

Predictions: Quantum AI, Graph Neural Networks, and Personal Data Pods



Dr. Jans Aasman, CEO, Franz Inc.

The new year will bring us many exciting developments in data-enabling technologies including the merging of artificial intelligence with quantum computing and graph neural networks, which will power extremely complex, next-generation algorithms. Knowledge graphs will become lego-like with the ability

to be plugged into diverse applications. With mounting concern over social media sites using personal data, expect new ways for users to regain control with “personal data pods.” Enterprises will interweave graphs with document and time-series databases to create a single enterprise-wide data fabric.

Quantum AI Environments Will Emerge

recent advances in quantum computing, in 2022, we will start to see the convergence of quantum computing with artificial intelligence, knowledge graphs, and programming languages. These distinct technologies will start to morph into a single computing environment operating in one memory space as a fully integrated solution. The separation between programming and AI/analytics will begin to blur as developers use Quantum-based computer languages to generate incredibly complex, next-generation AI algorithms and applications that result in new discoveries based on the quantum acceleration of machine learning and deep learning.

Graph Neural Networks (GNNs) Will Advance AI Reasoning

In the past few years, organizations have experienced the advantages of combining graphs with artificial intelligence. In 2022 and beyond, leading companies will apply machine learning's advanced pattern matching to graph neural networks (GNNs), which are complex high-dimensional, non-Euclidian datasets. By fusing GNN reasoning capabilities with classic semantic inferencing available in AI knowledge graphs, organizations will get two forms of reasoning in one framework. Automatically mixing and matching these two types of reasoning is the next level of AI and produces the best prescriptive outcomes. This "total AI" is swiftly becoming necessary to tackle enterprise-scale applications of mission-critical processes like predicting equipment failure, optimizing healthcare treatment, and maximizing customer relationships.

Read the full article to see the additional predictions:

Facebook Users Will Regain Control Through Personal Data Pods

Knowledge Graphs Will Become Composable

Graph, Document and Time-Series Databases will Dominate by 2030

AllegroGraph v7.2 – Now Available (GNN, Virtual Graphs, Spark, and Kafka)

AllegroGraph 7.2, provides organizations with essential Data Fabric tools, including Graph Neural Networks, Graph Virtualization, Apache Spark graph analytics, and streaming

graph pipelines. These new capabilities exemplify AllegroGraph's leadership in empowering data analytics professionals to derive business value out of Knowledge Graphs.

Graph Neural Networks

With AllegroGraph 7.2, users can create Graph Neural Networks (GNNs) and take advantage of a mature AI approach for Knowledge Graph enrichment via text processing for news classification, question and answer, search result organization, event prediction, and more. GNNs created in AllegroGraph enhance neural network methods by processing the graph data through rounds of message passing, as such, the nodes know more about their own features as well as neighbor nodes. This creates an even more accurate representation of the entire graph network. AllegroGraph GNNs advance text classification and relationship extraction for enhancing enterprise-wide Data Fabrics.

Graph Virtualization

AllegroGraph 7.2 allows users to easily virtualize data as part of their AllegroGraph Knowledge Graph solution. When graphs are virtual, the data remains in the source system and is easily linked and queried with other data stored directly in AllegroGraph.

Any data source with a supported JDBC driver can be integrated into an AllegroGraph Knowledge Graph, including Databases (i.e. Apache Cassandra, AWS Athena, Microsoft SQL Server, MongoDB, MySQL, Oracle Database); BI Tools (i.e. IBM Cognos, Microsoft PowerBI, RapidMiner, Tableau); CRM Systems (i.e. Dynamics CRM, Netsuite, Salesforce, SugarCRM); Cloud Services (i.e. Active Directory, AWS Management, Facebook, Marketo, Microsoft Teams, SAP, ServiceNow) and Shared Data Files (i.e. Box, Gmail, Google Drive, Office365).

Streaming Graph Pipelines using Kafka

Enterprises that need real-time experiences are starting to adopt streaming pipelines to provide insights that adapt to new data in real-time rather than processing data in batches. AllegroGraph is often used as an Entity Event Knowledge Graph platform in diverse settings such as call centers, hospitals, insurance companies, aviation organizations and financial firms.

AllegroGraph 7.2 can be used seamlessly with Apache Kafka, an open-source distributed event streaming platform for high-performance data pipelines, streaming analytics, data integration and mission-critical applications. By coupling AllegroGraph with Apache Kafka, users can create a real-time decision engine that produces real-time event streams based on computations that trigger specific actions. AllegroGraph accepts incoming events, executes instant queries and analytics on the new data and then stores events and results.

Graph Analytics with Apache Spark

AllegroGraph 7.2 enables users to export data out of the Knowledge Graph and then perform graph analytics with Apache Spark, one of the most popular platforms for large-scale data processing. Users immediately gain machine learning and SQL database solutions as well as GraphX and GraphFrames, two frameworks for running graph compute operations on data.

A key benefit of using Apache Spark for graph analytics within AllegroGraph is that it is built on top of Hadoop MapReduce and extends the MapReduce model to efficiently use more types of computations. Users can access interfaces (including interactive shells) for programming entire clusters with implicit data parallelism and fault-tolerance.

Availability of AllegroGraph 7.2

AllegroGraph 7.2 is immediately available directly from Franz Inc. For more information, visit the AllegroGraph Quick Start page for cloud and download options.

Examples

Visit our [Github AllegroGraph Examples](#) page.

Graph Neural Networks Combined with Semantic Reasoning Deliver 'Total AI'

The ability for machines to reason not just identify patterns in massive data amounts, but make rule or logic based inferences on domain specific knowledge is foundational to Artificial Intelligence. The growing momentum around Neuro-Symbolic AI and the increasing reliance on Graph Analytics demonstrate how important these developments are for the enterprise.

Combining AI's symbolic knowledge processing with its statistical branch (typified by machine learning) produces the best prescriptive outcomes, delivers total AI, and is swiftly becoming necessary to tackle enterprise scale applications of mission-critical processes like foretelling equipment failure, optimizing healthcare treatment, and maximizing customer relationships.

Graph Neural Networks (GNN) exemplify the confluence of machine learning and AI reasoning. Their underlying graph capabilities are ideal for applying machine learning's advanced pattern recognition to high-dimensional, non-Euclidian datasets that are too complex for other machine learning types.

Organizations get two forms of reasoning in one framework by fusing GNN reasoning capabilities around relationship predictions, entity classifications, and graph clustering, with classic semantic inferencing available in Knowledge Graphs. Automatically mixing and matching these two types of reasoning is next level AI and is the basis for predicting any business event based on context at scale.

Creating Knowledge from Knowledge

The reasoning furnished by semantic inferencing and Graph Neural Networks is complementary because it has similar functionality for distinct applications. Although semantic knowledge graphs support both, they offer native support for semantic inferencing. This formidable capability effectively allows organizations to derive new knowledge from existing facts about their data.

Read the full article at [Data Science Central](#).

The Next Step in Machine Learning's Evolution: Graph Neural Networks



Dr. Aasman recently publish this article on Graph Neural Networks at Toward Data Science.

The capacity to consistently attain enterprise value from mission critical machine learning deployments hinges on at least one of the following three applications: classifying

entities, predicting events, and understanding why events happened.

No matter which technique is used, whether it includes supervised, unsupervised, or reinforcement learning, or if the scale and compute of deep learning is involved, conventional machine learning has limitations for solving these business problems.

It works well for many types of data, but incurs difficulty and outright failure when applied to high dimensionality networked datasets. These limits demand a new approach for social network research, recommendation engines, biology, chemistry, computer vision, and Natural Language Processing deployments in which context is pivotal.

Graph Neural Networks excel at predicting events, explaining them, and classifying entities at scale to deliver striking accuracy for these and other pragmatic deployments. Pairing their reasoning with semantic inferences creates additional knowledge for predicting events based on the multifaceted, contextualized relationships in high dimensionality data.

Read the full Article.

The Amazing Applications of Graph Neural Networks

Dr. Jans Aasman, CEO, Franz Inc. was interviewed for this InsideBigData Article:

The predictive prowess of machine learning is widely hailed as the summit of statistical Artificial Intelligence. Vaunted for

its ability to enhance everything from customer service to operations, its numerous neural networks, multiple models, and deep learning deployments are considered an enterprise surety for profiting from data.

But according to Franz CEO Jans Aasman, there's just one tiny problem with this lofty esteem that's otherwise accurate: for the most part, it "only works for what they call Euclidian datasets where you can just look at the situation, extract a number of salient points from that, turn it into a number in a vector, and then you have supervised learning and unsupervised learning and all of that."

Granted, a generous portion of enterprise data is Euclidian and readily vectorized. However, there's a wealth of non-Euclidian, multidimensionality data serving as the catalyst for astounding machine learning use cases, such as:

Network Forecasting: Analysis of all the varying relationships between entities or events in complex social networks of friends and enemies yields staggeringly accurate predictions about how any event (such as a specific customer buying a certain product) will influence network participants. This intelligence can revamp everything from marketing and sales approaches to regulatory mandates (Know Your Customer, Anti-Money Laundering, etc.), healthcare treatment, law

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

Data Summit Connect – Graph

Neural Networks (Recording)



The Premier online data management and analytics conference. Dive into the latest strategies and technologies in data management and analytics at Data Summit Connect 2021. From modern data architecture and cloud migration, to DataOps and analytics, Data Summit 2021 provides technical know-how, practical advice and expert insights to succeed in this evolving space. Topics cover innovative approaches that the world's leading companies are taking to solve today's key challenges and emerging technologies revolutionizing how data is stored, protected, integrated, enhanced and acted upon.

Dr. Jans Aasman will be presenting – Graph Neural Networks – Transforming AI Knowledge Graphs