a breakthrough solution that allows infinite data integration through a patented approach unifying all data and siloed knowledge into an Entity-Event Knowledge Graph solution that can support massive big data analytics. AllegroGraph 7 utilizes unique federated sharding capabilities that drive 360-degree insights and enable complex reasoning across a distributed Knowledge Graph.

"AllegroGraph 7's support of Entity-Event Data Modeling is the most welcome innovation and addition to our arsenal in reimagining healthcare and implementing Precision Medicine," said Dr. Parsa Mirhaji, Director of Center for Health Data Innovations at the Albert Einstein College of Medicine and Montefiore Health System, NY. "Precision Medicine is about moving away from statistical averages and broad-based patterns. It is about connecting many dots, from different contexts and throughout time, to support precision diagnosis and to recommend the precision care that can take into account all the subtle differences and nuisances of individuals and their personal experiences throughout their life. This technology is about saving lives, by leveraging data, context and analytics and is what Franz's Entity-Event Data Modeling

brings to the table."

AllegroGraph 7 provides users with an integrated version of Gruff, a unique browser-based graph visualization software tool for exploring and discovering connections within enterprise Knowledge Graphs. Gruff enables users to visually build queries and visualize connections between data without writing code, which speeds discoveries and enhances the ability to uncover hidden connections within data.

"Few tools exist that can quickly turn arbitrary RDF graph pattern matches into clear visualizable results," said Michael Pool, Global Head of Semantic Modeling and Engineering, Senior Director at BNY Mellon Bank. "Gruff is invaluable in turning our knowledge graph data into useful and actionable analytic insights."

Louis Rumanes at UnitedHealth Group Research and Development recognizes the value of using Gruff as a browser-based app and commented, "Nice job on Gruff in a browser and I think this will be a gamechanger."

Gartner predicts "the application of graph processing and graph DBMSs will grow at 100 percent annually through 2022 to continuously accelerate data preparation and enable more complex and adaptive data science." In addition, Gartner named graph analytics as a "Top 10 Data and Analytics Trend" to solve critical business priorities." (*Source: Gartner, Top 10 Data and Analytics Trends, November 5, 2019)*

Data Fabrics and Knowledge

Graphs – A Symbiotic Relationship

Dr. Jans Aasman's recent article in Dzone.

The data fabric notion is gaining credence throughout the analyst community, in much the same way knowledge graphs have done so for years. Both technologies link all relevant data for a specific business purpose, which is why the most successful companies in the world employ them.

Amazon's knowledge graph retains metadata about its vast product array; Google's captures data about an exhaustive list of web entities of interest. Lesser-known organizations regularly deploy these mechanisms for everything from comprehensive customer views to manufacturing processes.

Data fabrics have a unique, symbiotic relationship with the knowledge graph movement because they substantially streamline the processes to extract data from the myriad sources that populate these platforms. In turn, knowledge graphs provide some of the fundamental capabilities enabling data fabrics to accomplish this objective.

Read the Full Article at Dzone.

RDF vs Property Graph – The Graph Show

The inaugural episode of The Graph Show, featured Josh Shinavier, Research Scientist at Uber, interviewing Franz Inc.’s CEO, Dr. Jans Aasman.

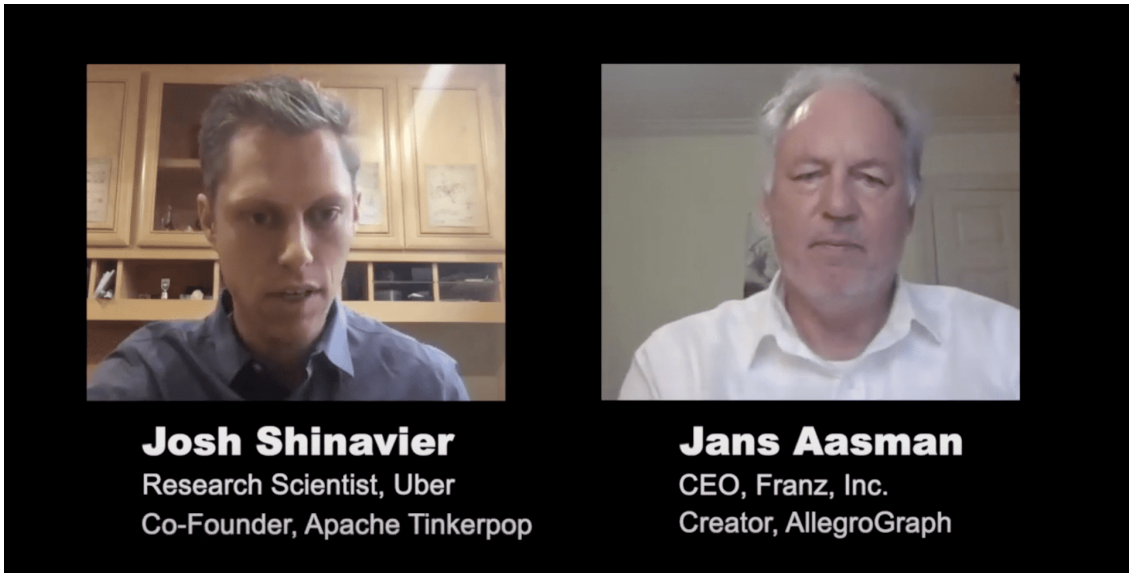

For more info on The Graph Show, visit: http://thegraphshow.com