

Gruff Time Machine Tutorial



Here is an example for trying out the new time slider in Gruff's graph view. It uses triples from crunchbase.com that contain a history of corporate acquisitions and funding events over several years. Gruff's time bar allows you to examine those events chronologically, and also to display only the nodes that have events within a specified date range.

- Download the Crunchbase triples from the bottom of the Gruff download page at <https://allegrograph.com/products/gruff/>
- Create a new triple-store and used "File | Load Triples | Load N-Triples" to load that triples file into the new triple-store. Use "File | Commit" to ensure that the loaded triples get saved.
- Select "Visual Graph Options | Time Bar | Momentary Time Predicates" and paste the following five predicate IRIs into the dialog that appears. The time bar will then work with the date properties that are provided by these predicates, whenever you are browsing this particular triple-store.

http://www.franz.com/hasfunded_at

http://www.franz.com/hasfirst_funding_at

http://www.franz.com/hasfounded_at
http://www.franz.com/haslast_funding_at
http://www.franz.com/hasacquired_at

- Select “View | Optional Graph View Panes | Show Time Bar” to reveal the time bar at the bottom of the graph view. The keyboard shortcut for this command is Shift+A to allow quickly toggling the time bar on and off.
- Select “Display | Display Some Sample Triples” to do just that. The time bar will now display a vertical line for each of the requested date properties of the displayed nodes. Moving the mouse cursor over these “date property markers” will display more information about those events.
- Click down on the yellow-orange rectangle at the right end of the time bar and drag it to the left. This will make the “time filter range” smaller, and nodes that have date properties that are no longer in this range will temporarily disappear from the display. They will reappear if you drag the slider back to the right or toggle the time bar back off.

For more information, the full time bar introduction is in the

Gruff documentation under the command “View | Optional Graph View Panes | Show Time Bar”.

Check out the “Chart Widget” for showing date properties of the visible nodes.



Why Smart Cities Need AI Knowledge Graphs

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



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

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

Despite these advantages, smart cities demand extensive data management. Consistent data integration from multiple

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

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

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



2019 Trends In The Internet

Of Things: The Makings Of An Intelligent IoT

AI Business – December 2018

2019 will be a crucial year for the Internet of Things for two reasons. Firstly, many of the initial predictions for this application of big data prognosticated a future whereby at the start of the next decade there would be billions of connected devices all simultaneously producing sensor data. The IoT is just a year away from making good on those claims.

Dr. Jans Aasman, Franz's CEO was quoted by the author:

The IIoT is the evolution of the IoT that will give it meaning and help it actualize the number of connected devices forecast for the start of the next decade. The IIoT will encompass smart cities, edge devices, wearables, deep learning and classic machine learning alongside lesser acknowledged elements of AI in a basic paradigm in which, according to Franz CEO Jans Aasman, "you can look at the past and learn from certain situations what's likely going to happen. You feed it in your [IoT] system and it does better... then you look at what actually happened and it goes back in your machine learning system. That will be your feedback loop."

Although deep learning relies on many of the same concepts as traditional machine learning, with "deep learning it's just that you do it with more computers and more intermediate layers," Aasman said, which results in higher accuracy levels.

The feedback mechanism described by Aasman has such a tremendous capacity to reform data-driven businesses because of the speed of the iterations provided by low latency IIoT data.

One of the critical learning facets the latter produces

involves optimization, such as determining the best way to optimize route deliveries encompassing a host of factors based on dedicated rules about them. “There’s no way in [Hades] that a machine learning system would be able to do the complex scheduling of 6,000 people,” Aasman declared. “That’s a really complicated thing where you have to think of every factor for every person.”

However, constraint systems utilizing multi-step reasoning can regularly complete such tasks and the optimization activities for smart cities. Aasman commented that for smart cities, semantic inferencing systems can incorporate data from traffic patterns and stop lights, weather predictions, the time of year, and data about specific businesses and their customers to devise rules for optimal event scheduling. Once the events actually take place, their results—as determined by KPIs—can be analyzed with machine learning to issue future predictions about how to better those results in what Aasman called “a beautiful feedback loop between a machine learning system and a rules-based system.”

In almost all of the examples discussed above, the IIoT incorporates cognitive computing “so humans can take action for better business results,” Aasman acknowledged. The means by which these advantages are created are practically limitless.

Read the Full Article at [AI Business](#).

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](#)

Knowledge Graphs – The path to true AI

Published in SD Times – December, 2018



Knowledge is the foundation of intelligence— whether artificial intelligence or conventional human intellect. The understanding implicit in intelligence, its application towards business problems or personal ones, requires knowledge of these problems (and potential solutions) to effectively overcome them.

The knowledge underpinning AI has traditionally come from two distinct methods: statistical reasoning, or machine learning, and symbolic reasoning based on rules and logic. The former approach learns by correlating inputs with outputs for increasingly progressive pattern identification; the latter approach uses expert, human-crafted rules to apply to particular real-world domains.

Read the full article at [SD Times](#).

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.



The system also acts as a giant feedback-response or learning loop which learns from the data collected via user/system interactions.

Montefiore Health, Franz, Intel and PwC research, 2017

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

.

Franz and Semantic Web Co. Partner to Create a Noam Chomsky Knowledge Graph

Press Release – September 10, 2018

First Semantic Knowledge Graph for a Public Figure will Semantically Link Books, Interviews, Movies, TV Programs and Writings from the Most Cited U.S. Scholar

OAKLAND, Calif. and VIENNA, Austria – Franz Inc., an early innovator in Artificial Intelligence (AI) and leading supplier of Semantic Graph Database technology, AllegroGraph, for Knowledge Graphs, and Semantic Web Company, developers of the PoolParty Semantic Suite and leading provider of Semantic AI solutions, today announced a partnership to develop the Noam Chomsky Knowledge Graph. This project is the first aimed at

connecting all the works from a public figure and turning the linked information into a searchable and retrievable resource for the public.

The Noam Chomsky Knowledge Graph project will organize and semantically link the vast knowledge domain surrounding Noam Chomsky, the founder of modern linguistics, a founder of cognitive science, and a major figure in analytic philosophy as well as an American linguist, philosopher, historian and social critic. Chomsky is currently an Institute Professor Emeritus at the Massachusetts Institute of Technology (MIT) and laureate professor at the University of Arizona. He has received many awards including Distinguished Scientific Contribution Award of the American Psychological Association, the Kyoto Prize in Basic Sciences, the Helmholtz Medal, the Dorothy Eldridge Peacemaker Award, and the Ben Franklin Medal in Computer and Cognitive Science.

“Noam Chomsky is one of the most brilliant minds of our generation,” said Fred Davis. Executive Director of the Chomsky Knowledge Graph project, “His body of work is tremendously valuable to people across many disciplines. Our goal is to make Chomsky’s work searchable in the context of topics and concepts, readable in excerpts, and easily available to journalists, scientists, technologists, students, philosophers, and historians as well as the general public.”

The Noam Chomsky Knowledge Graph will link to over 1,000 articles and over 100 books that Chomsky has authored about linguistics, mass media, politics and war. Hundreds of Chomsky’s media interviews, which aired on television, print and online will be part of the Knowledge Graph as well as more than a dozen Chomsky movies including: Is the Man who is Tall Happy?, Manufacturing Consent, Programming the Nation? Hijacking Catastrophe: 911 Fear and the Selling of American Empire. The content will be made available by searching the Knowledge Graph for specific titles, related topics and concepts.

Since the project is based on the latest and most advanced technologies, the data will be also available as machine-readable data (Linked Data) in order to be fed into smart applications, intelligent chatbots, and question/ answering machines – as well as other AI and data systems.

The Internet Archive, the world's largest digital lending library, will host Noam Chomsky's books, movies, and other content – enabling public access to his works and marking the first integration between the Internet Archive and a public Knowledge Graph.

"We are thrilled to be working on this momentous project," said Dr. Jans Aasman, CEO of Franz Inc. "Noam Chomsky is the ideal person to fulfill the vision of a Public Figure Knowledge Graph. We are looking forward to collaborating with the Semantic Web Company and Fred Davis on this exciting project."

"Knowledge Graphs are becoming increasingly important for addressing various data management challenges in industries such as financial services, life sciences, healthcare or energy," said Andreas Blumauer, CEO and founder of Semantic Web Company. "The application of Knowledge Graphs to public figures, such as Noam Chomsky, will offer a unique opportunity to link concepts and ideas to form new ideas and possible solutions."

About Knowledge Graphs

A Knowledge Graph represents a knowledge domain and connects things of different types in a systematic way. Knowledge Graphs encode knowledge arranged in a network of nodes and links rather than tables of rows and columns. People and machines can benefit from Knowledge Graphs by dynamically growing a semantic network of facts about things and use it for data integration, knowledge discovery, and in-depth analyses.

Gartner recently identified Knowledge Graphs as a key new technology in both their Hype Cycle for Artificial Intelligence and Hype Cycle for Emerging Technologies. Gartner's Hype Cycle for Artificial Intelligence, 2018 states, "The rising role of content and context for delivering insights with AI technologies, as well as recent knowledge graph offerings for AI applications have pulled knowledge graphs to the surface."

Knowledge Graphs are the Foundation for Artificial Intelligence

The foundation for AI lies in the facets Knowledge Graphs and semantic technology provided by Franz and Semantic Web Company. The Franz AllegroGraph Semantic Graph database provides the core technology environment to enrich and contextualized the understanding of data. The ability to rapidly integrate new knowledge is the crux of the Knowledge Graph and depends entirely on semantic technologies.

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 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. For more information, visit www.franz.com.

About Semantic Web Company

Semantic Web Company is the leading provider of graph-based metadata, search and analytic solutions. The company is the vendor of PoolParty Semantic Suite, one of the most renowned semantic software platforms on the global market. Among many

other customers, The World Bank, AT&T, Deutsche Telekom, and Pearson benefit from linking structured and unstructured data. In 2018, the Semantic Web Company has been named to KMWorld's "100 companies that matter in Knowledge Management." For more information about PoolParty Semantic Suite, please visit <https://ww.poolparty.biz>

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