

WHITE PAPER

Graph-Powered Recommendation Engines

How the Power of Suggestion Drives
Better Decisions & Higher Revenues

Nav Mathur, Senior Director, Global Solutions

TABLE OF CONTENTS

The Power of Suggestion	1
Recommendations Are Everywhere	1
Recommendations Add High Value	2
Graph Technology Is Superior for Recommendations	3
The Demands of Modern Recommendations Applications	4
Neo4j Recommendations Framework	5
Neo4j Graph Database	6
Flexible, Powerful Recommendation Engines	7
Developing & Deploying Recommendations Applications	9
Why Use Neo4j for Recommendations Applications	10

Graph-Powered Recommendation Engines

How the Power of Suggestion Drives Better Decisions & Higher Revenues

Nav Mathur, Senior Director, Global Solutions, Neo4j

Recommendation engines have become a crucial component of modern sales, finance, logistics and human resources applications. This core need has triggered a shift from relational and big-data approaches to graph-based technologies that are purpose-built to handle the rigorous demands of real-time recommendations.

The Power of Suggestion

In an increasingly competitive world, enterprises across the globe must differentiate their brands and streamline their business operations. But evaluating the overwhelming amount of information impacting each and every decision is an expensive, time-consuming and error-prone endeavor.

As a result, leading enterprises are turning to [artificial intelligence \(AI\) and machine learning \(ML\)](#) to make better decisions that improve enterprise efficiency and put real distance on competitors. At the core of this revolution is a shift from simple self-service to a more sophisticated approach – from the convenience of selection to the power of suggestion.

Recommendations Are Everywhere

As enterprises modernize their internal and customer-facing systems, they're building [smart, instant recommendations](#) that:

- Personalize user experiences to improve customer loyalty
- Help staff to make effective and impactful decisions
- Drive efficiency across enterprise workflows
- Increase revenues and eliminate costs

Recommendations give users a personalized experience that:

- Supports the context and objectives of their current workflow
- Takes non-linear paths through an application to suggest products or choices
- Enforces enterprise strategies for each persona and context

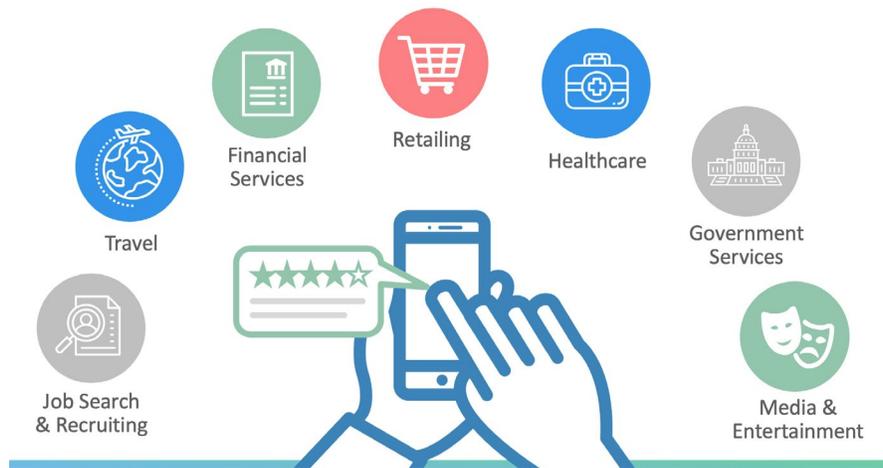
Graph-Powered Recommendations Engines

Product Recommendations

With online retail sites in ten countries, [Walmart](#) calls Neo4j “a perfect tool for real-time product recommendations.” Software developer Marcos Wada notes, “Neo4j helps us understand our online shoppers’ behavior and the relationship between our customers and products, providing a perfect tool for real-time product recommendations.”

Personalized and situational recommendations are used today in a wide spectrum of consumer and business applications including ecommerce, social networks, travel, entertainment, financial services, logistics, manufacturing, marketing, customer success and human resources.

Recommendations Are Everywhere



After ecommerce, recommendation engines have spread across the entire economy

What has ignited this explosion in the use of recommendations engines? Two factors:

- The sheer volume of data being collected from user visits and internal operations
- Modern technologies – specifically graph, AI and ML advancements – that discover and exploit otherwise-hidden data relationships

Recommendations Add High Value

Everyone is familiar with [ecommerce recommendations](#) that answer questions such as:

- Who else is similar to the current user and what have they purchased?
- Which products or services are most like the one the customer is currently viewing?
- Which overstock and promotional products are most appropriate to suggest to the current user?
- Which songs, events or other media do similar users find most appealing?

However, recommendation engines also answer business questions such as:

- Which employees have skills and experience that match open job opportunities?
- Which financial services are most attractive to this type of customer?
- Which prices need to be adjusted to attract more customer interest, and by how much?
- Which factory-supplier combination can produce products cheaply and ship them to Malaysia in four weeks?
- Which freight couriers are most cost-effective and dependable at delivering to Arizona?

Movie Recommendations

A movie website recommends current and upcoming releases based on fans' individual tastes. In turn, it gives studios insight into the preferences of movie viewers so they effectively target promotions. After considering MySQL for its recommendation system, they realized the sheer volume and connected nature of their data required Neo4j.

Music Recommendations

A large media company uses Neo4j's Recommendation Framework to provide artists with demographic and preference data for their listeners across multiple music-streaming networks. The artists then fine-tune their music and predict the outcomes of the tracks they produce.

Graph Technology Is Superior for Recommendations

The technology of choice for recommendation engines is [Neo4j](#), the world's leading graph database.

Graph technology is superior because it is specifically designed to manage not just mountains of data (of customers, products, services, prices, job openings, etc.) but also [data relationships](#) — the details of how and why people, products, services, preferences, likes, dislikes, jobs, events, and everything else in the real world are related.

Real-time recommendations must instantly assess the current business context, match user profiles and patterns, capture new behaviors for future recommendations, evaluate alternatives, and then recommend next steps. Matching business, demographic, historical and session data like this – in real time – is trivial for Neo4j.

Real-Time Recommendations

Consider user needs and your business strategies in your recommendations

User Perspective

Recommend similar and even surprising items by considering:

- Past purchases
- Similar users
- Related products



Business Perspective

Recommend items that promote your business strategies:

- High-margin items
- Overstocks
- Promotions

Relational Approaches

Legacy relational database technologies track only that two things are related, and don't track the context or direction of that relationship. It's ironic that relational databases are in fact, poor at capturing data relationships.

In sharp contrast, graph-powered recommendation engines aren't crippled by semantically poor relational data models and slow, expensive, one-way, unpredictable table JOINS. Graphs are rich in meaning, highly flexible and run at blazing speeds – fast enough to support real-time recommendations.

Data Lakes and Algorithmic Approaches

To address the problems of relational methods, data scientists have used data-lake and big-data technologies like Hadoop and HDFS. But these approaches must first interpret vast quantities of data into meaningful models, and then execute recommendation algorithms within the models. Since all this takes time, such approaches are not viable solutions for real-time recommendations in which busy users require split-second suggestions or move on to their next task.

Shopping Behavior and Recommendations

A top-tier ratings and data analysis firm uses Neo4j to analyze the shopping behavior of consumers at big-box stores compared to their purchase habits with other retailers. They are also looking for personalized product recommendations with complex scenarios of brand, size, ingredients, look-and-feel, price and more for a supermarket.

Native Graph Answers the Call

Neo4j's [native graph technology](#) is the right solution for real-time recommendations for many important reasons, including:

- Graph models of data entities and their relationships are already connected in the correct context when accessed by applications.
- Graph models allow relationships to be weighted.
- Queries traverse native-graph datasets instantaneously.
- Native-graph algorithms traverse large datasets to identify similar products and users, highly connected entities, and cohorts that help generate recommendations.
- Simple, native-graph queries perform collaborative, contextual and rule-based filtering with sub-second performance.
- Bespoke algorithmic models can also be integrated to improve the flexibility and quality of recommendations.

For all these reasons, Neo4j stands alone in its ability to produce meaningful, real-time recommendations across a full spectrum of enterprise applications.

The Demands of Modern Recommendation Engines

The need to provide smart, engaging and high-value customer experiences drive high-level business requirements for recommendation engines. In turn, the results of those recommendations should drive the strategic objectives of the enterprise while providing hyper-personalization for application users.

Explainable Results

Modern, online applications must include recommendations that are credible, contextual and reproducible – so they drive user confidence and adoption.

Real-Time Responsiveness

Online apps must make instant, real-time recommendations that maintain the momentum of online sessions and preserve user interest.

Comprehensive Visibility

Applications should leverage all available data and relationships in making in-depth, informed recommendations and predictions.

Advanced Algorithms & Techniques

Applications must provide scored and ordered recommendations generated by multiple algorithms and techniques to give users a spectrum of suggestions for moving forward.

Effective Business Results

All applications should meet business objectives, which can include higher revenues, lower costs, customer satisfaction, upsells, cross-sells, hires, reduced job flight and other goals.

Easy, Maintainable Solution

The final business requirement for most organizations is to develop and maintain a recommendations solution in a timely, cost-effective manner. This necessitates starting with a framework that is purpose-built to create and manage real-time recommendations across a variety of enterprise use cases.

Graph-Powered Recommendations Engines

Multi-Vendor Shopping

Cobrain makes personalized recommendations for products offered by more than 300 apparel merchants. After consumers tell Cobrain what they like, it uses Neo4j to find products loved by their anonymous cohorts and provide real-time recommendations. VP of Engineering and Technology Andy Rosenbaum said: "The graph database allows us to explore new connections between people, giving a very personalized experience of products we think they'll love."

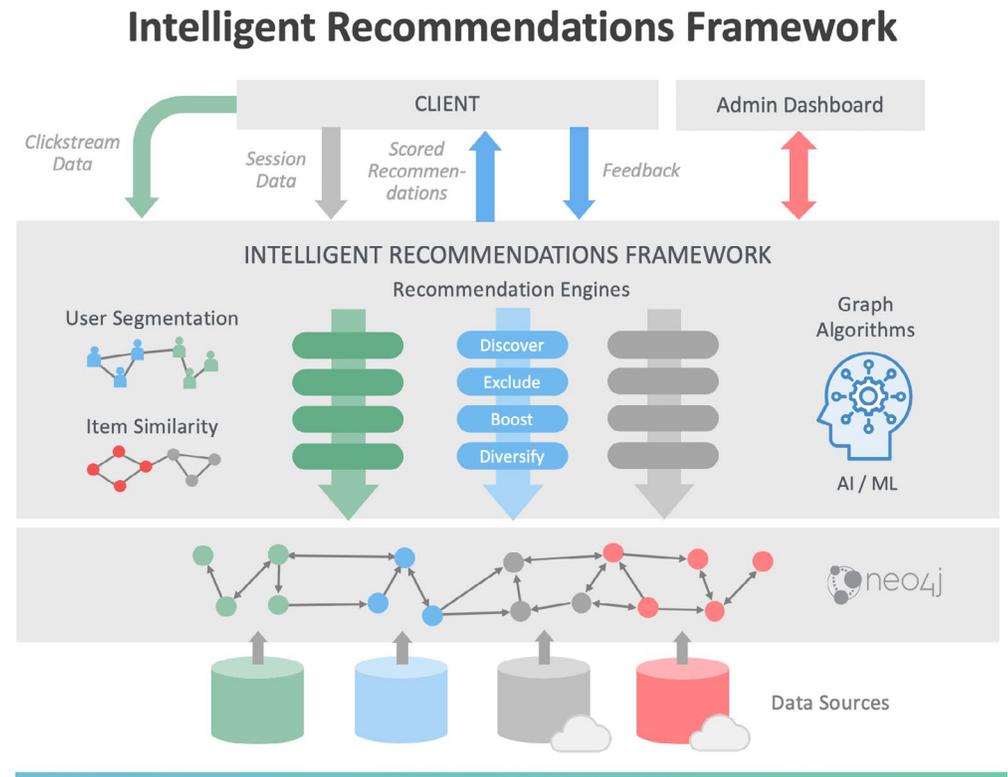
Neo4j Recommendations Framework

The [Neo4j Recommendations Framework](#) is an intelligent platform for making quality recommendations while significantly reducing the time required to develop associated applications.

The framework takes full advantage of Neo4j graph functionality and works with any data model implemented on a [Neo4j database](#). And unlike almost all recommendation technologies, the Neo4j framework is adaptable and customizable to all industries, functions and use cases – not just retail.

The framework includes the following components:

- A graph database that defines business relationships, user segments and item similarities
- Connections to enterprise databases and external sources
- Any number of hybrid-scoring recommendation engines powered by [graph and ML algorithms](#)
- A simple Node.js server with Apollo GraphQL APIs that interact with existing applications
- A dashboard for managing the recommendations environment



Neo4j provides an intelligent framework that significantly reduces the time required to develop recommendation applications

The Foundation: The Neo4j Graph Database

The foundation of the Neo4j Recommendations Framework is a [graph database](#) of the customers, products, services, vendors and suppliers that form the backbone of any enterprise.

Applications can connect to the Neo4j Recommendations Framework and traverse highly connected data from disparate data sources or use graph analytics and machine learning algorithms to produce desired recommendations. The Recommendations Framework leverages these advanced technologies to:

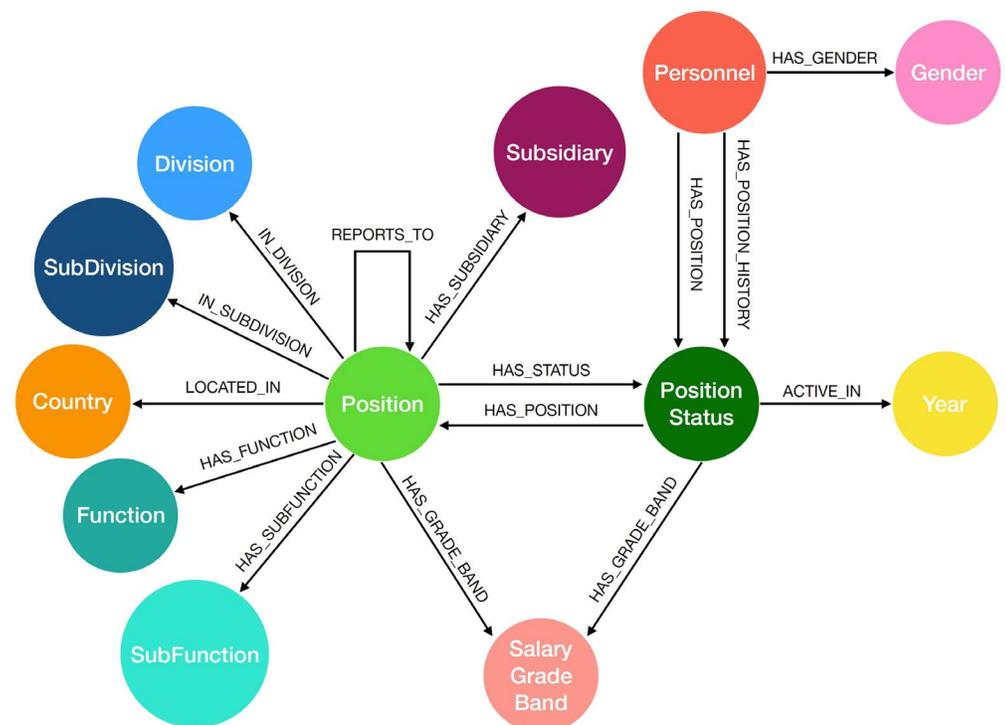
- Segment customers by demographics, preferences and other patterns
- Learn which offerings are similar or related to each other
- Identify alternate suppliers of various products and services
- Keep tabs on product pricing and promotions
- Monitor product and service availability and inventory levels
- Track order histories and preference profiles for every customer and product
- Recommend employees to fill open positions based on work histories and skill sets
- ...and much more

Human Resources Planning

A CPG company has their HR data in SAP HANA and SuccessFactors, and they used Tableau to visualize the information. They built a Neo4j dashboard that lets users build complex queries that utilize thirteen dimensions and nineteen filters. The Neo4j system includes a real-time recommendation engine that show managers who in their organizations is at high flight risk and which staff members are the best candidates to succeed them.

Spotlight: A CPG Company Uses Graph to Fill Open Positions with Internal Talent

A consumer packaged good (CPG) company tracks employee job histories, skills, reviews, locations and salaries in an extensive graph database. They use the data to analyze positions and people across their entire enterprise to fill open positions, plan corporate reorganizations, perform succession planning and identify employee flight risk. By accounting for the interrelation of all these foundational elements, Neo4j maximizes the accuracy and effectiveness of recommendations.



Corporate Teambuilding

Rather than following a traditional corporate hierarchy, Daimler Chrysler works in swarms of cross-functional teams. Daimler uses Neo4j to locate the right talent for each team, create org charts and track performance. Jochen Linkohr, the Manager of HR IT at Daimler notes that their multi-structured data is a great starting point for using graphs.

Flexible, Powerful Recommendation Engines

A Neo4j recommendations application contains one or many engines that build lists of recommendations for presenting to users.

Each recommendation engine has a database connection, a description and a pipeline of [Cypher](#) queries called phases that build recommendation lists. Each pipeline contains any of the following phases:

Discover Phase

A Discover phase always appears first in a pipeline and populates the recommendation list with items, each of which has a score. If a recommended item is re-discovered by a subsequent discovery phase, its new score is added to its existing score.

Boost Phase

A Boost phase modifies the scores of existing items in the recommendations list. The phase evaluates the score of each item in the list and modifies the score as needed.

Exclude Phase

An Exclude phase uses specific criteria to filter and remove items from the current recommendations list.

Diversify Phase

A Diversify phase widens the items in a recommendations list. It starts by returning an attribute for each item in the list. Then for each attribute, the phase keeps the top n items and removes any remaining items from the list.

These easy-to-use building blocks provide the power and flexibility required to:

- Deliver real-time recommendations
- Reduce query complexity
- Create easy-to-maintain recommendation engines quickly
- Score and weight phases depending on their relevance to the current context
- Enable (or disable) phases in the recommendations pipeline based on business rules (for promotions, inventory, etc.)
- Build powerful, efficient recommendation engines for virtually any business case



Graph-Powered Recommendations Engines

Pricing Recommendations

[Marriott Hotels](#) adopted Neo4j to power its real-time pricing recommendation engine after suffering significant slowdowns with its prior database architecture. Neo4j's graph approach has slashed response time and hardware requirements while driving 300% growth in the volume of recommended price changes.

Hybrid Scoring Improves Recommendation Quality

The Neo4j framework uses hybrid scoring, which permits the simultaneous use of multiple recommendation methods to select the one that best fits each user and situation in real time. In doing so, Neo4j enhances the quality and context of recommendations, achieving business objectives more quickly and effectively.

At the core of all recommendation engines are the graph queries and algorithms used to score people, products, services and other items as suggested connections or purchases. The Neo4j Recommendations Framework and its graph data science algorithms – including Similarity, Pathfinding, Centrality, Community Detection and Link Prediction algorithms – are used for scoring and filtering recommendations.

- **Collaborative scoring** methods are based on similar users or products.
- **Content scoring** methods are based on user histories and profiles.
- **Rules-based scoring** methods are based on predefined rules and criteria.
- **Strategy scoring** methods are based on promotions, margins, inventory and other business data.

Neo4j Graph Machine Learning Algorithms



Pathfinding & Search

Parallel Breadth *First Search and DFS*
Shortest Path
Single-Source Shortest Path
All Pairs Shortest Path
Minimum Spanning Tree
A* Shortest Path
Yen's K Shortest Path
K-Spanning Tree *MST*



Centrality & Importance

Degree Centrality
Closeness Centrality
Betweenness Centrality
PageRank
Harmonic Closeness Centrality
Dangalchev Closeness Centrality
Wasserman & Faust Closeness Centrality
Approximate Betweenness Centrality
Personalized PageRank



Community Detection

Triangle Count
Clustering Coefficients
Connected Components *Union Find*
Strongly Connected Components
Label Propagation
Louvain Modularity *1 Step*
Balanced Triad *Identification*
Louvain *Multi-Step*



Similarity & ML Workflow

Euclidean Distance
Cosine Similarity
Jaccard Similarity
Random Walk
One Hot Encoding

Material Management

A chemical company uses machine learning to predict material shortages and find other vendors that can develop the material quickly and at cheaper price points – and even predict which vendors and plants are most likely to produce specific amounts of chemicals and deliver them on time to selected markets. Neo4j also suggests alternate materials that are often several hops away from the original material.

Developing & Deploying Recommendation Applications

The Neo4j Recommendations Framework is the most productive, powerful and customizable environment for developing and deploying recommendation applications. The framework combines:

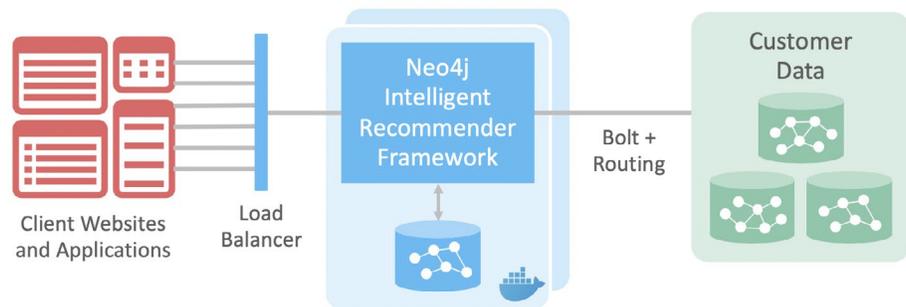
- The dependability and performance of the Neo4j graph database
- Powerful development and testing tools specifically designed for making intelligent recommendations
- Customizable deployment alternatives for application and organizational requirements

Development Advantages

The framework slashes development and testing times for recommendation applications by 50% by enabling developers to:

- Create recommendation engines using a zero-code approach
- Turn blocks of code on/off with a click to simplify testing and control phased rollouts
- Compile discrete engine phases at runtime to increase performance without losing maintainability
- Create fully functional GraphQL APIs automatically as they build recommendation engines and pipelines
- Improve maintainability by separating Cypher queries for engine phases rather than requiring gigantic Cypher statements
- Order recommendation results subjectively to promote selected products, bundles or pricing strategies

Neo4j Deployment Architecture



Deployment Advantages

Organizations have many options for deploying their Neo4j recommendation applications. Information managers maximize dependability, availability and performance by:

- Taking full advantage of Neo4j's High Availability, load balancing and server clustering architecture
- Configuring independent clusters for customer data and the Recommendations Framework to provide extreme scalability, balancing and fault tolerance
- Deploying recommendation apps with or without internet access, on local networks, or even on single, disconnected machines
- Scheduling recommendation engines to run at specific, predetermined times

Graph-Powered Recommendation Engines

Personal Finance Recommendations

Wobi provides best-value insurance offers to more than half a million users. At its core is a Neo4j recommendation engine that enables them to drill down into vendor offerings and customer profiles – and to add new information on the fly. With a database of more than four million nodes and 30 million relationships, Shai Bentin, Wobi's Chief Technology Officer says, "I feel safe with Neo4j."

Why Use Neo4j for Recommendations?

When you select Neo4j as the foundation for your recommendation systems, you transform the way your customers and employees interact with your online systems.

Real-Time Recommendations

When making recommendations, time is of the essence. The speed and flexibility of Neo4j's graph technology enables you to offer real-time recommendations for the first time. By reducing recommendation calculations from minutes to milliseconds, [eBay](#) cites that "Neo4j allows us to add functionality that was previously not possible."

Works with All Data Sources

As an open framework, Neo4j links to networked data sources and supports graph-based relationships that transcend all those sources. As a result, it maintains top performance for even your most complex recommendation algorithms and business environments.

Deployed Easily to All Applications

Neo4j recommendation engines can be independently deployed and scaled to meet the performance and load-balancing needs of existing applications and portals. Neo4j stores framework data logically separate from customer data to maintain the integrity, privacy and security requirements of your internal and customer information.

Widest Choice of Scoring Methods

The Neo4j framework combines graph-based queries and algorithms for scoring recommendations and enables you to create weighted scores based on multiple techniques in real time – resulting in more accurate, context-aware recommendations.

Flexible, Highly Configurable Engines

Neo4j gives you a fast and easy way to build modular, configurable pipelines that tune recommendations to the specific requirements of every use case and context. Their modularity gives them an unprecedented level of configurability and makes them trivial to modify and maintain.

Higher Quality Recommendations

Effective recommendations have obvious value to your users and are meaningful in their current context. Using Neo4j's flexible graph approach, your recommendations are valuable, explainable and relevant to your users.

Fast to Develop, Maintain and Expand

The application development framework in Neo4j enables you to focus on just the business logic of recommendations rather than writing infrastructure code. The result is applications with up to ten times less code – sometimes even code-free – decreasing your time to market with a recommendation engine that is easy to develop, support and maintain.

Achieve Business Objectives

The bottom line for all recommendation engines is how effectively they achieve your business objectives. The power, flexibility and high performance of Neo4j ensure your recommendations are on time and on target, so they meet the revenue and business goals of your organization.

Learn More about the Neo4j Recommendations Framework

To learn more about using Neo4j to add the power of suggestion to your business, contact us today.

Neo4j is the leader in graph database technology. As the world's most widely deployed graph database, we help global brands – including [Comcast](#), [NASA](#), [JBS](#), and [Volvo Cars](#) – to reveal and predict how people, processes and systems are interrelated.

Using this relationships-first approach, applications built with Neo4j tackle connected data challenges such as [analytics and artificial intelligence](#), [fraud detection](#), [real-time recommendations](#), and [knowledge graphs](#). Find out more at [neo4j.com](#).

Questions about Neo4j?

Contact us around the globe:
info@neo4j.com
neo4j.com/contact-us