

Text Analytics Forum 2020 – KMWorld Connect

Join us November 17, 2020 – Text Analytics has the ability to add depth, meaning, and intelligence to any organization's most under-utilized resource – text. Through text analytics, enterprises can unlock a wealth of information that would not otherwise be available. Join us as we explore the power of text analytics to provide relevant, valuable, and actionable data for enterprises of all kinds.

Jans Aasman to present – Analyzing Spoken Conversations for Real-Time Decision Support in Mission-Critical Applications

November 17, 2020 at 2PM Eastern

Sharing Ontologies Globally To Speed Science And Healthcare Solutions – OntoPortal

International Ontology Sharing Is Becoming A Reality

A consortium of researchers recently formed an organization dedicated to standardizing how scientists define their ontologies, which are essential for retrieving datasets as well as understanding and reproducing research. The group called OntoPortal Alliance is creating a public repository of internationally shared domain-specific ontologies. All the repositories will be managed with a common OntoPortal

appliance that has been tested with AllegroGraph Semantic Knowledge Graph software. This enables any OntoPortal adopter to get all the power, features, maintainability, and support benefits that come from using a widely adopted, state-of-the-art semantic knowledge graph database.

Read the full article at [HealthIT Outcomes](#) –

As Dr. Jans Aasman, CEO of Franz Inc. explains, “When building a Knowledge Graph as your enterprise’s single source of truth, it’s critical to include ontologies and taxonomies. AI applications and complex reasoning analytics require information from both databases and knowledge bases that contain domain information, taxonomies, and ontologies to solve complex questions. To make this possible, we developed a novel hybrid sharding technology called FedShard, which facilitates the combination of data and knowledge required by applications like Montefiore’s PALM. But this approach is not unique or specific to Healthcare, it is applicable in many other industries, which is why we are excited about OntoPortal’s plans to bring sharing of domain ontologies to a broad audience.”





Knowledge Graphs: A Single Source of Truth for the Enterprise



The notion of a “single source of truth” for the enterprise has been the proverbial moving goalpost for generations of CIOs, only to be waylaid by brittle technology and unending legacy systems. Truth-seeking visions rebuffed by technological trends have continuously confounded business

units trying to achieve growth and market penetration. But technology innovation has finally led us to a point where CIOs can now deliver that truth.

Graphing the Truth

Knowledge graphs possess the power to deliver a single source of truth by linking together any assortment of data sources required, standardizing their diversity of data elements, and eliminating silos. They support the most advanced analytics options and decentralized transactions, which is why they’re now deployed as systems of records for some of the most

significant, mission-critical use cases affecting our population.

Because they scale to include almost any number of applications – and link to other knowledge graphs as well – these repositories are the ideal solution for real-time information necessary to inform business users' performances with concrete, data-supported facts. Most importantly, users can get an exhaustive array of touchpoints pertaining to any customer, product, or interaction with an organization from the knowledge graph, making it a single source of truth.

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

Gartner Hype Cycle for AI – Knowledge Graphs

According to Gartner's 2020 Hype Cycle for Artificial Intelligence – Despite the global impact of COVID-19, 47% of artificial intelligence (AI) investments were unchanged since the start of the pandemic and 30% of organizations actually planned to increase such investments, according to a Gartner poll. Only 16% had temporarily suspended AI investments, and just 7% had decreased them.

“AI is starting to deliver on its potential and its benefits for businesses are becoming a reality”

Gartner's – AI Hype Cycle Article

The Hype Cycle growth is consistent with Franz's customer interest in our Enterprise Knowledge Graph Solutions – Read our recent White Paper.

Hype Cycle for Artificial Intelligence, 2020



gartner.com/SmarterWithGartner

Source: Gartner
 © 2020 Gartner, Inc. and/or its affiliates. All rights reserved. Gartner and Hype Cycle are registered trademarks of Gartner, Inc. and its affiliates in the U.S.

Gartner®

Franz Inc. Named to Big Data

50 – Companies Driving Innovation in 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 Big Data 50 – Companies Driving Innovation in 2020 by Database Trends and Applications. The COVID-19 crisis has presented some new hurdles—but they are ones that many innovative companies are actively working to overcome. Forward-looking companies aren't sitting the year out waiting for the business climate to improve. They are actively seeking ways to expand their reach and take advantage of new opportunities.

"Whether the goal is surfacing timely insights through better analytics, streamlining processes using automation, improving security and governance with better controls, taking advantage of evolving opportunities using cloud services—or any of a range of other possibilities—leading IT companies are continually helping to make these plans a reality for their customers," said Tom Hogan, Group Publisher at Big Data Quarterly. "To help organizations navigate through a rapidly changing big data ecosystem, Big Data Quarterly has developed the Big Data 50, a list of companies driving innovation."

"Franz Inc. has a rich, innovative history and we are honored to receive this acknowledgement for our efforts in delivering key innovations with our Knowledge Graph Solutions," said Dr. Jans Aasman, CEO, Franz Inc. "In the past year, we have seen demand for Intelligent Data Fabrics take off across industries along with recognition from top technology analyst firms that Knowledge Graphs provide the critical foundation for Enterprise Wide Data Fabrics. 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 Data Fabric deployments across the Enterprise.”

AllegroGraph – Trend Setting Product for 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 a 2020 Trend Setting Product by KMWorld. The annual list is designed to spread the word about new and noteworthy solutions that are helping to improve organizational systems. KMWorld compiles this list of Trend-Setting Product offerings that are pushing the limits of what is possible with knowledge management.

“The COVID-19 pandemic has impacted the way organizations and individuals work, as well as how they support and communicate with customers and partners – making it clear that knowledge management is essential, now more than ever. By transforming data into information and then getting it to the people who need it, the resulting knowledge can be used for decisions that can make a significant impact,” said Tom Hogan, Group Publisher at KMWorld. “To help spread the word about noteworthy products that are helping to improve organizational systems, KMWorld created the Trend-Setting Products list.”

“Franz Inc. is continually innovating and we are honored to receive this acknowledgement for our efforts in setting the pace for Knowledge Graph Solutions,” said Dr. Jans Aasman, CEO, Franz Inc. “In the past year, we have seen demand for Intelligent Data Fabrics take off across industries along with

recognition from top technology analyst firms that Knowledge Graphs provide the critical foundation for Enterprise Wide Data Fabrics. 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 Data Fabric deployments across the Enterprise.”

AllegroGraph product profile on KMWorld.

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.

Using Microsoft Power BI with AllegroGraph

There are multiple methods to integrate AllegroGraph SPARQL results into Microsoft Power BI. In this document we describe two best practices to automate queries and refresh results if you have a production AllegroGraph database with new streaming data:

The first method uses Python scripts to feed Power BI. The second method issues SPARQL queries directly from Power BI using POST requests.

Method 1: Python Script:

Assuming you know Python and have it installed locally, this is definitely the easiest way to incorporate SPARQL results into Power BI. The basic idea of the method is as follows: First, the Python script enables a connection to your desired AllegroGraph repository. Then we utilize AllegroGraph's Python API within our script to run a SPARQL query and return it as a Pandas dataframe. When running this script within Power BI Desktop, the Python scripting service recognizes all unique dataframes created, and allows you to import the dataframe into Power BI as a table, which can then be used to create visualizations.

Requirements:

1. You must have the AllegroGraph Python API installed. If you do not, installation instructions are here: <https://franz.com/agraph/support/documentation/current/python/install.html>
2. Python scripting must be enabled in Power BI Desktop.

Instructions to do so are here:
<https://docs.microsoft.com/en-us/power-bi/connect-data/desktop-python-scripts>

a) As mentioned in the article, pandas and matplotlib must be installed. This can be done with 'pip install pandas' and 'pip install matplotlib' in your terminal.

The Process:

Once these requirements have been met, create a Python file with whatever script editor you usually use. The following code will create a connection to your desired repository. For this example, we will be using the Kennedy dataset that is available with the AllegroGraph distribution (See the 'Tutorial' directory). Load the Kennedy.ntriples file into your running AllegroGraph. (Replace the '****' in the code with your corresponding username and password.)

#the necessary imports

```
import os

from franz.openrdf.connect import ag_connect

from franz.openrdf.query.query import QueryLanguage

import pandas as pd
```

#connect to your agraph repository

```
def setup_env_var(var_name, value, description):

    os.environ[var_name] = value

    print("{}: {}".format(description, value))

    setup_env_var('AGRAPH_HOST', 'localhost', 'Hostname')
```

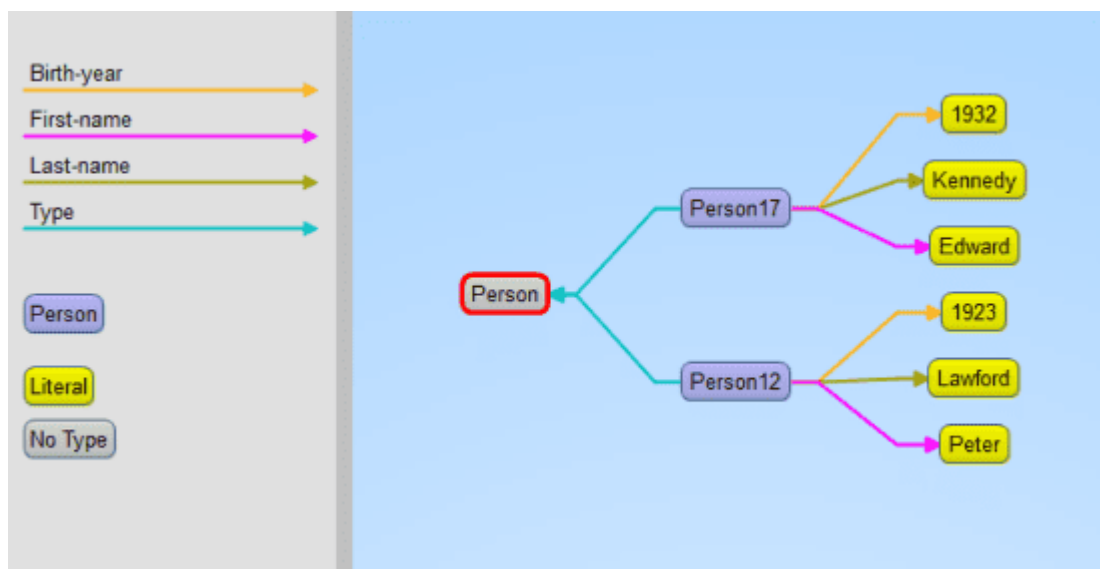
```
setup_env_var('AGRAPH_PORT', '10035', 'Port')

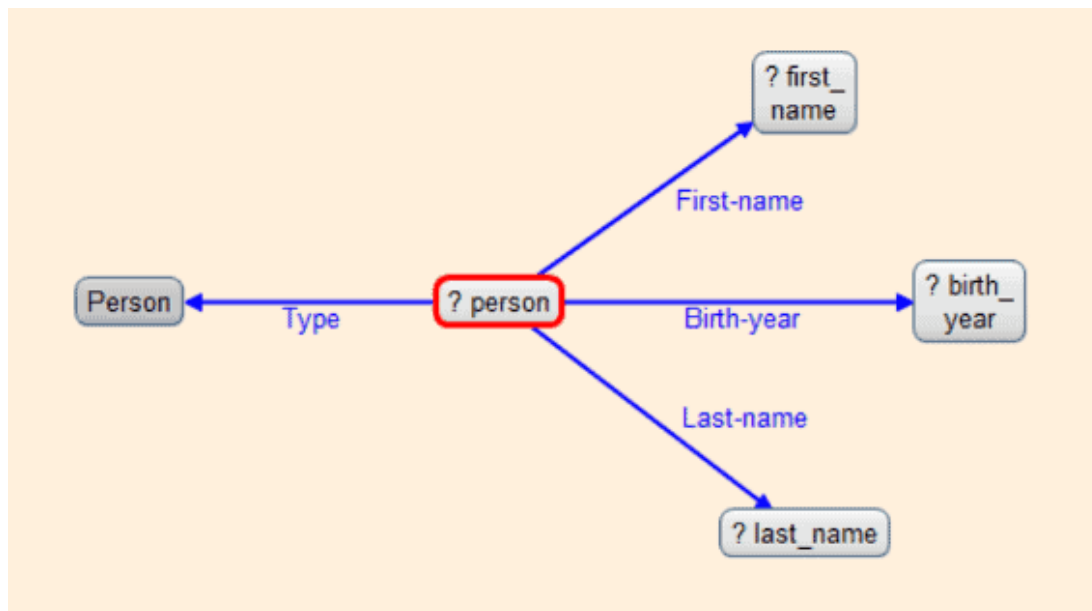
setup_env_var('AGRAPH_USER', '****', 'Username')

setup_env_var('AGRAPH_PASSWORD', '****', 'Password')

conn = ag_connect('kennedy', create=False, clear=False)
```

2. We then want to create a query. For this example, we will first show what our data looks like, what the visual query of the information is, and what the written query looks like. With the following query we want every person's first and last names, as well as their birth years. Here is a small portion of the data visualized in Gruff, and then the visualization of the query:





3. Then add the written query to the python script as a variable string (we added an additional line to the query to sort on birth year). Next use the API functionality to simply execute the query and turn the results into a pandas dataframe.

```
query = """select ?person ?first_name ?last_name ?birth_year
where
{ ?person <http://www.franz.com/simple#first-name> ?first_name
;
    <http://www.franz.com/simple#birth-year> ?birth_year
;
    rdf:type <http://www.franz.com/simple#person> ;
    <http://www.franz.com/simple#last-name> ?last_name .
}
order by desc(?birth_year)"""
```

```
with conn.executeTupleQuery(query) as result:
    df = result.toPandas()
```

When looking at the result, we see that we have a DataFrame!

```
[5] df.head()
```



	person	first_name	last_name	birth_year
0	<http://www.franz.com/simple#person70>	Kym	Smith	1972
1	<http://www.franz.com/simple#person63>	Molly	Stark	1968
2	<http://www.franz.com/simple#person64>	Rory	Kennedy	1968
3	<http://www.franz.com/simple#person62>	Douglas	Kennedy	1967
4	<http://www.franz.com/simple#person65>	Mark	Bailey	1967

4. Now we will use this script in Power BI. When in Power BI Desktop, go to 'Get Data' and look for the python script option. Then simply copy and paste your entire script into the text box, and run the script. In this case, our output looks like this:

Navigator

Display Options ▾

Python [1]

☒ df

df

person	first_name	last_name	birth_year
<http://www.franz.com/simple#person70>	Kym	Smith	
<http://www.franz.com/simple#person63>	Molly	Stark	
<http://www.franz.com/simple#person64>	Rory	Kennedy	
<http://www.franz.com/simple#person62>	Douglas	Kennedy	
<http://www.franz.com/simple#person65>	Mark	Bailey	
<http://www.franz.com/simple#person68>	Amanda	Smith	

5. Next simply 'Load' the data, and then you can use the Power BI Desktop interface to create whatever visualizations you want! If you do have a lot of additional operations to perform on your dataframe, we recommend doing these in your python script.

Method 2: POST Request:

For the SPARQL query via POST requests to work you need to url-encode the query. Every modern programming language will support that, but in our example we will be using Python

again. This method is better for when you do not have python locally installed or prefer a different programming language.

It is possible to send a GET request from Power BI, but once the results from the query reach a certain size, a POST request is required, which is confusing to do within the Power BI Desktop interface. The following steps will show you how to do SPARQL Queries using POST requests. It looks a bit odd but it works well.

The Process:

1. In your AG WebView create an 'anonymous' user. (Go to admin -> Users -> [add a user] -> and add 'anonymous' as username without adding a password). You can use these settings:

Users

anonymous [\[remove\]](#)

Roles: None

[\[suspend\]](#) [\[disable\]](#) [\[expire password\]](#)

☐ Superuser ☐ Start sessions ☐ Evaluate arbitrary code ☐ Control replication ☐ Two-phase commit

☒ Allow user attributes via HTTP header [x-user-attributes](#)

☐ Allow user attributes via SPARQL PREFIX [franzOption_userAttributes](#)

◦ [read/write on all](#) [\[remove\]](#)

Grant on catalog repository [\[ok\]](#)

Security Filters: [None](#) [\[add\]](#)

2. Go to your desired repository in WebView and Click on 'Queries' -> 'New'

3. Write a simple SPARQL query, and run it to make sure you get the correct response back.

4. In python create the following script: (Assuming your AllegroGraph is on your localhost port 10035 and your repo is called 'kennedy')

```
import urllib
```

```
def CreatePOSTquery(query):
    start = "http://anonymous:@localhost:10035/repositories/kennedy?queryL
n=SPARQL&limit=1000&infer=false&returnQueryMetadata=false&che
ckVariables=false&query="
    response = start + urllib.parse.quote(query)
    return response
```

This function url-encodes the query and attaches it to the POST request. Replace the 'localhost:10035' and 'kennedy' strings in the start variable with your corresponding data. Then, using the same query as our previous example, we create our url-encoded POST query:

```
query = """select ?person ?first_name ?last_name ?birth_year
where
{ ?person <http://www.franz.com/simple#first-name> ?first_name
;
    <http://www.franz.com/simple#birth-year> ?birth_year
;
    rdf:type <http://www.franz.com/simple#person> ;
    <http://www.franz.com/simple#last-name> ?last_name .
}
order by desc(?birth_year)"""

result = CreatePOSTquery(query)
print(result)
```

This gives us the following result:

[illegible]

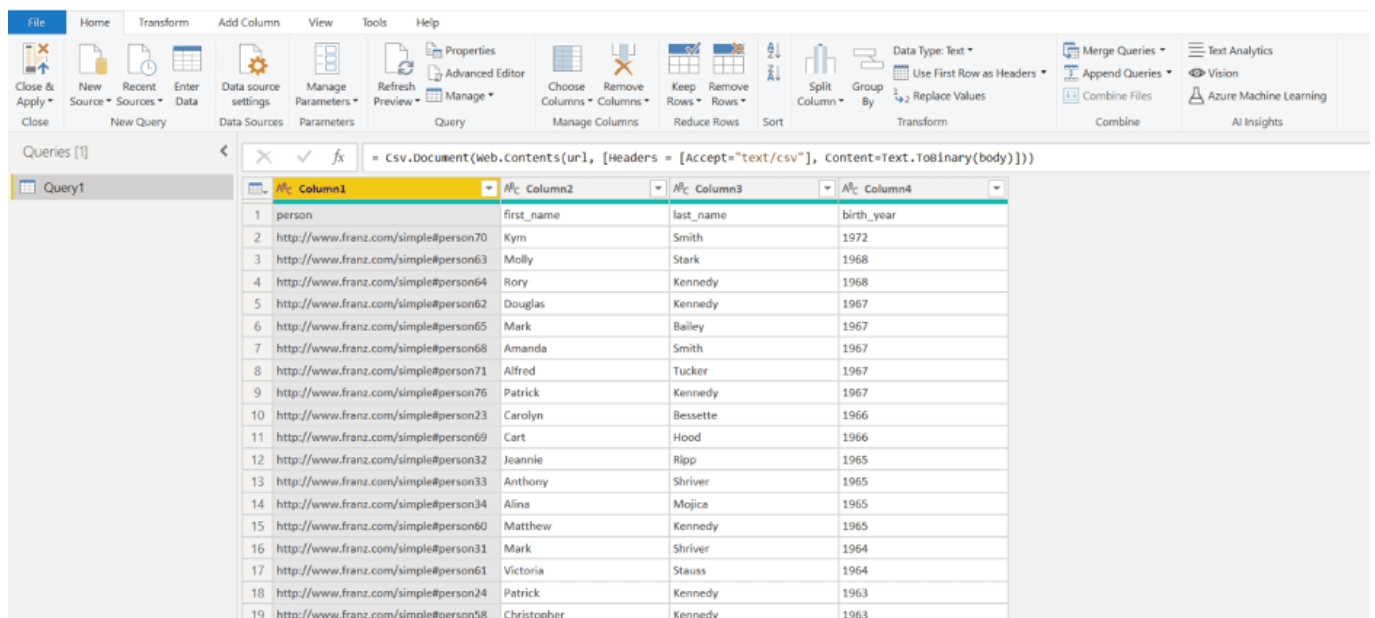
5. Within Power BI Desktop we go to 'Get data' and create a 'Blank query' and go into the 'Advanced Editor' window. Using the following format we will get our desired results (please note that due to the length of the url-encoded request, it did not all fit in the image. Copy and pasting into the url field works fine. The 'url' variable needs to be in quotes and have a comma at the end):

 Advanced Editor

Query1

```
let
    url = "http://anonymous:@localhost:10035/repositories/kennedy?queryLn=SPARQL&limit=1000&infer=false&returnQuery",
    body = "",
    Source = Csv.Document(Web.Contents(url, [Headers = [Accept="text/csv"], Content=Text.ToBinary(body)]))
in
    Source
```

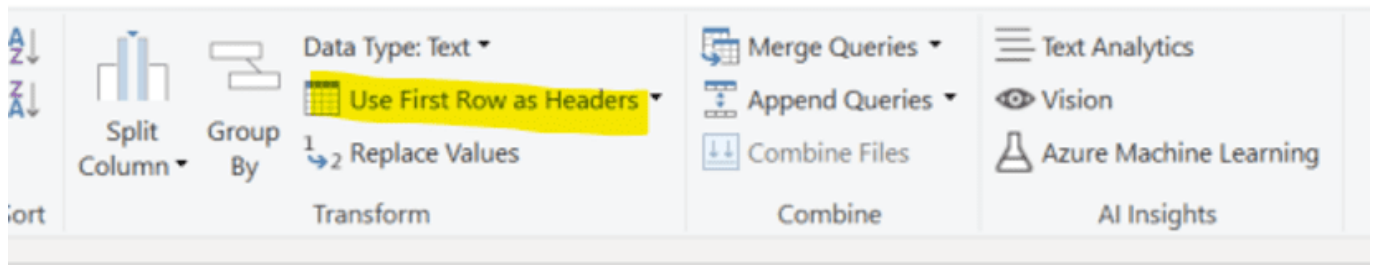
We see the following results:



The screenshot shows the Power BI Desktop interface with the Advanced Editor window open. The query in the editor is a CSV document query that fetches data from a local SPARQL endpoint. The results are displayed in a table with 4 columns: Column1, Column2, Column3, and Column4. The first row of the table contains the headers: person, first_name, last_name, and birth_year. The subsequent rows contain 19 data points, each representing a person with their ID, first name, last name, and birth year.

Column1	Column2	Column3	Column4
person	first_name	last_name	birth_year
http://www.franz.com/simple#person70	Kym	Smith	1972
http://www.franz.com/simple#person63	Molly	Stark	1968
http://www.franz.com/simple#person64	Rory	Kennedy	1968
http://www.franz.com/simple#person62	Douglas	Kennedy	1967
http://www.franz.com/simple#person65	Mark	Bailey	1967
http://www.franz.com/simple#person68	Amanda	Smith	1967
http://www.franz.com/simple#person71	Alfred	Tucker	1967
http://www.franz.com/simple#person76	Patrick	Kennedy	1967
http://www.franz.com/simple#person23	Carolyn	Bessette	1966
http://www.franz.com/simple#person69	Cart	Hood	1966
http://www.franz.com/simple#person32	Jeannie	Ripp	1965
http://www.franz.com/simple#person33	Anthony	Shriver	1965
http://www.franz.com/simple#person34	Alina	Mojica	1965
http://www.franz.com/simple#person60	Matthew	Kennedy	1965
http://www.franz.com/simple#person31	Mark	Shriver	1964
http://www.franz.com/simple#person61	Victoria	Stauss	1964
http://www.franz.com/simple#person24	Patrick	Kennedy	1963
http://www.franz.com/simple#person58	Christopher	Kennedy	1963

6. One last step is to turn the top row into the column names, which can be achieved by pressing the 'Use first row as headers':



The best part about both of these methods is that once the query has been created, Power BI can refresh the visuals using the same queries if your data changed. This can be achieved by scheduling refreshes within the Power BI Desktop interface (<https://docs.microsoft.com/en-us/power-bi/connect-data/refresh-data#configure-scheduled-refresh>)

Please send any questions or issues to: support@franz.com

Franz Inc. Named an AI 50 Company by KMWorld

AllegroGraph Powering Intelligent Knowledge Graph Solutions

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 AI 50 – Companies Empowering Intelligent Knowledge Management Companies 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.

This list recognizes companies based on their presence, execution, vision and innovation in delivering products and services to the marketplace.

“As the drive for digital transformation becomes an imperative for companies seeking to compete and succeed in all industry sectors, intelligent tools and services are being leveraged to enable speed, insight, and accuracy,” said Tom Hogan, Group Publisher at KMWorld. “To showcase organizations that are incorporating AI and an assortment of related technologies—including natural language processing, machine learning, and computer vision—into their offerings, KMWorld created the “AI 50: The Companies Empowering Intelligent Knowledge Management.”

“Franz Inc. has a rich history in AI and we are honored to receive this acknowledgement for our efforts in delivering AI Knowledge Graph Solutions,” said Dr. Jans Aasman, CEO, Franz Inc. “In the past year, we have seen demand for Intelligent Data Fabrics take off across industries along with recognition from top technology analyst firms that Knowledge Graphs provide the critical foundation for Enterprise Wide Data Fabrics. 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 Data Fabric deployments across the Enterprise.”

Gartner’s Top 10 Trends in Data and Analytics for 2020 noted “Relationships form the foundation of data and analytics value. By 2023, graph technologies will facilitate rapid contextualization for decision making in 30% of organizations worldwide. Graph analytics is a set of analytic techniques that allows for the exploration of relationships between entities of interest such as organizations, people and transactions. Data and analytics leaders need to evaluate opportunities to incorporate graph analytics into their analytics portfolios and applications to uncover hidden

patterns and relationships. In addition, consider investigating how graph algorithms and technologies can improve your AI and ML initiatives.” (Source: Gartner, Top 10 Trends in Data and Analytics for 2020, June 9, 2020).

“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 Data Fabrics 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 DBTA’s 100 Companies That Matter Most in

Data

AllegroGraph was also recently named to DBTA's 100 Companies That Matter Most in Data. The DBTA 100 showcases organizations that delivering solutions for customers to meet the need for real-time, data-driven insights.

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.

About Franz Inc.

Franz Inc. is an early innovator in Artificial Intelligence (AI) and leading supplier of Semantic Graph Database technology with expert knowledge in developing and deploying Knowledge Graph solutions. The foundation for Knowledge Graphs and AI lies in the facets of semantic technology provided by AllegroGraph with FedShard and Allegro CL. The ability to rapidly integrate new knowledge is the crux of the Knowledge

Graph and Franz Inc. provides the key technologies and services to address your complex challenges. Franz Inc. is your Knowledge Graph technology partner.

All trademarks and registered trademarks in this document are the properties of their respective owners.

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.”