



Top Trends in Data Science and Engineering 2022: Graph Databases + Graph Data Science

Jay Yu, Ph.D. - VP Product & Innovation

November 11, 2021



About Me



Jay Yu Ph.D.

VP of Product & Innovation

TigerGraph

Leading TigerGraph's product, solution and innovation teams.

Hands-on full-stack innovator, strategic thinker, leader and evangelist for new technology and product, with 25+ years of industry experience, including big tech company like Teradata and Intuit.

Prior to TigerGraph, served as a Distinguished Architect / Director at Intuit, with focus on enterprise data architecture and strategy, led 4 teams successfully adopted TigerGraph.

Hold 29 patents and received a PhD from University of Wisconsin - Madison's parallel database engine research group.

About TigerGraph



Our product: **the most scalable and performant graph analytics platform**

- MPP native graph database with speed **40-300x faster** than competition
- First and the only vendor completed **36TB** LDBC (SNB-SF30K BI) benchmark
- Bring AI/ML to your data via **In-database Machine Learning**
- Supports HTAP (**OLTP + OLAP**) for real-time and batch workloads
- SQL-like query language (**GSQL**) and delightful dev experience (**GraphStudio**)
- Available **on-premise** & on leading cloud platforms (**AWS, GCP, Azure**)



Our customers: enterprises at all sizes, including **Fortune 500 companies**, in

- financial services, healthcare, telecom, media, utilities and innovative startups in cybersecurity, ecommerce and retail ...



Our company: **innovative and disruptive** database technology company


- Founded in 2012, HQ in Redwood City, California. **Total Funding: \$170M+**.

2022 Top Trend in Data Engineering: Graph Databases



Rapid Growth of Graph Databases

Complete trend, starting with January 2013



By 2025, graph technologies will be used in **80%** of data and analytics innovations, up from 10% in 2021, facilitating rapid decision making across the enterprise.


Gartner[®]

Source: Gartner, "Top Trends in Data and Analytics for 2021", Rita Sallam et al, 2021



Why Graph is becoming more and more Important ?

Volume of data created and replicated worldwide (source: IDC)



Graph naturally connects data together in a simple and efficient way that other big data technology cannot do!

Data Silo

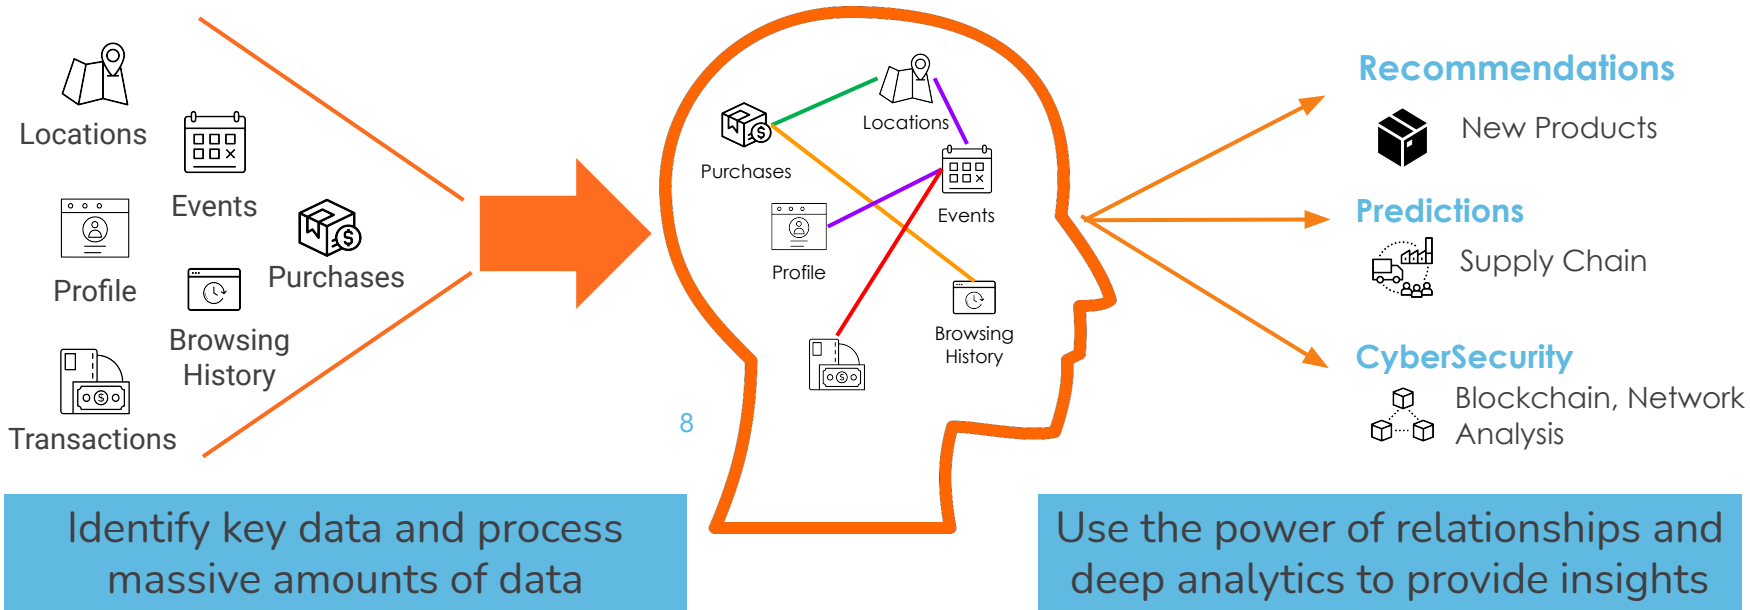


Data Swamp




Graph Is How WE THINK

Graph is a natural, adaptable model for transforming data into knowledge and insights.




How is Graph Model different than other Data Models ?


Relational




Analytical (OLAP)




Key-Value




Column-Family



Graph



Document



Graph Model can represent all other Data Models **Naturally** !

Source: Dan McCreary <https://www.slideshare.net/Dataversity/nosql-now-nosql-architecture-patterns-23589170>

Adoption by Fortune 500 Enterprises

CONNECT ALL DATASETS AND PIPELINES

Friction-free scale up from GB to TB to Petabyte with lowest cost of ownership

UNITEDHEALTH GROUP®

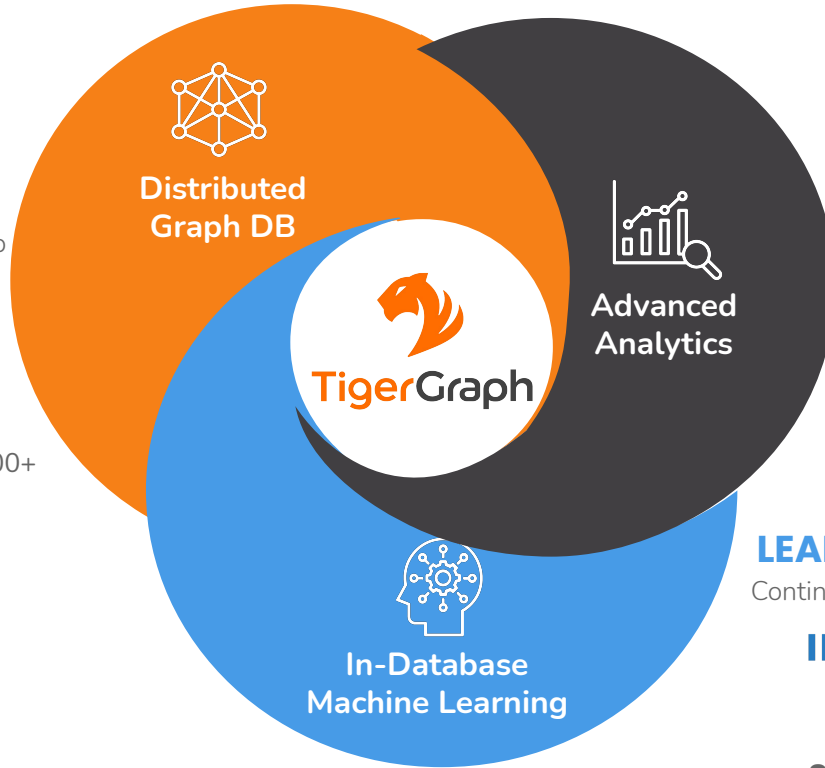
Customer 360 connecting 200+ datasets and pipelines

Fortune 50 Retailer

Item 360 for eCommerce across 100+ datasets



Identity graph connecting multiple data pipelines



ANALYZE CONNECTED DATA

10-100X faster than current solutions

Jaguar Land Rover

Supply chain planning accelerated from 3 weeks to 45 minutes



Fraud Detection - batch to real-time for 750 million calls/day

LEARN FROM CONNECTED DATA

Continuous graph-based feature generation and training

intuit.

AI-based Customer 360 for entity resolution, recommendation engine, fraud detection

8 out of top 10 global banks

Real-time fraud detection and credit risk assessment

Industry Standard Benchmark Reaches New Scale !



INTRODUCTION

Home / Introduction

The Linked Data Benchmark Council (LDBC) is processing technologies. LDBC consists of me

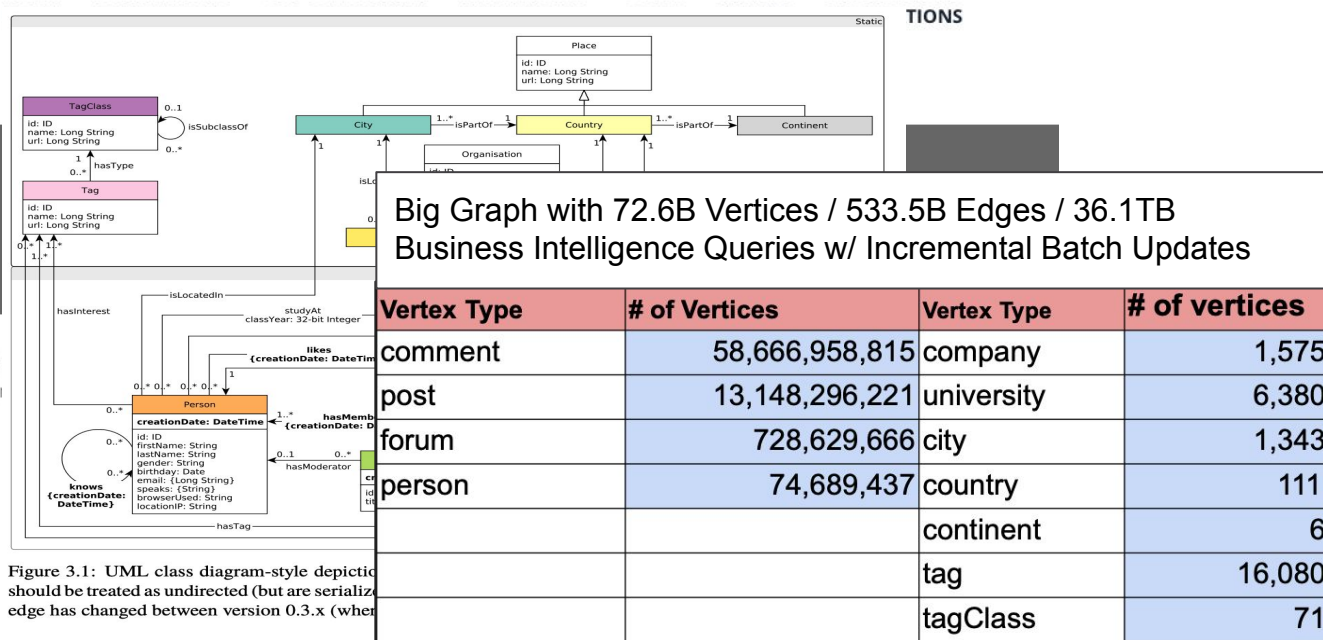


Figure 3.1: UML class diagram-style depiction of the LDBC SNB-SF30K BI graph. The diagram shows classes such as TagClass, Tag, Person, City, Country, Organisation, and Place. Relationships include isSubclassOf, hasType, isLocatedIn, likes, hasMember, hasModerator, hasTag, and knows. Multiplicities and directions are indicated on the association lines.

TigerGraph: **the first and only vendor** completed LDBC SNB-SF30K BI at 36TB!

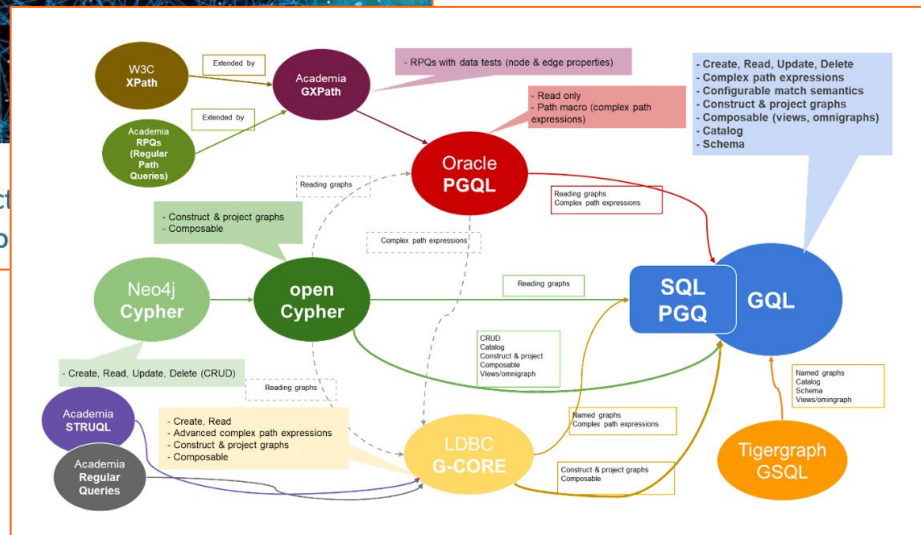


Graph Query Languages are Converging: GQL !



NEWS: GQL was just inaugurated as an official ISO project for Graph Databases to Become International Standard for

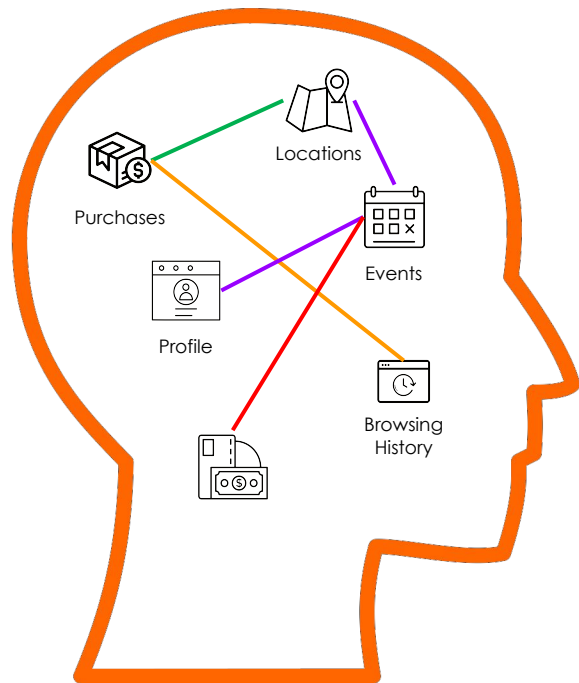
Source: <https://www.gqlstandards.org/home>



2022 Top Trend in Data Science: Graph Data Science



Rise of Graph Data Science

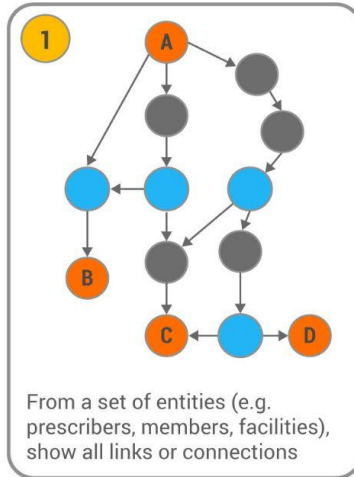


- Natural Merging of Data + Learning: Relationships are Fundamental
- Meaningful Patterns
- Enhances Conventional ML
- Explainable Results

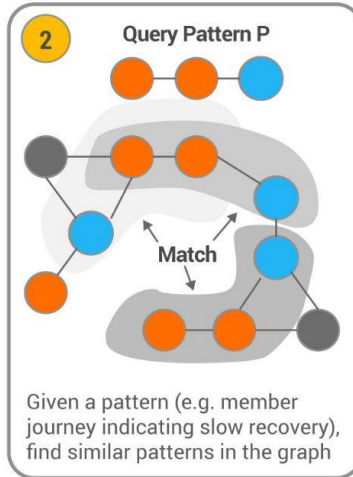
Graph Data Science brings Expert Knowledge to ML via Graph Analytics !

Graph Analytics: Unsupervised Learning

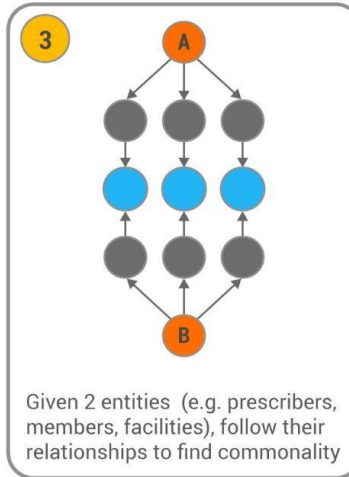
Deep Link Analysis



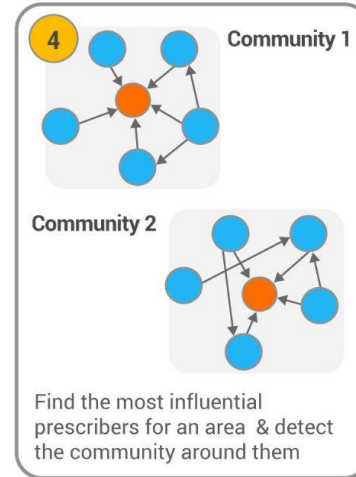
Multi-dimensional Entity & Pattern Matching



Relational Commonality Discovery & Computation



Hub & Community Detection



5 Geospatial Graph Analysis

Analyze changes in entities (prescribers, members) & their relationships with location data

6 Temporal (Time-Series) Graph Analysis

Analyze changes in (prescribers, members) & their relationships over time

7 Machine Learning Feature Generation & Explainable AI

Extract graph-based features to feed as training data for machine learning; Power Explainable AI

Graph Data Science by TigerGraph

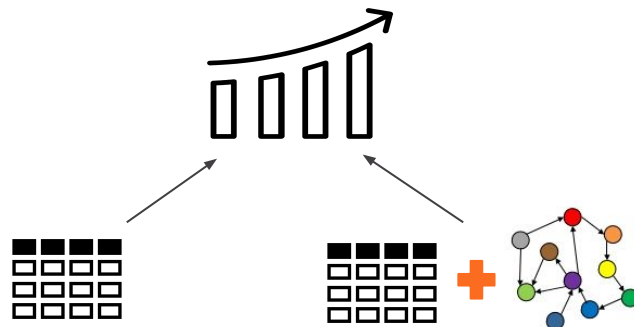
In-Database Graph Algorithms

- 50+ algorithms for similarity, clustering, best paths, community
- No export needed. Live, updatable data
- Open-source, customizable
- Scalable, Ultra-fast MPP graph engine
- Toolkit, building blocks for problem-solving



Graph Features Enriching ML Quality

- Graph enriches ML training data: quality in \Rightarrow quality out
- Use graph algorithms, pattern-matching queries, or graph embedding



Solving Enterprise-Scale Business Problems

Sample Graph Data Science Use Cases in Production

Healthcare:


Real-time recommendations



- 1.3TB graph brain
- Real-time care recommendations
- Improving healthcare, lowering cost

Industrial Supply Chain:


Analytics for decisions



- Analytics: weeks → minutes
- Reveal opportunities, optimize tactical & strategic decisions
- Saving \$100M+/yr


Financial Services:

Real-time fraud detection




- Integrates multiple tools
- "Magical" real-time visual results for investigators
- Scalable for growth

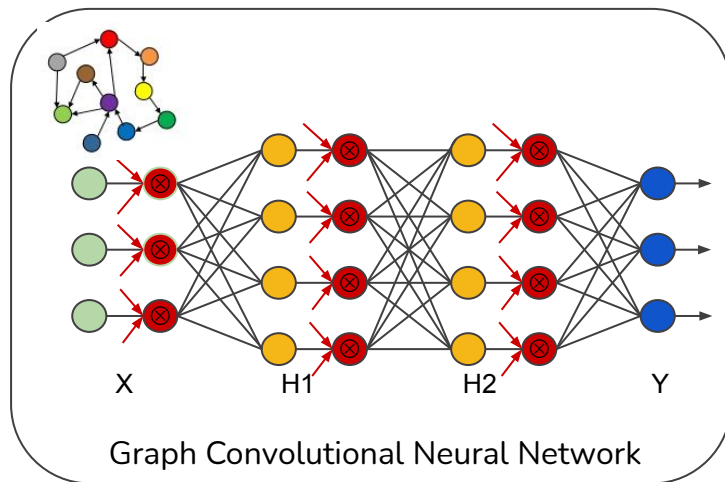

Advanced R&D: Graph Neural Networks



Powerful Machine Learning to
Predict and Classify



Insight from Connected Data



Graph Convolutional Neural Network

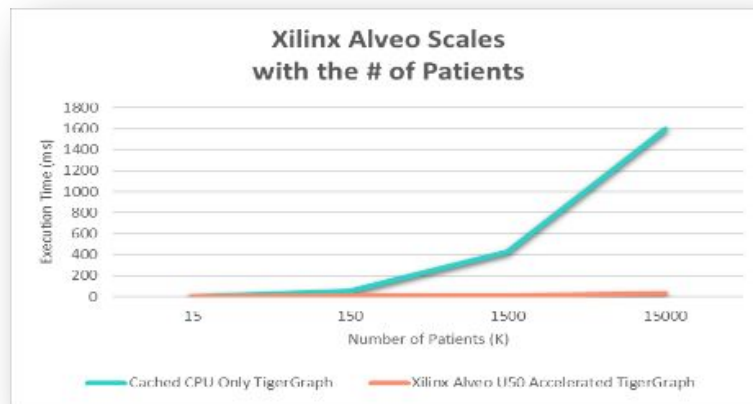
- Combines the added insight from connected data with the modeling power of neural networks
- Uses the graph during training; in-database training is the ideal.

HW Acceleration for Graph Data Science

HOME > STORE > The Accelerated Path To Petabyte-Scale Graph Databases

THE ACCELERATED PATH TO PETABYTE-SCALE GRAPH DATABASES

October 28, 2021 Timothy Prickett Morgan



Special HW Integration accelerating compute intensive graph algorithms by **20X to 48X !**



Summary



Key Takeaways: in 2022, we will see ...

The Year of Graph Databases !

- Scale: 10TB → 100TB
- Standardization: Siloed → GQL (Query) + LDBC (Benchmark)
- Deployment: On-Prem → Cloud (Private + Public)

Rapid Adoption of Graph Data Science !

- Setting: Exploration → Production
- Approach: ML Only → In-Database Machine Learning (Graph + ML)
- Infrastructure: Commodity HW → + Specialized HW Acceleration



Try TigerGraph Cloud for free:

- tigergraph.com/cloud

Learn more from Graph + AI Summit

- tigergraph.com/graphaisummit

Thank You

