TigerGraph

TigerGraph Increases Revenue with Hyper-Personalized Product Recommendations

RECOMMENDATIONS INCREASE REVENUE

Businesses with a wide range of products or services have an opportunity to boost revenue by recommending complementary options to consumers. The impact of these recommendations can be huge: Amazon estimates that 35% of its revenue comes from cross-selling and up-selling. Purchases with recommendation clicks result in a 10% higher average order value and the per visit spend of a shopper who clicks a recommendation is five times higher.

An ability to offer increasingly personalized recommendations can be a significant competitive differentiator. According to one study, 82% of consumers report being influenced by a personalized shopping experience. Moreover, relevant product recommendations help retailers build a deep relationship with their customers as they give them a sense of being understood and properly served. Now, more than ever, businesses are looking to capture the business moments into revenue and win market share with the personalized offers.

LEGACY SYSTEMS CANNOT KEEP UP

Traditional recommendation engines built on the relational databases aren't able to keep up with these requirements. Legacy recommendation system performs global statistical computations offline, using snapshots of data that can be days old. It doesn't have the real-time modeling and nuanced profiling that is needed today.

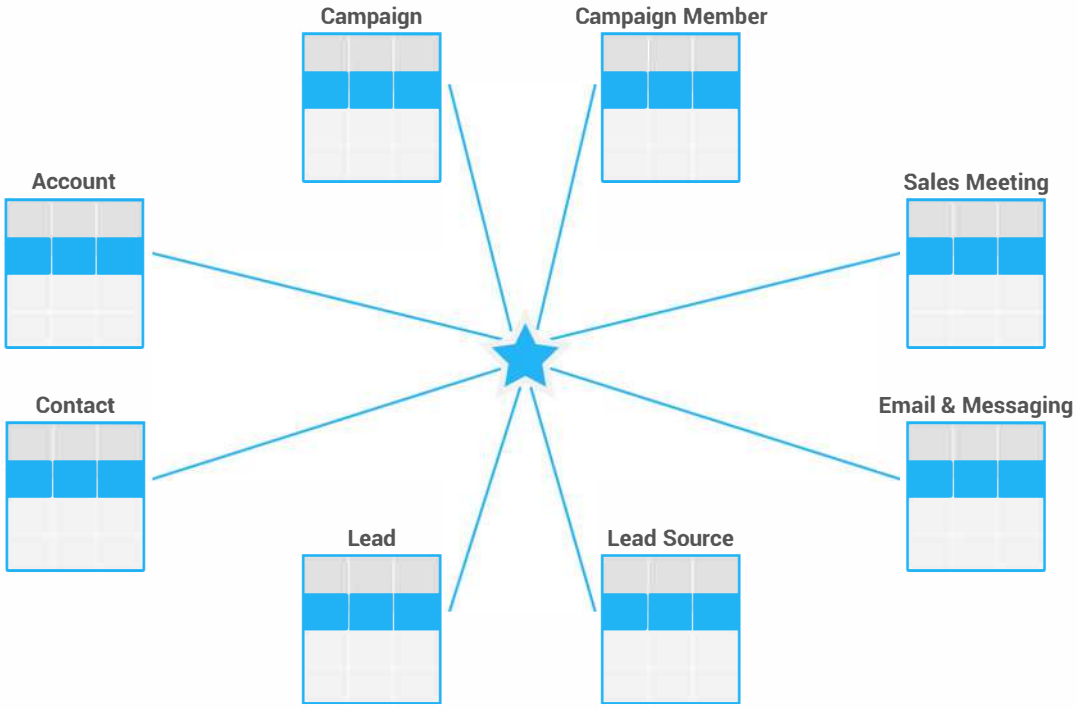

Figure 1: Businesses relying on a relational database are limited in their ability to provide sub-second product and service recommendations to potential buyers

An ability to offer complementary product or service recommendations instantly is essential in many scenarios. In eCommerce environments, for example, recommendations need to be presented to consumers within 500ms. Recommendation systems need to quickly understand the profile of their client, align that with the rapidly changing profiles of the larger customer base and product catalog, and produce engaging, personalized recommendations.

TigerGraph

TigerGraph Increases Revenue with Hyper-Personalized Product Recommendations

Using a graph for recommendation analytics is the first step towards faster product and service recommendations. Native parallel graphs, such as TigerGraph, are built to understand, explore and analyze the complex relationships in the eCommerce data, allowing data scientists and business users, to go 10 or more levels deep into the data, across billions of orders and millions of customers, in real-time.

Consider that the standard collaborative filtering algorithm is a three hop or connection graph traversal such as People ➪ who bought the product ➪which you just bought ➪also bought these other products. While other approaches struggle with more than two hops or connections, TigerGraph handles 10 or more hops. TigerGraph's deep link analytics enables businesses to customize and extend their analytics, adding hops to consider product features, customer demographics, and the context of the current situation, resulting in more personalized recommendations.

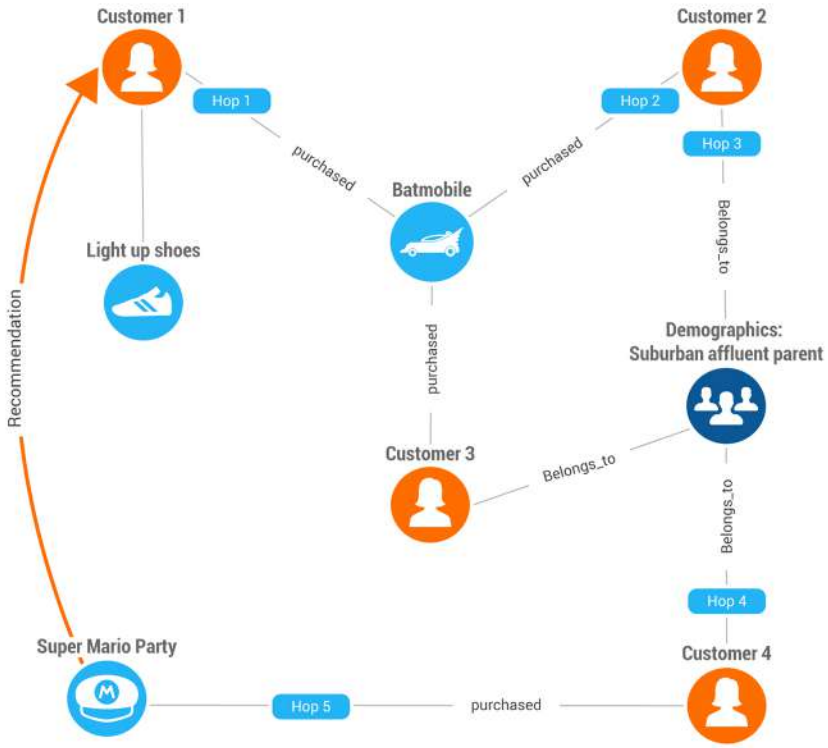

Figure 2: Businesses can use deep link analytics to identify complementary products and services

It takes two hops to find similar shoppers: Shopper ➪ (Demographics) ➪ Similar Shoppers and demographic-aware collaborative filtering can be implemented in five hops: Shopper ➪Products purchased ➪ Other Shoppers who purchased the product ➪ Demographics ➪ Other Shoppers that belong to the Demographics ➪ Other Products Purchased.

Consider the example shown in figure 2 where Customer 1 has purchased a toy batmobile and light up shoes in prior visits. Customer 2 and Customer 3 have both purchased the batmobile toy, and belong to the demographics of “Suburban affluent parent”. Customer 4 belongs to that demographic as well and has purchased the video game “Super Mario Party for Nintendo Switch”. Based on the five-hop analysis, the video game, “Super Mario Party for Nintendo Switch” is the recommended product for Customer 1.

TigerGraph

TigerGraph Increases Revenue with Hyper-Personalized Product Recommendations

CONSUMERS EXPECT PERSONALIZATION

As the digital era has given consumers unprecedented choices, all just a click away, simple recommendations such as "customers who bought this item also bought" or "these products are often bought together" are no longer enough. To foster brand loyalty, businesses must gain a more sophisticated understanding of the unique, varied, and complex characteristics of every customer and deliver, on the fly, offers and recommendations that truly speak to them.

A standard collaborative filtering algorithm, used by popular sites such as Amazon.com, creates a recommendation based on the buying behavior of other users delving only two or three connections or hops into the data. The first hop is from the customer to the product they are reviewing or have purchased. The second hop is from the product to other customers who have bought that product. The third hop is from these customers to other products they have bought over time that aren't yet purchased by the customer. Simple next best offer recommendations—such as "Customers who viewed/bought this product, also bought (these other products)" or "Frequently bought together"—are losing out to more personalized recommendations built on advanced technologies capable of going beyond three hops.

KICKDYNAMIC GETS PERSONAL WITH TIGERGRAPH

Kickdynamic is an example of a company that is outperforming its competition by offering hyper-personalized product recommendations. Kickdynamic, a marketing platform used by businesses to personalize email content is using TigerGraph to help its more than 200 global customers in fashion, retail, travel, and other sectors achieve maximum email personalization and relevancy.

By using TigerGraph's native parallel graph database, Kickdynamic can instantly understand the recipients' profiles and deliver relevant recommendations showing products that are more likely to capture their interest, exploiting past onsite browsing behavior, similarities across products and users, and cross-user behavioral patterns to discover and extract similarities between users, products, or both.

Kickdynamic is able to go 10 or more hops deeper and provide real-time deep link analytics that opens up vastly more information about a customer's likes, needs, desires, and intent. This supports their advanced personalization capabilities that can display products that customers have browsed on-site and live in email, automatically at open time, and provides them with a significant competitive edge.

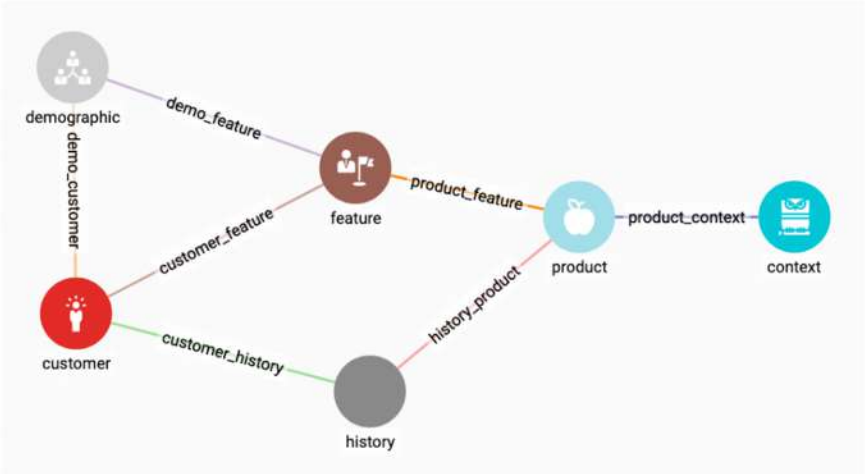

Figure 3 - The recommendation graph updates in real-time in response to changes in customer and product data.

TigerGraph CLOUD

The Real-Time Native Parallel Graph

CUSTOMERS

Real-time fraud detection at 4 out of the world's top 5 banks

Care path recommendations for 50 million patients

Personalized offers for 300 million consumers

Energy infrastructure optimization for 1 billion people

Learn more at tigergraph.com/customers

CUSTOMER QUOTES

"It's huge data (*terabytes*) and finding influencers in that data, it's not easy, but TigerGraph has scaled for us."

- Vishnu Maddileti
Director of Data Sciences and Analytics
Amgen

"Some of the questions that graph databases answer are hard to come to conclusion with in RDBMS or it takes forever. We needed a better tool to find relationships and TigerGraph was just that."

- Ely Turkenitz, IS Manager
Santa Clara County

GET STARTED FOR FREE AT TIGERGRAPH.COM/CLOUD

TigerGraph Cloud graph database as a service is built for agile teams who'd rather be building innovative applications to deliver new insights than managing databases.

CLOUD STARTER KITS

TigerGraph Cloud Starter Kits are built with sample graph data schema, dataset, and queries focused on specific use cases such as fraud detection, recommendation engine, supply chain analysis and/or a specific industry such as healthcare, pharmaceutical or financial services.

Starter Kit	Overview
Customer 360 – Attribution and Engagement Graph	Real-time 360-degree view of the customer journey for attribution and engagement
Cybersecurity Threat Detection-IT	Block cybersecurity threats by detecting interconnected events, devices and people
Enterprise Knowledge Graph (Corporate Data)	Analysis of corporate data including investors and key stakeholders
Enterprise Knowledge Graph (Crunchbase)	Knowledge Graph example featuring Crunchbase data
Entity Resolution (MDM)	Identify, link and merge entities with analysis of attributes and relationships
Financial Services (Payments) – Fraud Detection	Detect and stop fraudulent payments in real-time
Fraud and Money Laundering Detection (Financial Services)	Multiple types of fraud and money laundering patterns
GSQL 101	Introduction to TigerGraph's powerful graph query language
Healthcare Graph (Drug Interaction /FAERS)	Focused on public (FAERS) and private data for pharmaceutical drugs
Healthcare – Referral Networks, Hub (PageRank) & Community Detection	Analyze patient claims to establish referral networks, identify influential doctors
Machine Learning and Real-time Fraud Detection	Mobile industry example for detecting fraud in real-time and for machine learning
Network and IT Resource Optimization	Network and IT resource graph for analyzing the impact of hardware outages
Recommendation Engine (Movie Recommendation)	Graph-based movie recommendation engine built with public data
Social Network Analysis	Social network example for understanding and analyzing relationships
Supply Chain Analysis	Example covering inventory planning and impact analysis

About TigerGraph

TigerGraph is the only scalable graph database for the enterprise. TigerGraph's proven technology connects data silos for deeper, wider and operational analytics at scale. Four out of the top five global banks use TigerGraph for real-time fraud detection. Over 50 million patients receive care path recommendations to assist them on their wellness journey. 300 million consumers receive personalized offers with recommendation engines powered by TigerGraph. Energy infrastructure for 1 billion people is optimized by TigerGraph for reducing power outages. TigerGraph's proven technology supports applications such as fraud detection, customer 360, MDM, IoT, AI and machine learning. The company is headquartered in Redwood City, California, USA. Follow TigerGraph on Twitter at @TigerGraphDB or start free at tigergraph.com/cloud.

CONTACT

TigerGraph
3 Twin Dolphin Drive, Suite 225
Redwood City, California 94065
United States

www.tigergraph.com