ENTERPRISE **KNOWLEDGE**

Global Knowledge & Information Management Services

Using Knowledge Graph Data Models to Solve Real Business Problems

Lulit Tesfaye

A successful business today must possess the capacity to quickly glean valuable insights from massive amounts of data and information coming from diverse sources. The scale and speed at which companies are generating data and information, however, often makes this task seem overwhelming.

An [Enterprise Knowledge Graph](#) allows organizations to connect and show meaningful relationships between data regardless of type, format, size, or where it is located. This allows us to view and analyze an organization's knowledge and data assets in a format that is understood by both humans and machines.

> *While most organizations have been built to organize and manage data by **department** and **type,** a Knowledge Graph allows us to view and connect data the way our brain relates and infers information to answer specific business questions, **without** making another copy of the data from its original sources.*

How much of your data and information is currently dispersed across business units, departments, systems, and knowledge domains? How many clicks or reports do you currently navigate to find an answer to a single business problem or get relevant results to your search?

Below, I will share a selection of real business problems that are able to be tackled more efficiently by using knowledge graph data models, as well as examples of how we have used knowledge graphs to better serve our clients.

Aggregate Disparate Data in Order to Spot Trends and Make Better Investment Decisions

A vast amount of the data we create and work with is unstructured, in the form of emails, webpages, video files, financial reports, images, etc. Our own wide assessment of organizations finds that as much as 85% of an organization's information exists in an unstructured

form. Organizing all of these data proves to be a necessary undertaking for many large and small institutions in order to extract meaning and value from their organization's information. One way we have found to make this manageable is through leveraging semantic models and technologies to automatically extract and classify unstructured text to make it machine readable/processable. This allows us to further relate this classified content with other data sources to be able to define relationships, understand patterns, and quickly obtain holistic insights on a given topic from varied sources, despite where the data and content lives (business units, departments, systems, and locations).
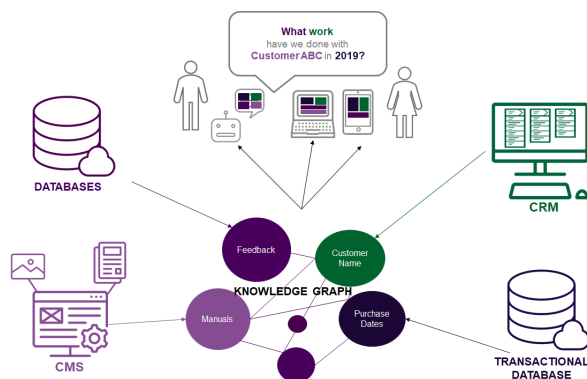
One of the largest supply chain clients we work with needed to provide its business users a way to obtain quick answers based on very large and varied data sets. The goal was to bring meaningful information and facts closer to the business to make funding and investment decisions. By extracting topics, places, people, etc. from a given file, we were able to develop an ontology to describe

the key types of things business users were interested in and how they relate to each other. We mapped the various data sets to the ontology and leveraged semantic Natural Language Processing (NLP) capabilities to recognize user intent, link concepts, and dynamically generate the data queries that provide the response. This enabled non-technical users to uncover the answers to critical business questions, such as:

- Which of your products or services are most profitable and perform better?

- What investments are successful, and when are they successful?

- How much of a given product did we deliver in a given timeframe?

- Who were my most profitable customers last year?

- How can we align products and services with the right experts, locations, delivery method, and timing?

Discover Hidden Facts in Data to Predict and Reduce Operational Risks

By allowing organizations to collect, integrate, and identify user interest and intent, ontologies and knowledge graphs build the foundations for Artificial Intelligence (AI) to allow organizations to analyze different paths

jointly, describe their connectivity from various angles, and discover hidden facts and relationships through inferences in related content that would have otherwise gone unnoticed.

What this means for our large engineering and manufacturing partner is that, by connecting internal data to analyze relationships and further mining external data sources (e.g. social media, news, help-desk, forums, etc.), they were able to gain a holistic view of products and services to influence operational decisions.

Examples include the ability to:

- Predict breakdown and detect service failures in early stages to schedule preventive maintenance, minimize downtime or maximize service life;

- Maintain right level of operator, experts and inventory;

- Find remaining life of an asset or determine right warranty period; and

- Prevent risk of negative brand image and impacts to lifetime loyalty of customers.

Facilitate Employee Engagement, Knowledge Discovery, and Retention

Most organizations have accumulated vast amounts of structured and unstructured data that are not easy to share, use, or reuse among staff. This difficulty leads to diminished retention of institutional knowledge and rework. Our work with a global development bank, for instance, was driven by the need to find a better way to disseminate information and expertise to all of their staff so that projects would be more efficient and successful and employees would have a simple knowledge sharing tool to solve complex project challenges without rework and knowledge loss. We developed a semantic hub, leveraging a knowledge graph, that collects organizational content, user context, and project activities. This information then powers a recommendation engine that suggests relevant articles and information when an email or a calendar invite is sent on a given topic or during searches on that topic. This will eventually power a chatbot as part of a larger AI Strategy. These outputs were then published on the bank's website to help improve knowledge retention and to showcase expertise via Google recognition and search optimization for future reference.

Using knowledge graphs based on this linked data strategy enables the organization to connect all of their knowledge assets in a meaningful way to:

- Increase the relevancy and personalization of search;

- Enable employees to discover content across unstructured content types, such as webinars, classes, or other learning materials based on factors like location, interest, role, seniority level, etc.; and

- Further facilitate connections between people who share similar interests, expertise, or location.
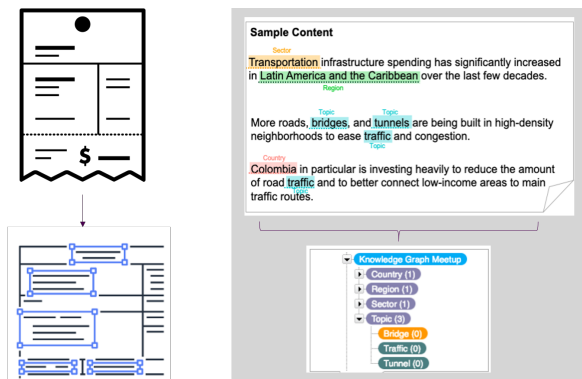
## Implement Scalable Data Management and Governance Models

The size, variety, and complexity by which businesses are integrating data is making it difficult for IT and data management teams to keep up with a traditional data management structure. Currently, most organizations are facing challenges to efficiently and properly map ownership and usage of various data sources, to track changes, and to determine the right level of access and security. For example, we worked with a large US Federal Agency that needed a better way to manage the thousands of data sets that they use to develop economic models that drive US policy. We developed a semantic data model that captures data sets, metadata, how they relate to one another, and information about how they are used. This information is stored in a triple store that sits behind a custom web application. Through this ontology, a front-end semantic search application, and an administrative dashboard, the model provided Economists and the Agency:

- A single unified data view that clearly represents the knowledge domains of the organization without copying or duplicating data from the authoritative and managed sources;

- A machine-processable metadata from documents, images, and other unstructured content as well as from relational databases, data lakes, and NoSQL sources to see relationships across data sets; and

- An automated log of usage so that the Agency can better evaluate the internal value of a data set which then allows it to better negotiate deals with data vendors.

Lay the Foundation for AI Strategy

Considered to be the primary catalyst for what is now being termed the 4th industrial revolution, Artificial Intelligence (AI) is expected to drive half of our world's economic gains within a decade. Where KM and AI meet, we discuss [Knowledge Artificial Intelligence](#), a key element of this forthcoming revolution. For organizations looking to dabble in a pragmatic and scalable AI strategy, their enterprise needs to have a solid data practice that lays down the infrastructure necessary to sow and reap the advantages of data across disparate sources, as well as drive scale and efficient governance through graph-based machine learning.
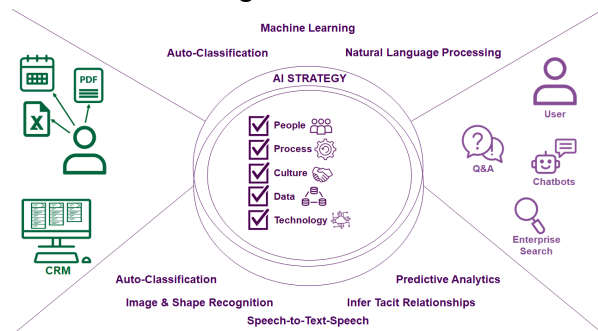
The key approach that the business problems above share is that the knowledge graph modeling is enabling the application of some form of AI to transform the productivity and growth of these organizations with "hard" and "soft" business returns. These range from improvements to processes through staff augmentation and task automation, employee and customer satisfaction through personalized sales and marketing, risk avoidance through anomaly detection and predictive analytics, and facilitating staff intelligence through the enablement of natural ways to ask the toughest business questions.

Conclusion

Data lakes and data warehouses have gotten us this far by allowing organizations to integrate and analyze large data to drive business decisions, but the practicality of data consolidation will always be a limiting factor for the enterprise implementation of these technologies. As business agility and high volumes of data are becoming the ingredients for success, the need for speed, exploration, scalability, and optimization of data is becoming an undertaking that traditional and relational data models are struggling to keep up.

Backed by operational convenience, semantic data models provide organizations the power to synthesize real-time decisions, make relevant recommendations and facilitate knowledge sharing with limited administrative burden and a proven potential for scale. What further sets graph models apart is that they rely on context from human knowledge, structure, and reasoning that are necessary to relate knowledge to language in a natural way.

Graph data models are able to leverage machine learning to apply this knowledge by collecting and automatically classifying knowledge from various sources into machine readable formats, and allowing the organization to benefit from the deeper layers of AI, such as natural language processing, image recognition, predictive analytics, and so much more.

If your organization, like many, is facing challenges in surfacing, exploring, and managing data in an understandable and actionable manner, contact us to learn more on how we can help.

Enterprise Knowledge (EK) is a services firm that integrates Knowledge Management, Information Management, Information Technology, and Agile Approaches to deliver comprehensive solutions. Our mission is to form true partnerships with our clients, listening and collaborating to create tailored, practical, and results-oriented solutions that enable them to thrive and adapt to changing needs.

Our core services include strategy, design, and development of Knowledge and Information Management systems, with proven approaches for Taxonomy Design, Project Strategy and Road Mapping, Brand and Content Strategy, Change Management and Communication, and Agile Transformation and Facilitation. At the heart of these services, we always focus on working alongside our clients to understand their needs, ensuring we can provide practical and achievable solutions on an iterative, ongoing basis.

info@enterprise-knowledge.com | 571-403-1109 | @EKConsulting