


SHAKUDO

# Knowledge Graphs

How Knowledge Graphs Power Explainable  
AI Decisions



---

# Knowledge Graphs

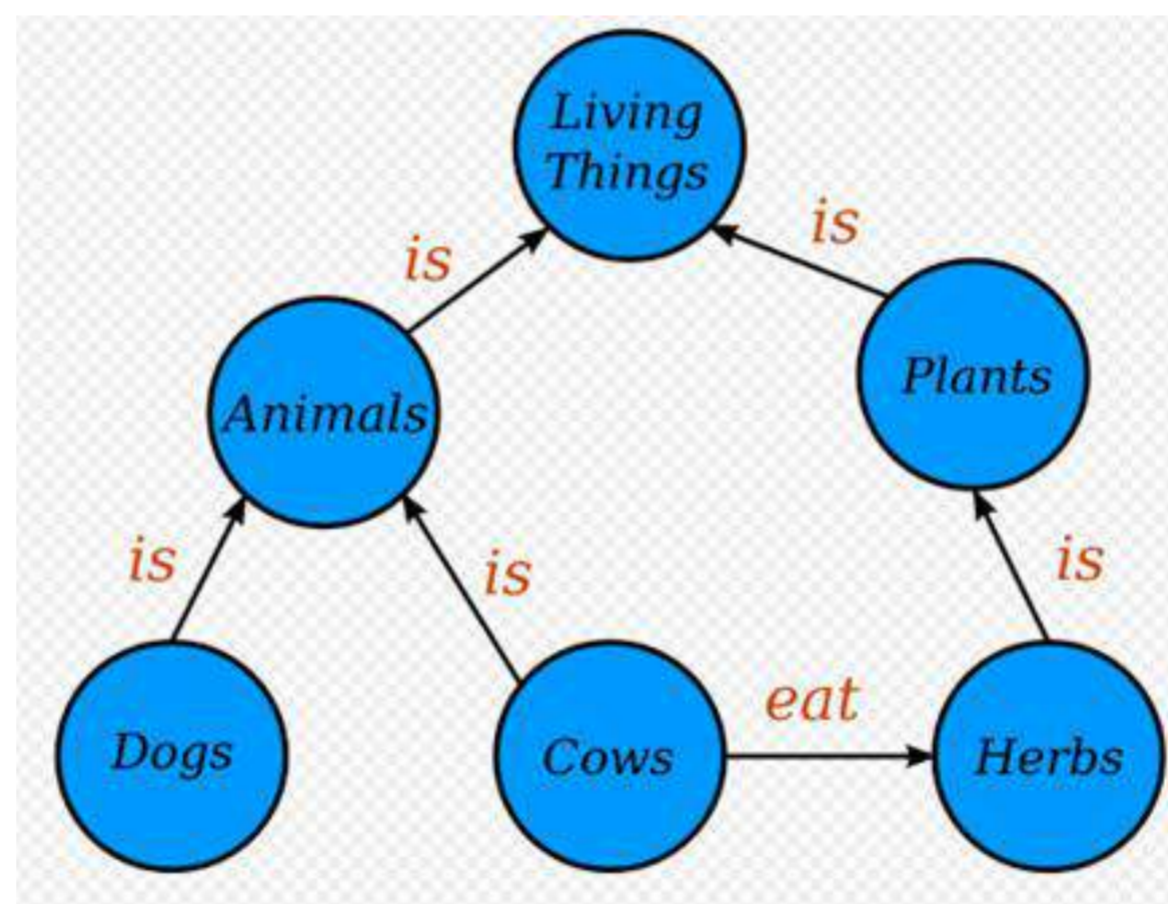
*How Knowledge Graphs Power  
Explainable AI Decisions*

---

# What are Knowledge Graphs?

The digital age has succeeded in producing more data than humans can ever hope to analyze. We need innovative ways to manage this data. This is where the concept of Knowledge Graphs comes in.

A knowledge graph can take an immense amount of data and determine its dependencies, correlations, and other connections between data points in a set. Unlike traditional databases that store data in rows and columns, knowledge graphs organize data as nodes (entities), edges (relationships), and properties (attributes), creating a network of information that is presented as a graph. Here is a simple example of a knowledge graph:



Source: [https://en.wikipedia.org/wiki/Knowledge\\_graph](https://en.wikipedia.org/wiki/Knowledge_graph)

This graph-based approach facilitates analysis and allows users to quickly derive meaningful insights from their massive amounts of data.

Google coined the term "knowledge graph" in 2012 when it started using semantic search information to enhance its search engine results. Since then, the concept has expanded beyond search engines to many industries and business functions. Gartner forecasts global investment in knowledge graphs will grow to \$3.2 billion by 2025, with a CAGR of 28.1%.

---

# Why Knowledge Graphs Are Important

Knowledge graphs depict relationships within data, helping make complex data sets more understandable and usable. They are especially useful in scenarios where relationships and the context of data points are needed to provide comprehensive insights. For example, in healthcare, knowledge graphs can link seemingly-unrelated data from patient records, clinical studies, and healthcare databases to provide personalized treatment recommendations.

Knowledge graphs also allow businesses to leverage their data more effectively to make informed decisions. For example, a business can improve product recommendations and customer service by understanding customer preferences, behavior, and relationships between products.

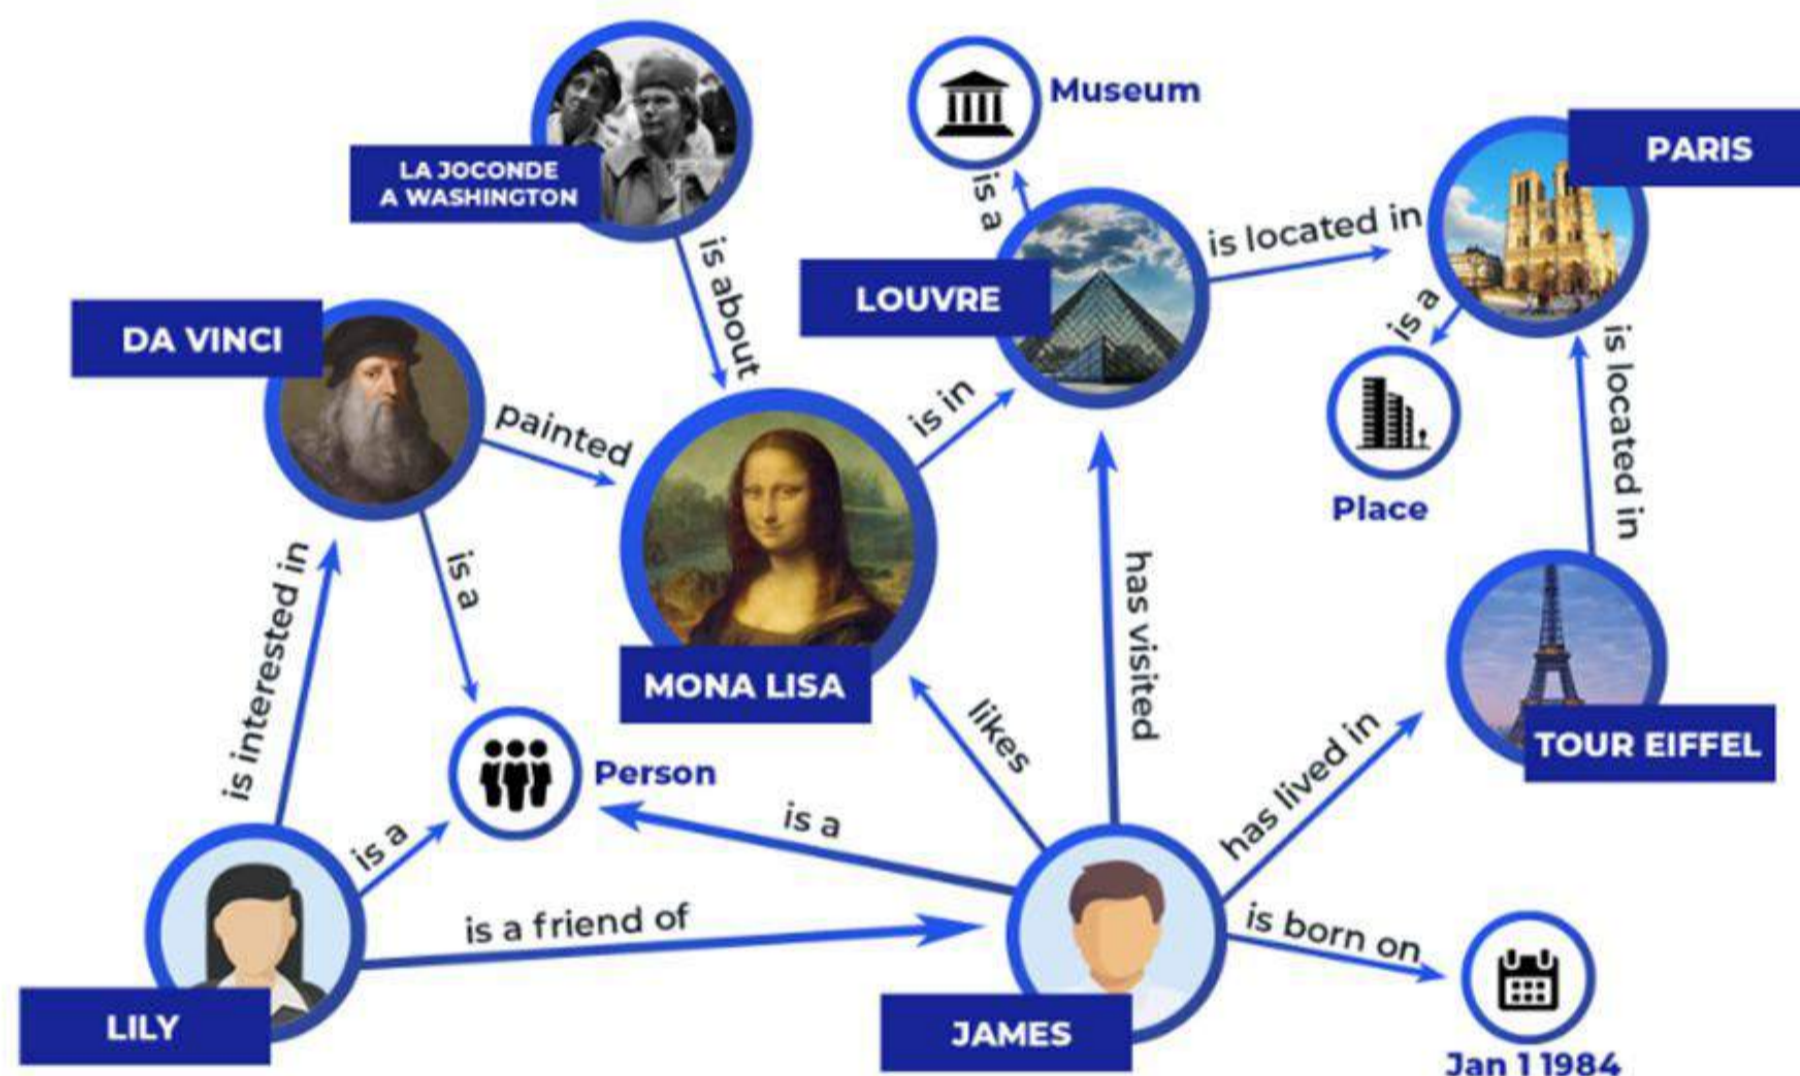
As data continues to grow in volume and variety, the importance of knowledge graphs in managing complex datasets efficiently and accurately becomes increasingly critical.

---

# What are Knowledge Graphs Used For?

## Data Integration

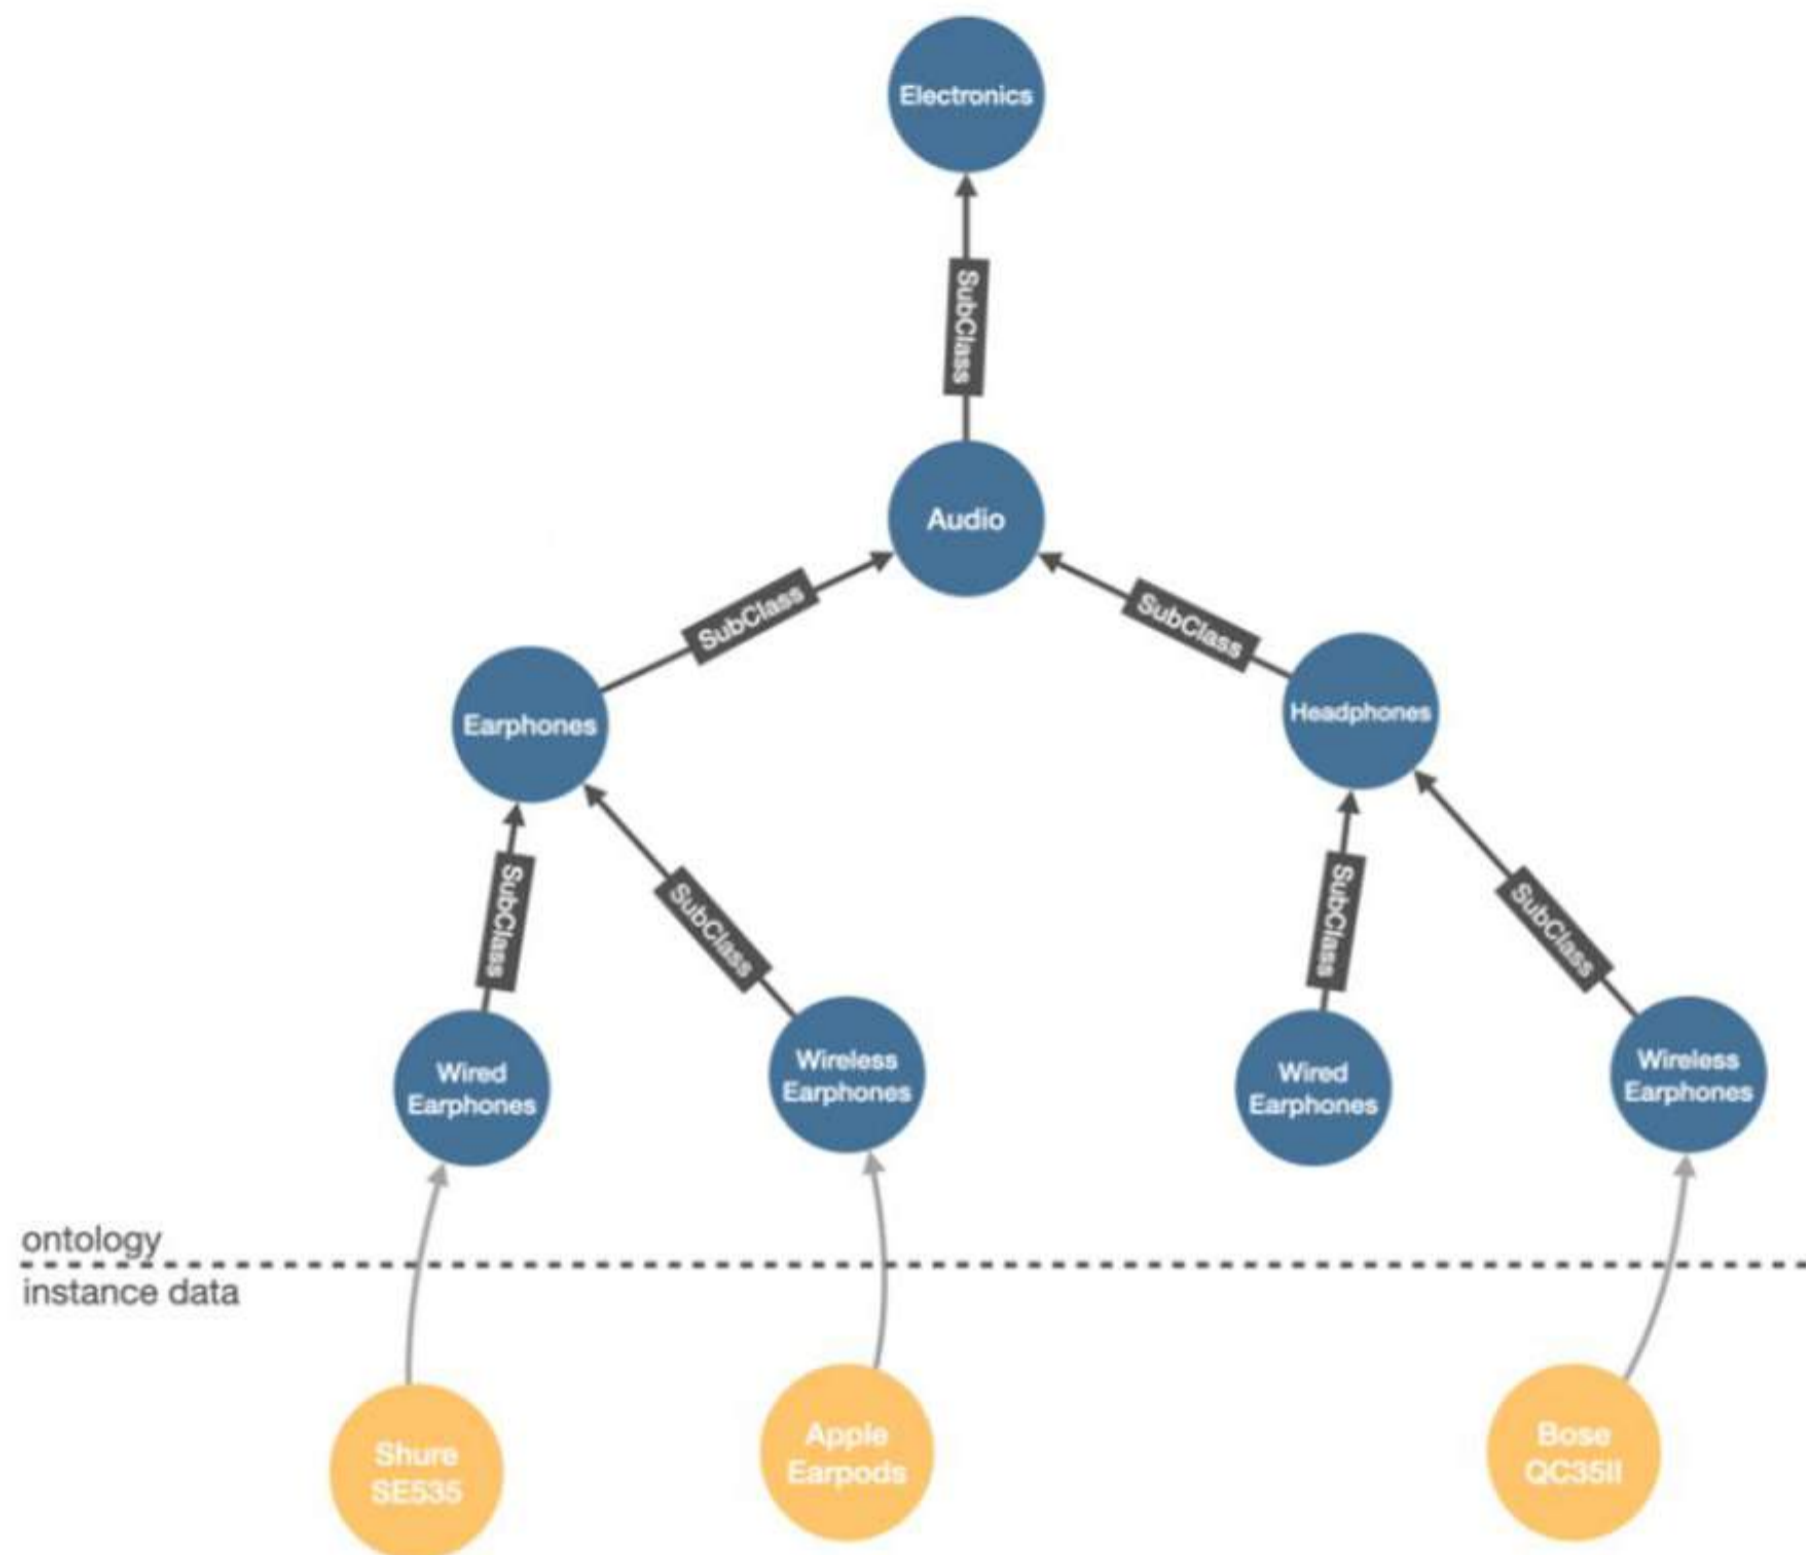
Knowledge graphs connect related pieces of data across different systems, creating a unified view that can be more easily queried and analyzed. This works well based on the semantic nature of knowledge graphs, where data is linked through meaningful relationships. Here is an example:



Source: <https://www.atulhost.com/what-is-knowledge-graph>

## Semantic Querying

With semantic querying, users can ask for information based on the meaning of the data rather than the structure. An example of a semantic query is finding all products that are substitutes for a discontinued line, considering both product characteristics and historical sales data. Here is an example where a customer searching for Apple AirPods might be offered Bose or Shure alternatives if the Apple option is out of stock.



Source: <https://neo4j.com/blog/what-is-knowledge-graph/>

## AI and Machine Learning

Knowledge graphs are also increasingly used in AI and machine learning to enhance the models. By providing a richly structured data environment, knowledge graphs allow AI to understand the relationships and contexts of data sets, leading to more accurate predictions and recommendations. For example, recommendation systems in e-commerce platforms use knowledge graphs to suggest products that are not only similar but also complement each other based on user purchase history and product characteristics. Think of all the right-on-target Netflix or Amazon recommendations you've received.

---

# How do Knowledge Graphs Work?

Knowledge graphs are built on three foundational elements: nodes, edges, and properties.

**Nodes:** These are the primary entities in a knowledge graph, such as people, places, organizations, or events. Each node represents a unique entity or concept.

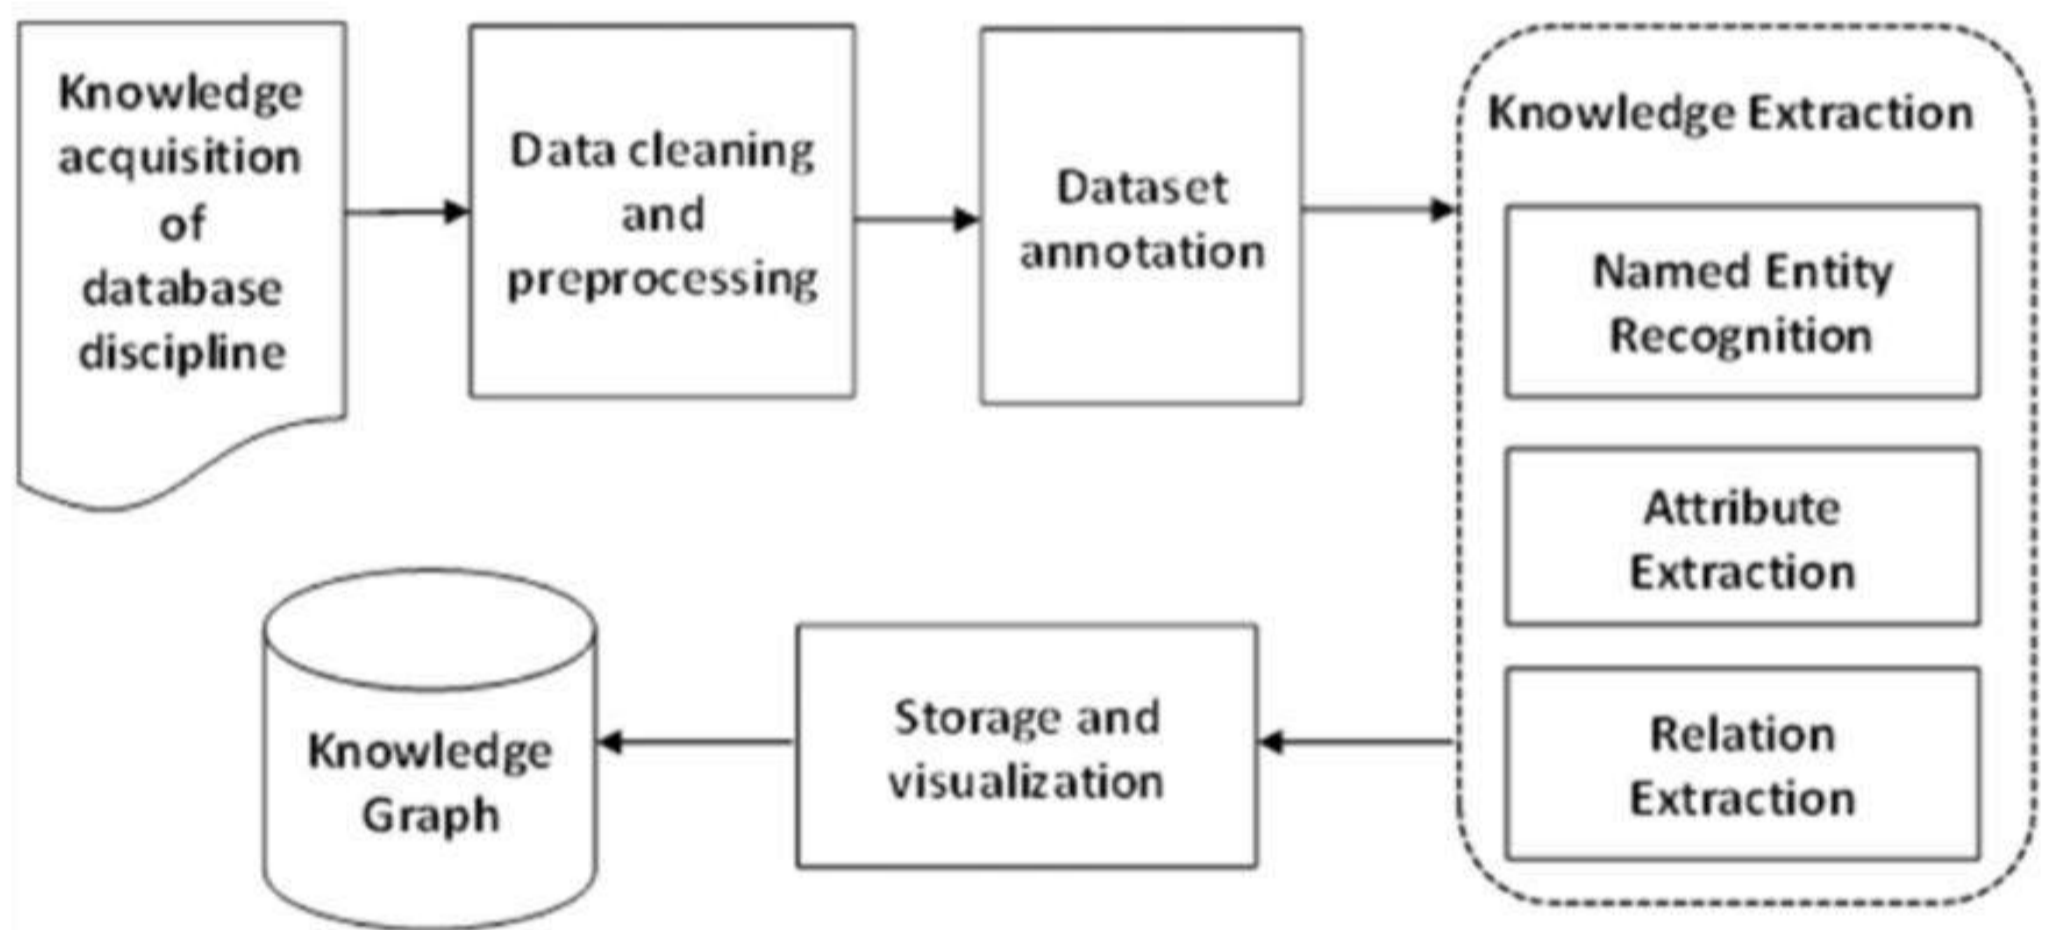
**Edges (Relationships):** Edges connect nodes and show the relationships between them, such as "works at," "enjoys doing," "located in," or "married to."

**Properties:** Both nodes and edges can have properties that provide additional details about the entities or their connections. For example, a person node might have properties like age, address, gender, and profession.

The architecture of a knowledge graph is designed to ensure efficient data integration, querying, and management. Knowledge graphs can be stored in specialized graph databases like Neo4j, which are optimized for handling connected data. Other common open-source knowledge graphs are DBpedia, Freebase, OpenCyc, Wikidata, and YAGO. In addition, indexing is crucial for quick retrieval of nodes and relationships, especially in large graphs.

Processing data within a knowledge graph involves several key operations:

- **Data Ingestion:** Integrating data from various sources into the graph, often involving data transformation and normalization.
- **Data Enrichment:** Adding new nodes and relationships to the graph, including the use of external APIs or datasets to enhance the graph's information.
- **Query Execution:** Executing complex queries that cross multiple nodes and relationships to extract insights or identify patterns.



Source: [https://www.researchgate.net/figure/Knowledge-graph-construction-steps\\_fig1\\_343718462](https://www.researchgate.net/figure/Knowledge-graph-construction-steps_fig1_343718462)

---

# How to Build and Manage Knowledge Graphs

Constructing a knowledge graph requires careful planning and execution. Here are the essential steps:

## **Step 1: Define the Scope and Schema**

Identify the types of entities, relationships, and properties that are relevant to the project context.

## **Step 2: Select the Technology Stack**

Choose appropriate graph database technologies and tools based on the project's requirements.

## **Step 3: Data Integration and Modeling**

Design the graph model and integrate data from various sources. Ensure the model reflects the relationships and properties needed to support future queries.

## **Step 4: Graph Population**

Populate the graph with data, ensuring that entities are linked correctly and that the data adheres to the defined schema.

Once built, the knowledge graph must be managed on an ongoing basis. This involves tracking the performance of the graph database and query execution; ensuring the accuracy, completeness, and consistency of the graph data; and implementing access controls and encryption to protect sensitive information in the graph. The graph must be kept current by regularly adding new data and updating existing entities and relationships. It is also necessary to plan for growth -- both in terms of data volume and query complexity -- to ensure that the graph remains high-performing and relevant.

---

# How to Query, Analyze, and Visualize Your Knowledge Graph

Knowledge graphs use specialized query languages that enable users to efficiently explore complex relationships and retrieve relevant data. The two most widely used graph query languages are:

**Cypher:** Primarily used with Neo4j, Cypher is designed to be easily readable by humans and focuses on pattern matching within the graph structure.

Example of a query: *Find all colleagues of a person named "Alice" within a company.*

```
MATCH (p:Person {name: "Alice"})-[:WORKS_AT]->(c:Company)<-[:WORKS_AT]-(colleague)
RETURN colleague.name
```

**SPARQL:** Used with RDF-based graph databases, SPARQL allows for querying distributed data and is very flexible in aggregating, filtering, and transforming graph data.

Example of a query: *Find all products made by a company that are located in a specific country.*

```
SELECT ?product
WHERE {
  ?company a :Company ;
           :makes ?product .
  ?product :locatedIn :USA .
}
```

## **Knowledge Graph Analysis**

There are many different advanced analysis techniques. Pathfinding looks for the shortest paths or connections between nodes. This is ideal for use cases such as network analysis and logistics, where understanding the most efficient routes is important for optimizing operations.

Centrality analysis can identify the most influential nodes within a graph. This is particularly useful in social network analysis and recommendation systems, where selecting the right influencers is essential for target marketing efforts. Community detection groups nodes into clusters based on their connections, which is important for market segmentation and detecting fraud patterns.

## **Visualization Tools**

Making sense of these complex relationships is difficult, yet essential, especially for presentations to stakeholders who may not be familiar with raw graph data. There is a short list of tools that facilitate the clear visualization of knowledge graphs. Neo4j Bloom is a popular choice that offers a straightforward and user-friendly interface for exploring and visualizing Neo4j graph databases. For those dealing with large graphs, Gephi, an open-source network analysis and visualization software, offers robust tools for deep data exploration. Finally, Graphviz provides highly customizable graph drawing tools, allowing for the creation of detailed diagrams that can visually represent intricate graph structures and insights. These visualization tools help users understand and present their data clearly.

---

# How to Use Machine Learning to Enrich Your Knowledge Graph

Integrating machine learning with knowledge graphs can go a long way toward automating data enrichment and enhancing the decision-making process. For example:

**External Data Integration:** Organizations can automatically enrich knowledge graph nodes with additional data from external APIs or databases. For example, this type of data integration has been proven helpful for know-your-customer (KYC) and anti-money laundering initiatives within the finance industry.

**Continuous Learning:** Supervised learning systems or unsupervised feedback mechanisms help the knowledge graph learn from new data over time. For example, DevOps teams can adjust load distribution on the fly based on real-time traffic.

**Predictive Modeling:** Machine learning models can use knowledge graph-based features to predict outcomes. For example, it can detect patterns of fraud or suspicious behavior in user accounts.

**Recommendation Systems:** Organizations can leverage the relationships in a knowledge graph to recommend products or services. This personalization helps online sites recommend new content for users to read or watch based on engagement.

---

# Challenges of Implementing Knowledge Graphs in Real-World Scenarios

Implementing knowledge graphs in real-world business scenarios presents several challenges to organizations.

**Data Silos:** In many organizations, data is fragmented across various departments and stored in different formats, making it difficult to achieve a unified view of information. To address this, companies can implement data integration strategies that consolidate various data sources into a single knowledge graph. For example, an IT department might integrate data from their various infrastructure components -- such as servers, storage, network devices, and applications -- into a single knowledge graph to better understand dependencies, and to assess the impact of a planned change across the system to minimize downtime and ensure smooth operations.

**Scalability:** As knowledge graphs grow to include millions or even billions of relationships, maintaining performance becomes increasingly complex. Organizations can mitigate these issues by choosing scalable knowledge graph databases and technologies that support horizontal scaling and efficient querying mechanisms even during peak periods.

**Data quality:** The knowledge graph's data accuracy is essential for reliable analytics and decision-making. Regular data validation and cleaning processes—along with automated checks and balances—go a long way to maintaining data integrity. An example here could be a financial institution that employs automated scripts to periodically check and rectify inconsistencies in transactional data across different accounts in its knowledge graph, ensuring that financial advisors have access to reliable information for client consultations.

---

# Why Using Knowledge Graphs in Your VPC Makes Sense

Despite the increasing popularity of knowledge graphs, many organizations hesitate to use them due to security concerns. They are often wary of sending their most sensitive data, like CRM or financial information, to the cloud. Instead, some may choose to try to build their own infrastructure for knowledge graph – an approach that has proven to be costly, as it requires significant DevOps resources to create and maintain an in-house knowledge graph infrastructure. The good news is that organizations no longer have to choose between security and high development costs: new advancements now allow organizations to use knowledge graphs within their own Virtual Private Cloud (VPC) using data and AI operating systems. To find out more about the benefits of using knowledge graphs through a data and AI operating system like Shakudo, [click here](#).

---

# Closing Words

As we've explored, knowledge graphs transform massive and disparate data sets into organized, interconnected networks that allow more intelligent and effective use of data. From healthcare providing personalized treatment options to e-commerce platforms like Amazon and Netflix delivering perfectly-targeted recommendations, knowledge graphs offer significant advantages. These benefits include enhanced integration, richer semantic querying capabilities, and the ability to leverage AI for advanced data analysis and effective decision-making.

The growing integration of AI and machine learning with knowledge graphs is pushing the boundaries of what these tools can achieve. By automating data enrichment and enhancing predictive modeling, knowledge graphs are set to become even more integral to business strategies and decision-making processes. The continuous learning capabilities ensure that businesses remain at the forefront of innovation and efficiency.

As technology evolves, knowledge graphs will become even more critical in harnessing the full potential of big data, with expanding applications in AI, machine learning, and beyond. The future of knowledge graphs certainly looks promising.



## ABOUT SHAKUDO

Shakudo creates compatibility across the best-of-breed data tools for a more reliable, performant, and cost effective data and AI operating system. As an operating layer on top of your cloud Shakudo allows you to pick the best-of-breed data tools for your needs, while providing a platform with fully automated DevOps experience. This combines the best of both worlds in data stack practices so you can focus on delivering business value with data.

Shakudo is the most **easy, secure, cost-effective, scalable** way to bring the most advanced data and AI tools to your data. Find out more at **[shakudo.io](https://shakudo.io)**.