Why Young Developers Don’t Get Knowledge Graphs

Dr. Aasman recently interviewed for this Datanami article.

Business is booming these days for graph databases—maybe it took COVID to show us how connected everything is—and that’s good news for Franz, which develops a semantic graph database called AllegroGraph. Just the same, you won’t find CEO Jans Aasman spending much time convincing developers of a certain age to use it.

“If you live in our world of semantic graph databases, I only talked to people over 35, 40,” Aasman tells *Datanami*. “I never talk to young developers.”

The problem with younger developers, he explains, is that they’re usually interested in using the graph database to build point solutions to solve specific problems, as opposed to creating a wide base of knowledge that can not only solve a specific problem, but be used with future solutions too. Plus, building point solutions exacerbates the data silo problem, he says.

“In our community of the semantic graph databases, literally everything is about integration and making sure that everything can interoperate,” the Franz CEO continues. “And there’s not a single young programmer that cares about that. Seriously. You’re young, you want to do a fun project, your managers are saying, in three months I need this thing done. You do whatever you want to do. Well, they get it done. And then you have new a data silo.”

Read the full article at Datanami.

AiThority Interview with Dr. Jans Aasman

Jans Aasman, please tell us about your current role and the team / technology you handle at Franz.

As CEO of Franz Inc., I drive the overall technology vision for our Enterprise Knowledge Graph solutions and ensure our customer projects deliver the ROI results expected with graph based architectures.

Franz Inc. is composed of an expert team with skills in Graph Databases, Semantic technologies, Graph Visualization, AI, NLP and Machine Learning. Our domain knowledge encompasses large enterprises in Healthcare, Pharma, Customer Support, and Intelligence Agencies.

Our main business today revolves around AllegroGraph, a Semantic Graph platform 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's FedShard feature utilizes patented federated sharding capabilities that drive 360-degree insights and enable complex

reasoning across a distributed Knowledge Graph. AllegroGraph is utilized by dozens of the top Fortune 500 companies worldwide.

We also offer a popular data visualization and no-code query builder called Gruff – the most advanced Knowledge Graph visualization application on the market, which we recently integrated into Franz AllegroGraph. Gruff enables users to create visual Knowledge Graphs that display data relationships in views that are driven by the user. Ad hoc and exploratory analysis can be performed by simply clicking on different graph nodes to answer questions. Gruff's unique 'Time Machine' feature provides the capability to explore temporal context and connections within data. The visual query builder within Gruff empowers both novice and expert users to create simple to highly complex queries without writing any code.

Read the full interview at AIThority.

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.

Connected Data London – The Future of AI in the Enterprise

The Future of AI in the Enterprise:

Entity-Event Knowledge Graphs for Data-Centric Organizations

Presented by: Dr. Jans Aasman

Register: https://enterprise-kg-cdl-online-meetup.heysummit.com/

Personalized medicine. Predictive call centers. Digital twins for IoT. Predictive supply chain management, and domain-specific Q&A applications. These are just a few AI-driven applications organizations across a broad range of industries are deploying.

Graph databases and Knowledge Graphs are now viewed as a must-have by Enterprises serious about leveraging AI and predictive analytics within their organization.

See how Franz Inc. is helping organizations deploy novel

Entity-Event Knowledge Graph Solutions to gain a holistic view of customers, patients, students or other important entities, and the ability to discover deep connections, uncover new patterns and attain explainable results.

Description:

To support ubiquitous AI, a Knowledge Graph system will have to fuse and integrate data, not just in representation, but in context (ontologies, metadata, domain knowledge, terminology systems), and time (temporal relationships between components of data). Building from ‘Entities’ (e.g. Customers, Patients, Bill of Materials) requires a new data model approach that unifies typical enterprise data with knowledge bases such as industry terms and other domain knowledge.

Entity-Event Knowledge Graphs are about connecting the many dots, from different contexts and throughout time, to support and recommend industry-specific solutions that can take into account all the subtle differences and nuisances of entities and their relevant interactions to deliver insights and drive growth. The Entity-Event Data Model we present puts core entities of interest at the center and then collects several layers of knowledge related to the entity as ‘Events’.

Franz Inc. is working with organizations across a broad range of industries to deploy large-scale, high-performance Entity-Event Knowledge Graphs that serve as the foundation for AI-driven applications for personalized medicine, predictive call centers, digital twins for IoT, predictive supply chain management and domain-specific Q&A applications—just to name a few.

During this presentation we will explain and demonstrate how Entity-Event Knowledge Graphs are the future of AI in the Enterprise.

Knowledge graphs enhance customer experience through speed and accuracy

KMWorld's recent article covers AllegroGraph and Franz's customer N3 Solutions.

The Full Article – KMWorld

Knowledge graphs are a way to model enterprise knowledge and represent complex interrelationships in data. Information stored in a graph database can enable rapid retrieval of well-targeted results and provide insights into customers' interests and needs. Gartner predicts a 100% per-year growth in applications for graph analytics and databases for the next several years. Although knowledge graphs have been deployed by major companies such as Google, Amazon, and LinkedIn due to their ability to incorporate relationships in their analyses as well as their speed, only in the last 5 years has their use become more widespread.

N3 is an outsourced sales company for major organizations that sell complex B2B software, hardware, and tech solutions. It supports businesses in 92 countries, provides services in 25 languages, and holds thousands of hours of conversations every month with customers and prospects. "In today's world of complex products, it takes a well-educated team to tell the story about how this technology can help a company become more

competitive," said Shannon Copeland, COO of N3. "The sales team needs to be able to instantly access the information they need to do their job."

Faster insights

The company has been operating for 16 years, and in the last few years began an initiative to manage its knowledge in a more intentional way. "We generate a great deal of data," noted Copeland, "and we wanted to make more effective use of it to understand our customers. And because of the speed at which business is transacted now, we needed to get insights right away, not a month later in a report."

N3 built a data model to reflect the essential data elements and the associations among them and decided that a knowledge graph was the best way to represent the information. After looking into partner options, N3 chose to work with Franz, Inc., which provides a semantic graph database called AllegroGraph. "We decided to work with Franz because of its extensive experience and the fact that it had worked with a variety of industries," Copeland said.

The system built by N3 allows sales teams to organize signals from the market in a way that allows them to better explain the products to prospective buyers. "We build relationships with tech buyers on behalf of our clients," continued Copeland. "Our employees are typically college graduates who would like to begin their careers in sales and marketing in tech solutions. They take ownership of their territory and we help them be as sophisticated as a future CMO would be." The resources supplied by the knowledge graph provide the support the sales team needs to tailor information to each prospective customer.

The specific expertise required by the team varies depending on the products being sold, the geographic region, and other factors, and the knowledge graph supports these needs. For

example, if a team in southern Portugal needs to know the preferences of that market, the associations built into the graph database can provide the information that is essential for them. “The information we can access helps customers understand the answers to their questions very quickly,” Copeland commented. “We believe the experience that the customers have helps them scope out what they need and what the road map might be.”

The strength of graph databases

A graph databases is a type of NoSQL database that stores data according to associations among data elements rather than in the rows and columns of a relational database. Because graph databases use a dynamic schema rather than a fixed, table-based one, adding new data types and categories is much easier. And because they are semantics-based, graph databases have strengths in inferring intent, producing answers to questions, and making recommendations. They can also make inferences about possible associations from existing associations.

A graph database also provides much more context than a relational database and therefore can return more relevant results when a user is searching; they also integrate data from multiple sources. “At one telecom company we worked with, customer service reps might have [had] to open 15 databases to find out what went wrong and what the solution was,” said Jans Aasman, CEO of Franz. “We took their core customer data, billing information, every CRM call, and every action and put them into AllegroGraph, and the customer service reps were finally able to respond in a meaningful way, whether that was to make an offer to the customer or provide appropriate technical support.” The capability of graph databases to overcome silos and provide an integrated view of the customer is one of its strengths.

In order to create the graph database on which the knowledge

graph is built, the relationships among entities need to be mapped. In the case of a hospital patient, the patient is the core entity, and the events are medical encounters or lab results, which may come out of different databases or a data warehouse. “The mapping is a major project, but it only needs to be done once,” Aasman pointed out. “After that, the relationships do not need to be regenerated during the search because they are indexed in AllegroGraph, which makes retrieval very rapid.”

AllegroGraph Named to 100 Companies That Matter Most in Data

Franz Inc. Acknowledged as a Leader for Knowledge Graph Solutions

Lafayette, Calif., June 23, 2020 – Franz Inc., an early innovator in Artificial Intelligence (AI) and leading supplier of Semantic Graph Database technology for Knowledge Graph Solutions, today announced that it has been named to The 100 Companies That Matter in Data by Database Trends and Applications. 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. This list recognizes companies based on their presence, execution, vision and innovation in delivering products and services to the marketplace.

“We’re excited to announce our eighth annual list, as the

industry continues to grow and evolve," remarked Thomas Hogan, Group Publisher at Database Trends and Applications. "Now, more than ever, businesses are looking for ways transform how they operate and deliver value to customers with greater agility, efficiency and innovation. This list seeks to highlight those companies that have been successful in establishing themselves as unique resources for data professionals and stakeholders."

"We are honored to receive this acknowledgement for our efforts in delivering Enterprise Knowledge Graph Solutions," said Dr. Jans Aasman, CEO, Franz Inc. "In the past year, we have seen demand for Enterprise Knowledge Graphs take off across industries along with recognition from top technology analyst firms that Knowledge Graphs provide the critical foundation for artificial intelligence applications and predictive analytics.

Our recent launch of AllegroGraph 7 with FedShard, a breakthrough that allows infinite data integration to unify all data and siloed knowledge into an Entity-Event Knowledge Graph solution will catalyze Knowledge Graph deployments across the Enterprise."

Gartner recently released a report "How to Build Knowledge Graphs That Enable AI-Driven Enterprise Applications" and have previously stated, "The application of graph processing and graph databases will grow at 100 percent annually through 2022 to continuously accelerate data preparation and enable more complex and adaptive data science." To that end, 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)*.

"Graph databases and knowledge graphs are now viewed as a must-have by enterprises serious about leveraging AI and predictive analytics within their organization," said Dr. Aasman "We are working with organizations across a broad range

of industries to deploy large-scale, high-performance Entity-Event Knowledge Graphs that serve as the foundation for AI-driven applications for personalized medicine, predictive call centers, digital twins for IoT, predictive supply chain management and domain-specific Q&A applications – just to name a few."

Forrester Shortlists AllegroGraph

AllegroGraph was shortlisted in the February 3, 2020 Forrester Now Tech: Graph Data Platforms, Q1 2020 report, which recommends that organizations "Use graph data platforms to accelerate connected-data initiatives." Forrester states, "You can use graph data platforms to become significantly more productive, deliver accurate customer recommendations, and quickly make connections to related data."

Bloor Research covers AllegroGraph with FedShard

Bloor Research Analyst, Daniel Howard noted "With the 7.0 release of AllegroGraph, arguably the most compelling new capability is its ability to create what Franz refers to as "Entity-Event Knowledge Graphs" (or EEKGs) via its patented FedShard technology." Mr. Howard goes on to state "Franz clearly considers this a major release for AllegroGraph. Certainly, the introduction of an explicit entity-event graph is not something I've seen before. The newly introduced text to speech capabilities also seem highly promising."

AllegroGraph Named to KMWorld's 100 Companies That Matter in Knowledge Management

AllegroGraph was also recently named to KMWorld's 100 Companies That Matter in Knowledge Management. The KMWorld 100 showcases organizations that are advancing their products and capabilities to meet changing requirements in Knowledge Management.

Franz Knowledge Graph Technology and Services

Franz’s Knowledge Graph Solution includes both technology and services for building industrial strength Entity-Event Knowledge Graphs based on best-of-class tools, products, knowledge, skills and experience. At the core of the solution is Franz’s graph database technology, AllegroGraph with FedShard, which is utilized by dozens of the top F500 companies worldwide and enables businesses to extract sophisticated decision insights and predictive analytics from highly complex, distributed data that cannot be uncovered with conventional databases.

Franz delivers the expertise for designing ontology and taxonomy-based solutions by utilizing standards-based development processes and tools. Franz also offers data integration services from siloed data using W3C industry standard semantics, which can then be continually integrated with information that comes from other data sources. In addition, the Franz data science team provides expertise in custom algorithms to maximize data analytics and uncover hidden knowledge.

Ubiquitous AI Demands A New Type Of Database Sharding

Forbes published the following article by Dr. Jans Aasman, Franz Inc.’s CEO.

The notion of sharding has become increasingly crucial for selecting and optimizing database architectures. In many cases, sharding is a means of horizontally distributing data; if properly implemented, it results in near-infinite scalability. This option enables database availability for business continuity, allowing organizations to replicate databases among geographic locations. It's equally useful for load balancing, in which computational necessities (like processing) shift between machines to improve IT resource allocation.

However, these use cases fail to actualize sharding's full potential to maximize database performance in today's post-big data landscape. There's an even more powerful form of sharding, called "hybrid sharding," that drastically improves the speed of query results and duly expands the complexity of the questions that can be asked and answered. Hybrid sharding is the ability to combine data that can be partitioned into shards with data that represents knowledge that is usually un-shardable.

This hybrid sharding works particularly well with the knowledge graph phenomenon leveraged by the world's top data-driven companies. Hybrid sharding also creates the enterprise scalability to query scores of internal and external sources for nuanced, detailed results, with responsiveness commensurate to that of the contemporary AI age.

Read the full article at Forbes.

Document Knowledge Graphs with NLP and ML

A core competency for Franz Inc is turning text and documents into Knowledge Graphs (KG) using Natural Language Processing (NLP) and Machine Learning (ML) techniques in combination with AllegroGraph. In this document we discuss how the techniques described in [NLP and ML components of AllegroGraph] can be combined with popular software tools to create a robust Document Knowledge Graph pipeline.

We have applied these techniques for several Knowledge Graphs but in this document we will primarily focus on three completely different examples that we summarize below. First is the Chomsky Legacy Project where we have a large set of very dense documents and very different knowledge sources, Second is a knowledge graph for an intelligent call center where we have to deal with high volume dynamic data and real-time decision support and finally, a large government organization where it is very important that people can do a semantic search against documents and policies that steadily change over time and where it is important that you can see the history of documents and policies.

Example [1] Chomsky Knowledge Graph
The Chomsky Legacy Project is a project run by a group of admirers of Noam Chomsky with the primary goal to preserve all his written work, including all his books, papers and interviews but also everything written about him. Ultimately students, researchers, journalists, lobbyists, people from the AI community, and linguists can all use this knowledge graph for their particular goals and questions.

The biggest challenges for this project are finding causal relationships in his work using event and relationship extraction. A simple example we extracted from an author quoting Chomsky is that neoliberalism ultimately causes childhood death.

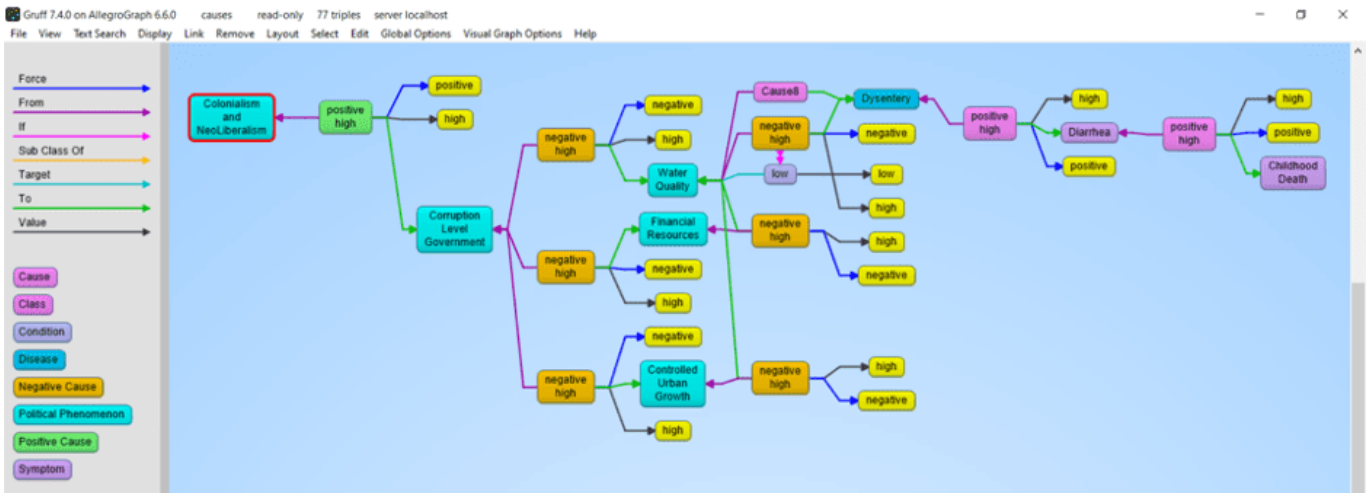

Example 2: N3 Results and the Intelligent Call Center

This is a completely different use case (See a recent KMWorld Articlehttps://allegrograph.com/knowledge-graphs-enhance-customer-experience-through-speed-and-accuracy/). Whereas the previous use case was very static, this one is highly dynamic. We analyze in real-time the text chats and spoken conversations between call center agents and customers. Our knowledge graph software provides real-time decision support to make the call center agents more efficient. N3 Results helps big tech companies to sell their high tech solutions, mostly cloud-based products and services but also helps their clients sell many other technologies and services.

The main challenge we tackle is to really deeply understand what the customer and agent are talking about. None of this can be solved by only simple entity extraction but requires elaborate rule-based and machine learning techniques. Just to give a few examples. We want to know if the agent talked about their most important talking points: that is, did the agent ask if the customer has a budget, or the authority to make a decision or a timeline about when they need the new technology

or whether they actually have expressed their need. But also whether the agent reached the right person, and whether the agent talked about the follow-up. In addition, if the customer talks about competing technology we need to recognize that and provide the agent in real-time with a battle card specific to the competing technology. And in order to be able to do the latter, we also analyzed the complicated marketing materials of the clients of N3.

Example 3: Complex Government Documents
Imagine a regulatory body with tens of thousands of documents. Where nearly every paragraph has reference to other paragraphs in the same document or other documents and the documents change over time. The goal here is to provide the end-users in the government with the right document given their current task at hand. The second goal is to keep track of all the changes in the documents (and the relationship between documents) over time.

The Document to Knowledge Graph Pipeline

Process Name	Input	Output
1. Custom Taxonomy Creation	Corpus Analytics, Taxonomy tool	A SKOS taxonomy containing concepts, concept hierarchy, prefLabels, altLabels.
2. Document Preparation	Documents (pdf, word, ppt, xlsx), Apache Tika, Spacy for XML cleanup	An XML version of each document
3. Extract Document Meta Data	Document + Apache Tika	JSON dictionary of the Document MetaData
4. XML-to-Triples	XML+JSON dictionary, XMLToTriples.py	Graph-based document tree with chapters, sections, and paragraphs as triples. Also includes meta data as triples
5. Entity-Extraction	Paragraphs + taxonomies + AllegroGraph Entity extract or external extractors	Concepts, persons, places, currencies. Connected to paragraphs
6. LOD Enrichment	Paragraphs + IBM Natural Language Understanding.	Concept categories and links to DBpedia and GeoNames, etc.
7. Complex Relationship and Event extraction.	Paragraphs + Taxonomy + Rules in Spacy or AllegroGraph	Complex events and relationships, References to other document sections.
8. NLP and ML	Chapters and paragraphs + all the tools described [here], but also using Spacy, Gensim, BERT, SciKit Learn.	Similarities, sentiment, query answering, smart search, text classification, word embeddings, abstracts
9. Versioning and Document tracking	Old + New document, compare.py	Old document in historic repository, new document in current, changed graph.
10. Statistical Relationships	Concepts + OddRatio.py or OddsRatio.cl	Statistical relationships between concepts.

Let us first give a quick summary in words of how we turn documents into a Knowledge Graph.

[1] Taxonomy Creation

Taxonomy of all the concepts important to the business using open source or commercial taxonomy builders. An available industry taxonomy is a good starting point for additional customizations.

[2] Document Preparation

We then take a document and turn it into an intermediate XML using Apache Tika. Apache Tika supports more than 1000 document types and although Apache Tika is a fantastic tool, the output is still usually not clean enough to create a graph from, so we use Spacy rules to clean up the XML to make it as uniform as possible.

[3] Extract Document MetaData

Most documents also contain document metadata (author, date, version, title, etc) and Apache Tika will also deliver the metadata for a document as a JSON object.

[4] XML to Triples

Our tools ingest the XML and metadata and transform that into a graph-based document tree. The document is the root and from that, it branches out into chapters, optionally sections, all the way down to paragraphs. The ultimate text content is in the paragraphs. In the following example we took the XML version of Noam Chomsky's book Media Control and turned that into a tree. The following shows a tiny part of that tree. We start with the Media Control node, then we show three (of the 11) chapters, for one chapter we show three (of the 6) paragraphs, and then we show the actual text in that paragraph. We sometimes can go even deeper to the level of sentences and tokens but for most projects that is overkill.

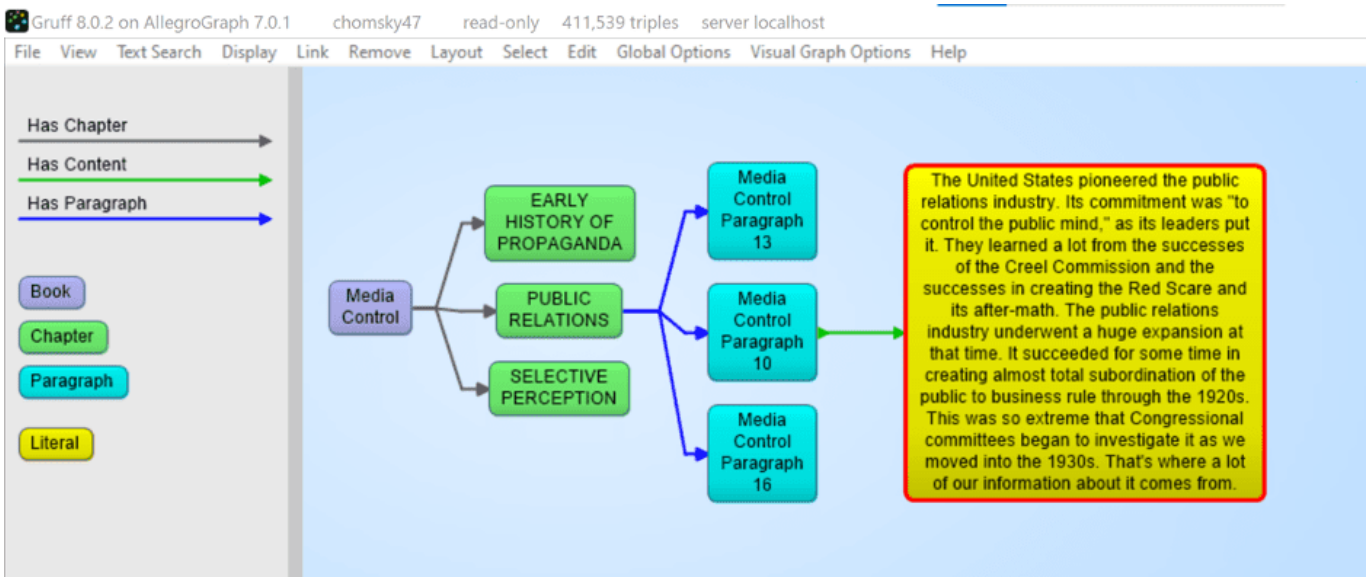

[5] Entity Extractor

AllegroGraph's entity extractor takes as input the text of each paragraph in the document tree and one or more of the taxonomies and returns recognized SKOS concepts based on prefLabels and altLabels. AllegroGraph's entity extractor is state of the art and especially powerful when it comes to complex terms like product names. We find that in our call center a technical product name can sometimes have up to six synonyms or very specific jargon. For example the Cisco product Catalyst 9000 will also be abbreviated as the cat 9k. Instead of developing altLabels for every possible permutation that human beings *will* use, we have specialized heuristics to optimize the yield from the entity extractor. The following picture shows 4 (of the 14) concepts discovered in paragraph 16. Plus one person that was extracted by IBM's NLU.

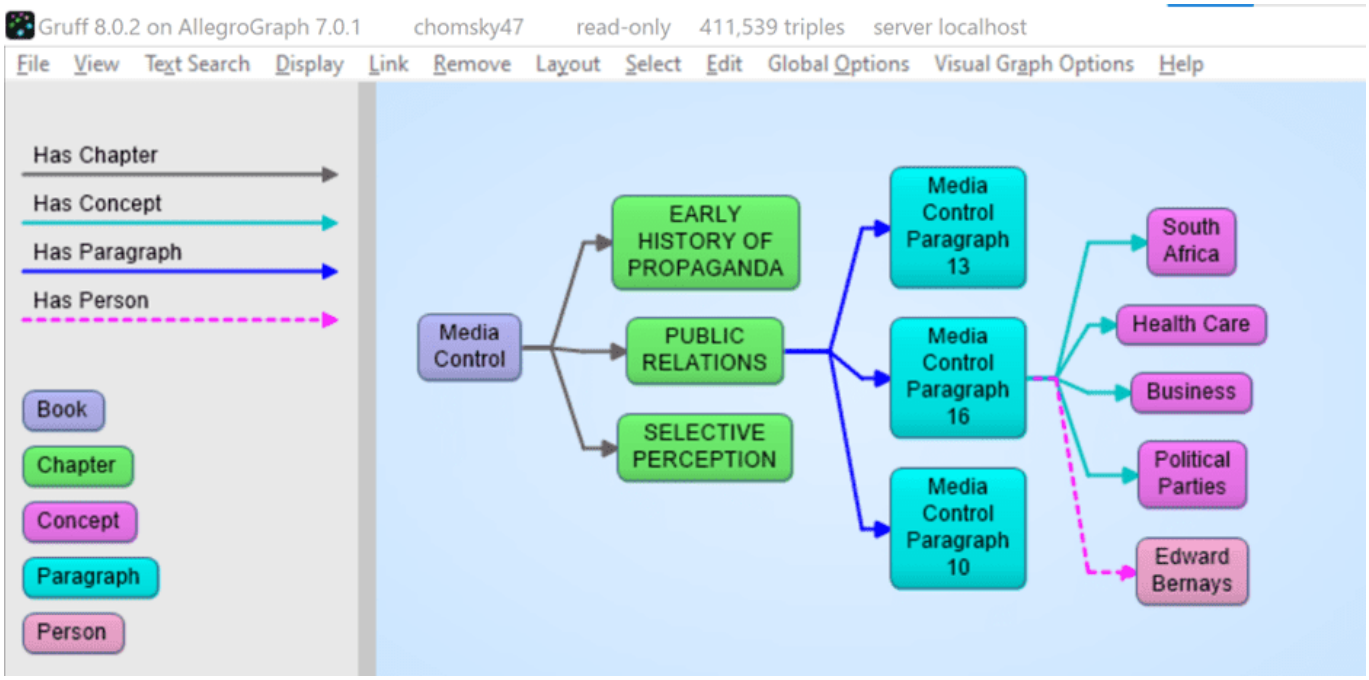

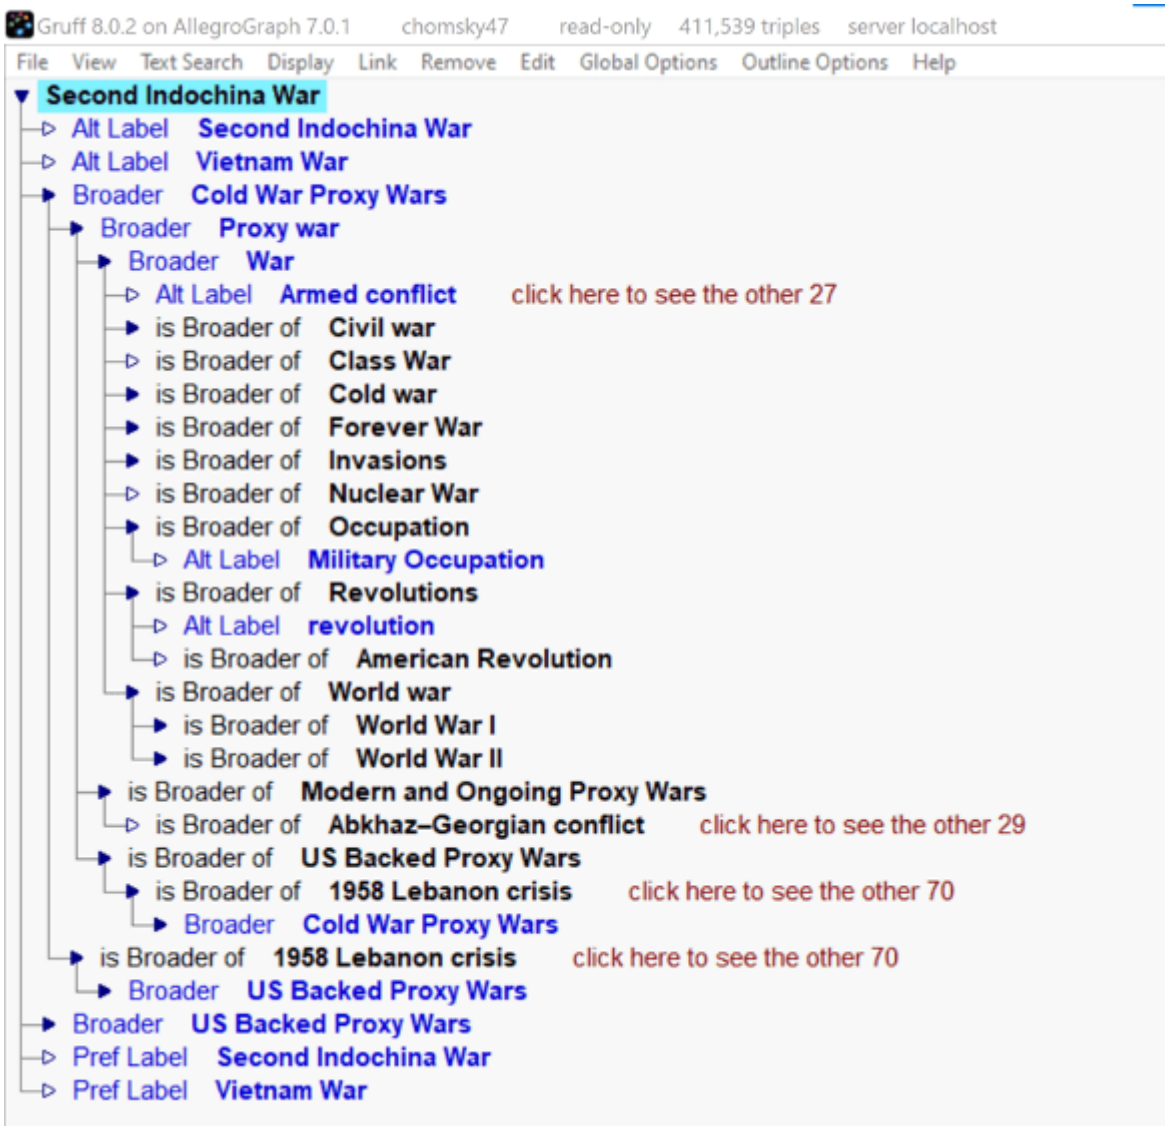

[6] Linked Data Enrichment

In many use cases, AllegroGraph can link extracted entities to concepts in the linked data cloud. The most prominent being DBpedia, wikidata, the census database, GeoNames, but also many Linked Open Data repositories. One tool that is very useful for this is IBM's Natural Language Understanding program but there are others available. In the following image we see that the Nelson Mandela entity (Red) is linked to the dbpedia entity for Nelson Mandela and that then links to the DBpedia itself. We extracted some of his spouses and a child with their pictures.

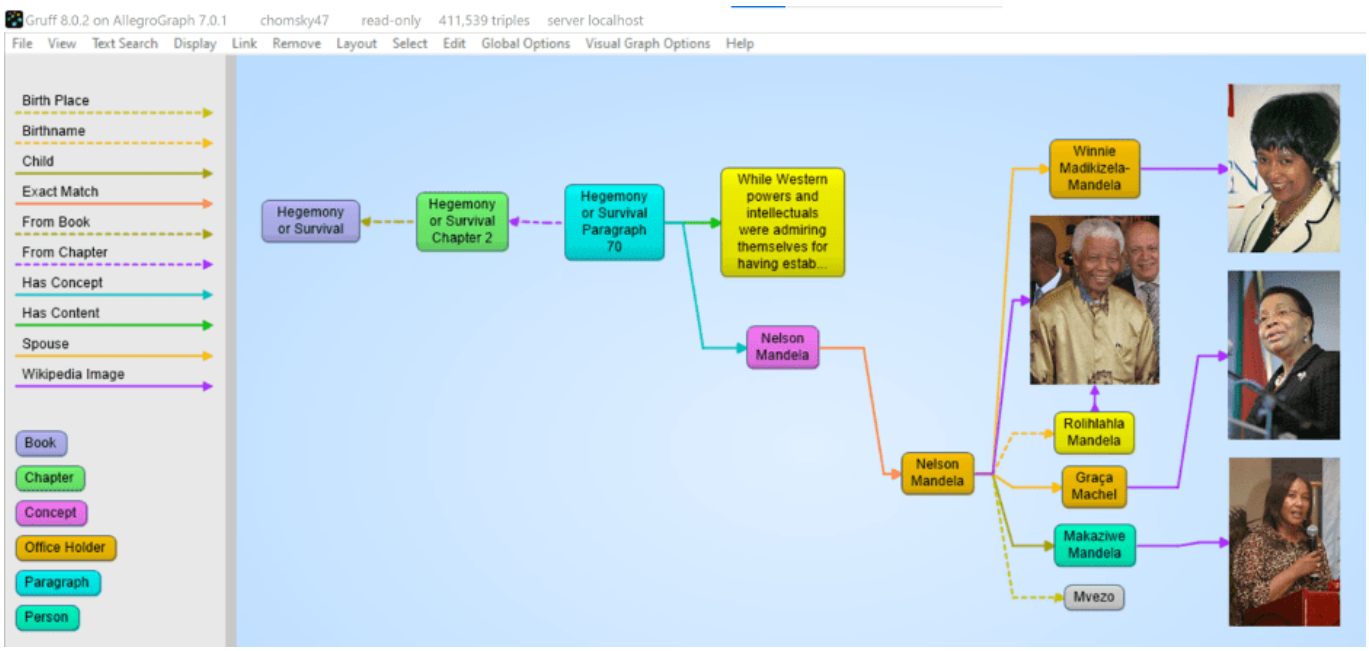

[7] Complex Relationship and Event Extraction

Entity extraction is a first good step to 'see' what is in your documents but it is just the first step. For example: how do you find in a text whether company C1 merged with company C2. There are many different ways to express the fact that a company fired a CEO. For example: Uber got rid of Kalanick, Uber and Kalanick parted ways, the board of Uber kicked out the CEO, etc. We need to write explicit symbolic rules for this or we need a lot of training data to feed a machine learning algorithm.

[8] NLP and Machine Learning

There are many many AI algorithms that can be applied in Document Knowledge Graphs. We provide best practices for topics like:

[a] Sentiment Analysis, using good/bad word lists or training data.
[b] Paragraph or Chapter similarity using statistical techniques like Gensim similarity or symbolic techniques where we just the overlap of recognized entities as a function of the size of a text.
[c] Query answering using word2vec or more advanced techniques like BERT
[d] Semantic search using the hierarchy in SKOS taxonomies.
[e] Summarization techniques for Abstractive or Extractive abstracts using Gensim or Spacy.

[9] Versioning and Document tracking

Several of our customers with Document Knowledge Graphs have noted the one constant in all of these KGs is that documents change over time. As part of our solution, we have created best practices where we deal with these changes. A crucial first step is to put each document in its own graph (i.e. the fourth element of every triple in the document tree is the document id itself). When we get a new version of a document the document ID changes but the new document will point back to the old version. We then compute which paragraphs stayed the same within a certain margin (there are always changes in whitespace) and we materialize what paragraphs disappeared in the new version and what new paragraphs appeared compared to the previous version. Part of the best practice is to put the old version of a document in a historical database that at all times can be federated with the ‘current’ set of documents.

Note that in the following picture we see the progression of a document. On the right hand side we have a newer version of a document 1100.161 with a chapter -> section -> paragraph -> contents where the content is almost the same as the one in

the older version. But note that the newer one spells ‘decision making’ as one word whereas the older version said ‘decision-making’. Note that also the chapter titles and the section titles are almost the same but not entirely. Also, note that the new version has a back-pointer (changed-from) to the older version.

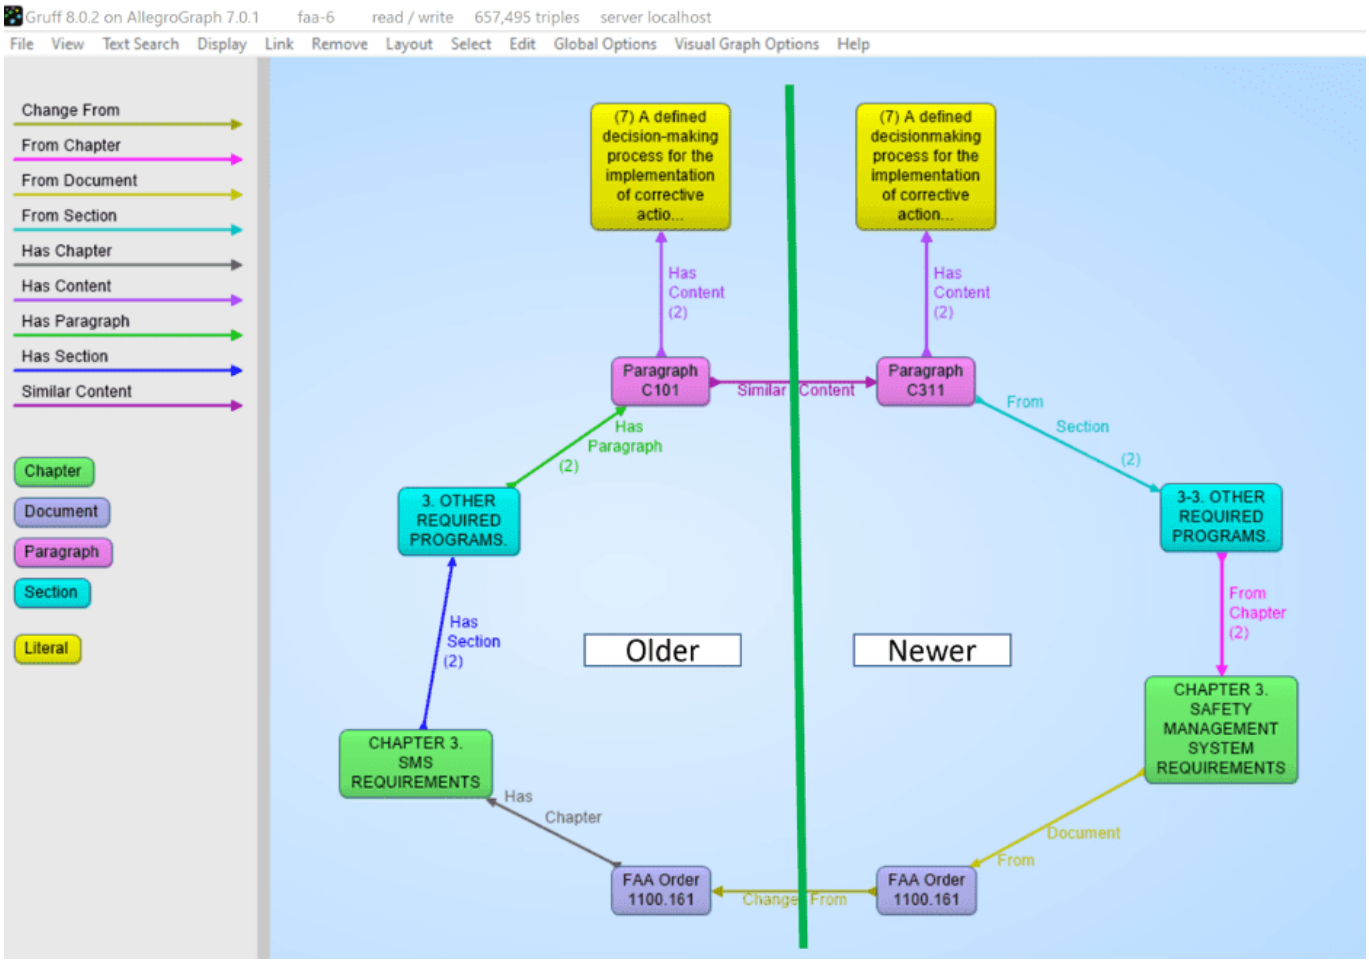

[10] Statistical Relationships

One important analytic one can do on documents is to look at the co-occurrence of terms. Although, given that certain words might occur more frequently in text, we have to correct the co-occurrence between words for the frequency of the two terms in a co-occurrence to get a better idea of the ‘surprisingness’ of a co-occurrence. The platform offers several techniques in Python and Lisp to compute these co-occurrences. Note that in the following picture we computed the odds ratios between recognized entities and so we see in

the following gruff picture that if Noam Chomsky talks about South Africa then the chances are very high he will also talk about Nelson Mandela.

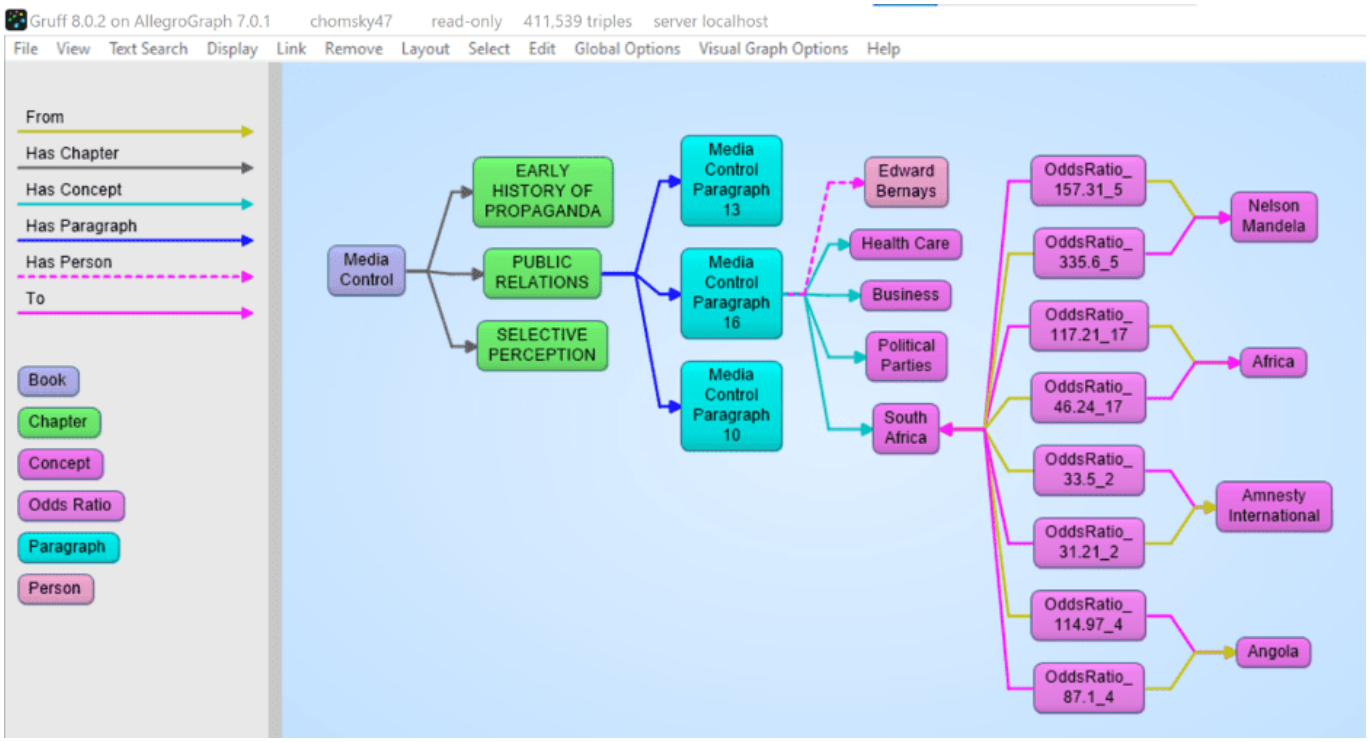

The Knowledge Graph Cookbook

Recipes for Knowledge Graphs that Work:

- Learn why and how to build knowledge graphs that help enterprises use data to innovate, create value and increase revenue. This practical manual is full of recipes and knowledge on the subject.
- Learn more about the variety of applications based on knowledge graphs.
- Learn how to build working knowledge graphs and which technologies to use.
- See how knowledge graphs can benefit different parts of your organization.

- Get ready for the next generation of enterprise data management tools.

Dr. Jans Aasman, CEO, Franz Inc. is interviewed in the Expert Opinion Section.

> **“KNOWLEDGE GRAPHS AREN’T WORTH THEIR NAME IF THEY DON’T ALSO LEARN AND BECOME SMARTER DAY BY DAY”** – Dr. Aasman

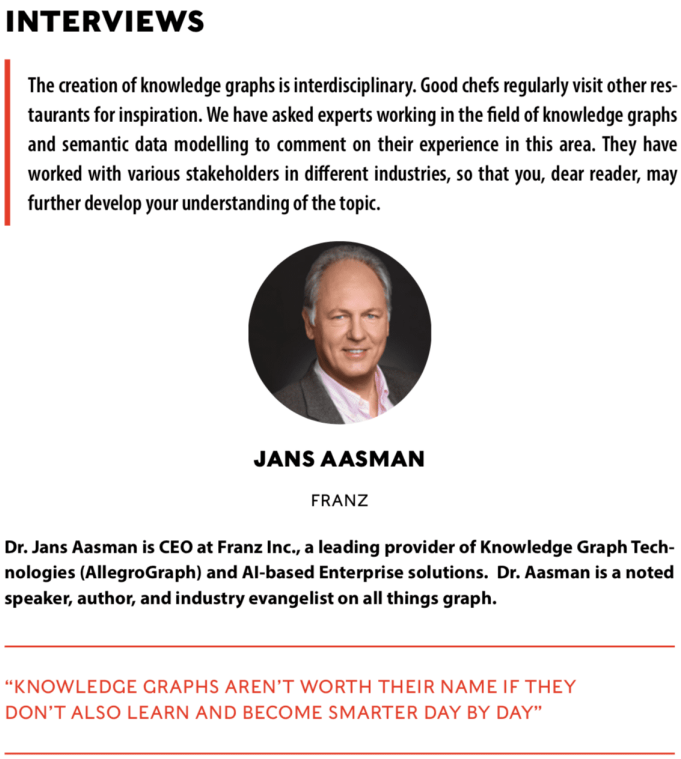

INTERVIEWS

The creation of knowledge graphs is interdisciplinary. Good chefs regularly visit other restaurants for inspiration. We have asked experts working in the field of knowledge graphs and semantic data modelling to comment on their experience in this area. They have worked with various stakeholders in different industries, so that you, dear reader, may further develop your understanding of the topic.

JANS AASMAN

FRANZ

Dr. Jans Aasman is CEO at Franz Inc., a leading provider of Knowledge Graph Technologies (AllegroGraph) and AI-based Enterprise solutions. Dr. Aasman is a noted speaker, author, and industry evangelist on all things graph.

“KNOWLEDGE GRAPHS AREN’T WORTH THEIR NAME IF THEY DON’T ALSO LEARN AND BECOME SMARTER DAY BY DAY”

Click here to get the book as free PDF or Kindle version.

100 Companies That Matter in Knowledge Management

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.

“Knowledge management software and services providers are embracing a fresh wave of technological innovation to address heightened expectations—among both customers and employees—for the right information to be delivered to the right people at the right time, said Tom Hogan, Group Publisher at KMWorld. “To showcase organizations that are advancing their products and capabilities to meet changing requirements, KMWorld created the annual list of 100 Companies That Matter in Knowledge Management.”

“We are honored to receive this acknowledgement for our efforts in delivering Enterprise Knowledge Graph Solutions,” said Dr. Jans Aasman, CEO, Franz Inc. “In the past year, we have seen demand for Enterprise Knowledge Graphs take off across industries along with recognition from top technology analyst firms that Knowledge Graphs provide the critical foundation for artificial intelligence applications and predictive analytics. Our AllegroGraph Knowledge Graph Platform Solution offers a unique comprehensive approach for helping companies accelerate the creation of Enterprise Knowledge Graphs that deliver new value to their organization.”