Deloitte.

Wisdom of Enterprise Knowledge Graphs

The path to collective intelligence within your company

“Intelligence is the ability to adapt to change.”

Stephen Hawking

Preface

Today, companies are relying more and more on artificial intelligence in their decision-making processes. Semantic AI enables machines to solve complex problems and explain to users how their recommendations were derived. It is the key to collective intelligence within a company – manifested in Knowledge Graphs.

Unleashing the power of knowledge is one of the most crucial tasks for enterprises striving to stay competitive. Knowledge, however, is not just data thrown into a database. It is a complex, dynamic model that puts every piece of information into a larger frame, builds a world around it and reveals its connections and meaning in a specific context.

In reality,a database can only accommodate a small fraction of a company's knowledge, much less the combined knowledge of a group of experts. A lot of vital information is stored in unstructured documents, images and people's minds.

For decades, companies have been trying to figure out an intelligent way to integrate all expert knowledge in one magic place, driven by a belief in the power concentrated in the combined experience of thousands of employees.

That said, it is extremely time consuming to share your domain knowledge. First you have to structure and adapt the information to fit into a pre-defined data model.

Knowledge Graphs connect knowledge from different domains, data models and heterogeneous data formats without changing their initial form. They enable industrial enterprises to harness the potential of collective intelligence.

By design, an Enterprise Knowledge Graph will easily connect all knowledge sources within a company, thereby enabling their AI technologies to go beyond conventional Machine Learning.

The search for information takes 14–30 percent of the engineers' time.

Fig. 1 – Knowledge Graphs support highly complex decision-making by considering expert knowledge from different domains. Real world dependencies and cross-correlations are taken into account before recommendations are derived and explained to humans.

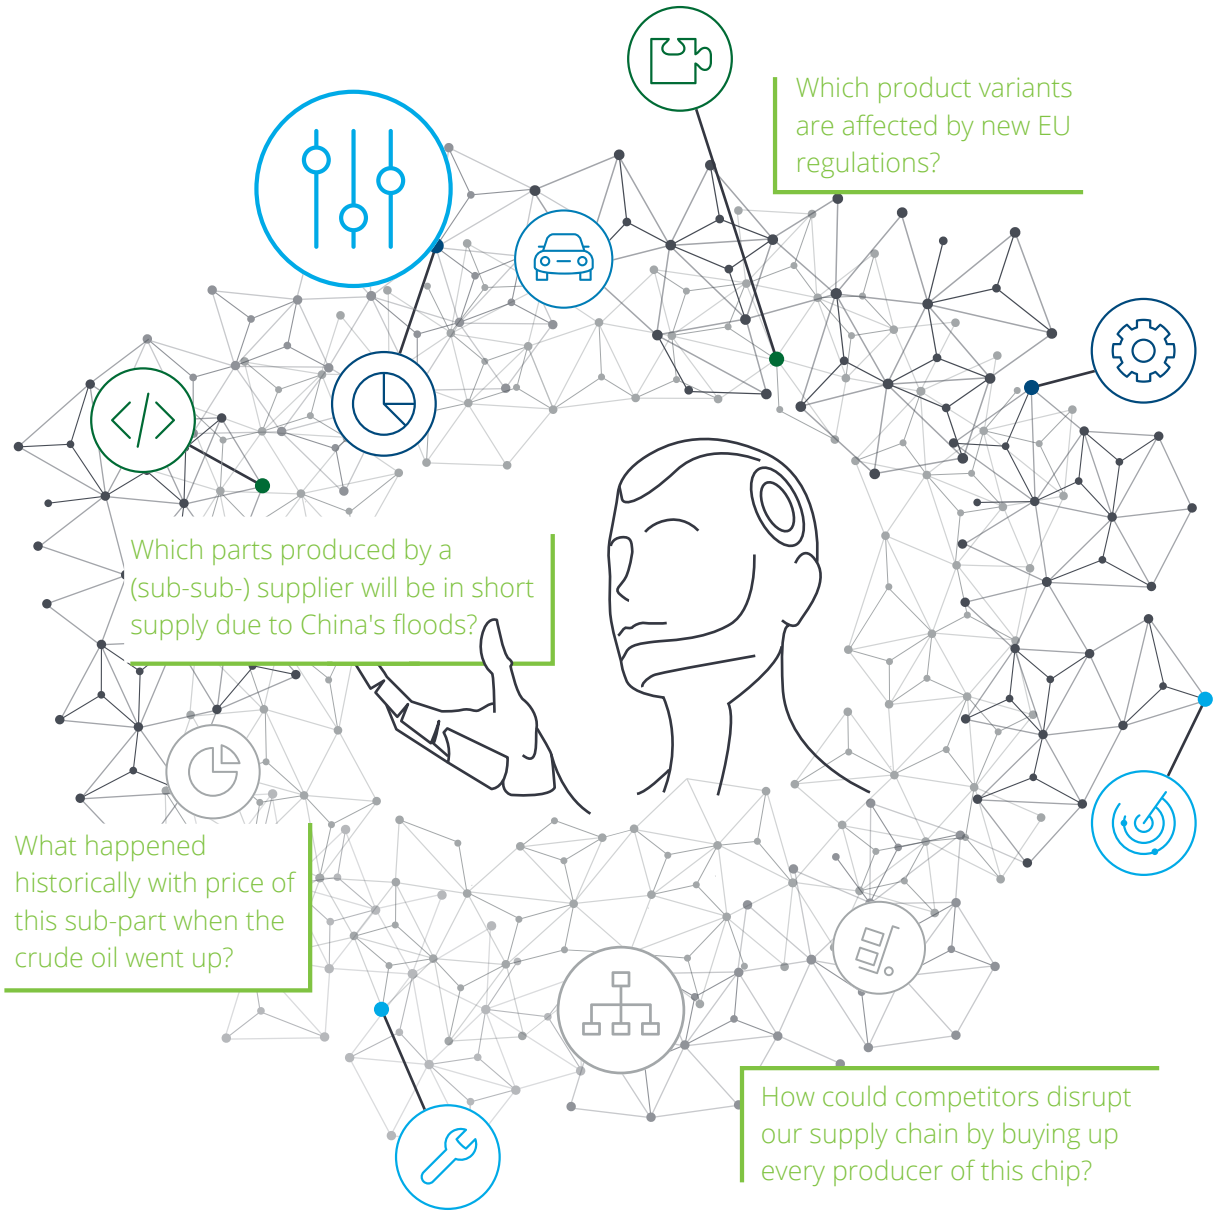

Use Cases

Impact analysis

The Enterprise Knowledge Graphs reveal the far-reaching impacts of seemingly small decisions.

There are literally thousands of decisions being made across a company every day in areas like product modifications, production plans or the supply chain, but even small decisions can have huge consequences. Though science has effectively debunked the well-known butterfly effect, it is an undisputed fact that business decisions can have a far-reaching impact.

Our ability to grasp the range of risks arising from a decision may, however, be limited when we use an opportunistic mindset. It is much too complex for us to imagine all of the possible future scenarios arising from every single factor across an enterprise in our intuitive risk calculations, and we may fail to take key issues into account in our decision-making.

An Enterprise Knowledge Graph can assess the far-reaching impacts of decisions and their consequences for any area of the company. It combines the human understanding of the real-world interdependencies with the power of computation. As a result, the system can detect cross correlations that humans cannot.

Fig. 2 – The Butterfly Effect after a seemingly small decision: Knowledge Graphs reveal the far-reaching consequences of a product change.

Example
To reduce internal friction in a combustion engine, the engineer reduces the surface roughness by 10 µm without realizing that the n-tier supplier does not have the machinery required to produce this tolerance. The result: a massive investment and a delayed start of production.

Using a Knowledge Graph to connect the engineering information with the manufacturing information across the value chain alerts companies to all possible risks arising from a change of tolerance particularly in terms of the supplier's production capability. An AI-based analysis of internal and external sources can also assess the risk of sub-suppliers going out of business and how that would impact requirements. Even though this is common practice for OEMs – the only way to achieve truly exhaustive results is to make the connection with technical product information.

Use Cases

AI-based decision-making

Integration of all product data within a knowledge graph leads to process cost savings of up to 65 percent and enables managers to make well-founded decisions.

Today's companies are relying more and more on artificial intelligence in their decision-making process. Too often, these systems do not allow users to determine why the AI recommends a certain course of action – even though traceability can be essential for critical cases like product liability provided by the AI.

A Knowledge Graph is a new form of AI that allows us to answer the question "Why". It contains all of a company's business logic and allows to question, backtrack and explain any recommendation proposed by the AI.

Fig. 3 – A knowledge-based representation of real-world interdependencies enables machines to solve complex problems and explain to users how it derived its recommendations.

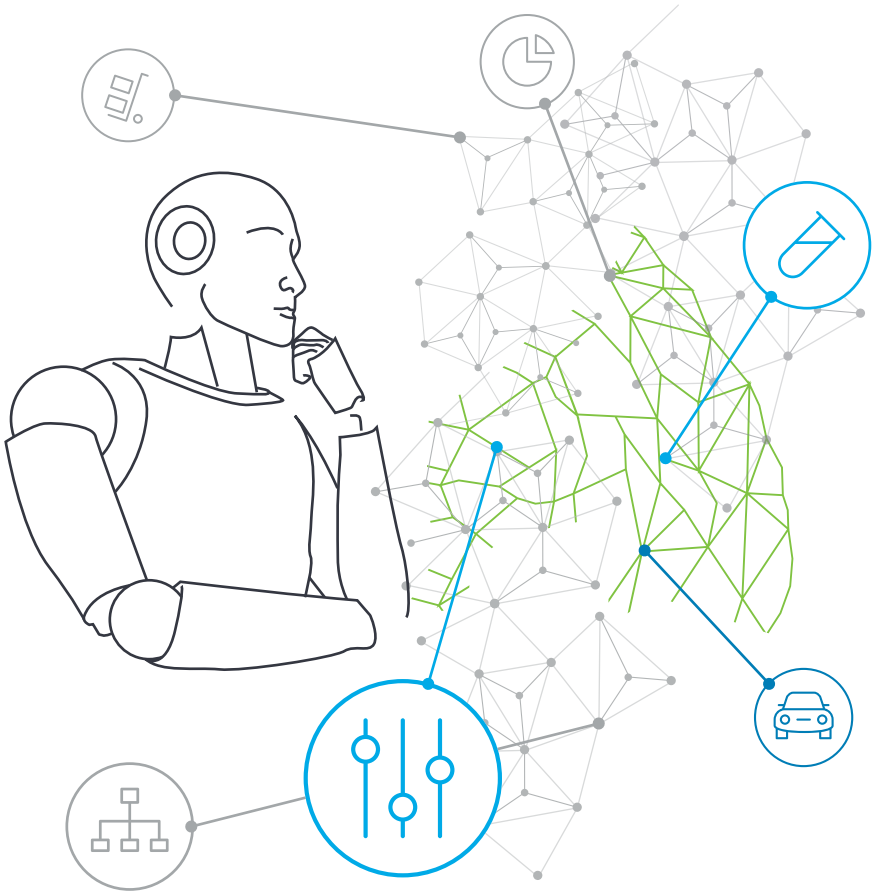

Example
Complex mechatronic products such as cars, airplanes and engineering machinery. usually have a large number of variants. Each variant has its own set of configuration rules. In cars, for example, a single component can impact 90% of the other roughly 50,000 components on average. It takes tens of thousands of such configuration rules to map these relationships.

The components within the configurations are described by a limited list of structured attributes. Human intelligence can associate thousands of partially unstructured pieces of context information on each part in terms of its function, geometry, logistics, quality and many more areas. This further increases the complexity of the product structure.

A knowledge-based representation of configuration rules allows us to automate our analysis of the consistency, completeness and redundancy within the configuration rule set. The system also automatically recommends how to resolve a conflict. As it is visualized in an intuitive graph, the information is easily accessible for human workers. Even unskilled users can trace back any recommendation to individual rules and conflicts. Being able to enrich the configuration knowledge with information about logistics, finance and production enables AI powered systems to derive even more complex recommendations, such as removing a configuration from the product portfolio or reorganising the assembly plan.

Use Cases

Product DNA and root-cause analysis

WIthin seconds, we can resolve complex product structures that are configurable on many levels and compare different states.

The Enterprise Knowledge Graph not only connects information using enterprise business logic but also adds temporal context. Thus, any part, service, money transfer or test result can be traced back in time through the different versions, transformations and operations that caused them. We can create – in milliseconds – a complete DNA profile for any object that comprises all of the data associated with it throughout its lifecycle, including the resolution of the entire production structure.

Example
Where product failure is caused by a combination of ambient conditions, the batch of raw materials used or the supplier production site, a classic root-cause analysis can be very time consuming and may only produce limited results.

Knowledge Graphs can statistically analyze root causes by connecting data from different domains without creating a hypothesis in advance. The statistical analysis uses the intelligence within the Knowledge Graph to avoid spurious correlations.

Fig. 4 – Semantic root-cause analysis based on a Knowledge Graph.

Use Cases

Intelligent data governance

Knowledge Graphs store the entire business logic of an enterprise and enable intelligent master data governance.

Given the changing market conditions, digitalization of product development, and increasing product complexity and customization, companies are faced with new challenges that require advanced data interaction.

Knowledge Graphs enable semantic technologies to automate processes like data structuring, text analysis and data model merging. They can also integrate structured and unstructured data and facilitate data clean up and enrichment based on reasoning. Merging data models from different sources and creating rules to avoid inconsistencies will become an intuitive and efficient process.

The semantic backbone relies on customized metalevel ontologies based on the graphs, which help to connect data from different systems and guarantee consistency through logical inferencing and similarity calculations. At the same time, the machine learning component derives rules to automatically detect redundancies and inconsistencies in the data.

Example
The system landscape at an automotive OEM has evolved over time and presents data inconsistencies between different systems. Materials with the same UUIDs have different attributes, and materials with same attributes often have different identifiers. Integrating the sources within a semantic backbone based on a Knowledge Graph enables the company to monitor data quality very quickly using semantic similarity. Redundant, inconsistent and incomplete data is quickly identified and automatically corrected. At the same time, the machine learning component derives rules based on input data and cleans up data even as it improves the data correction process.

Fig. 5 – Connecting data-sources by storing the transformation logic within a Knowledge Graph.

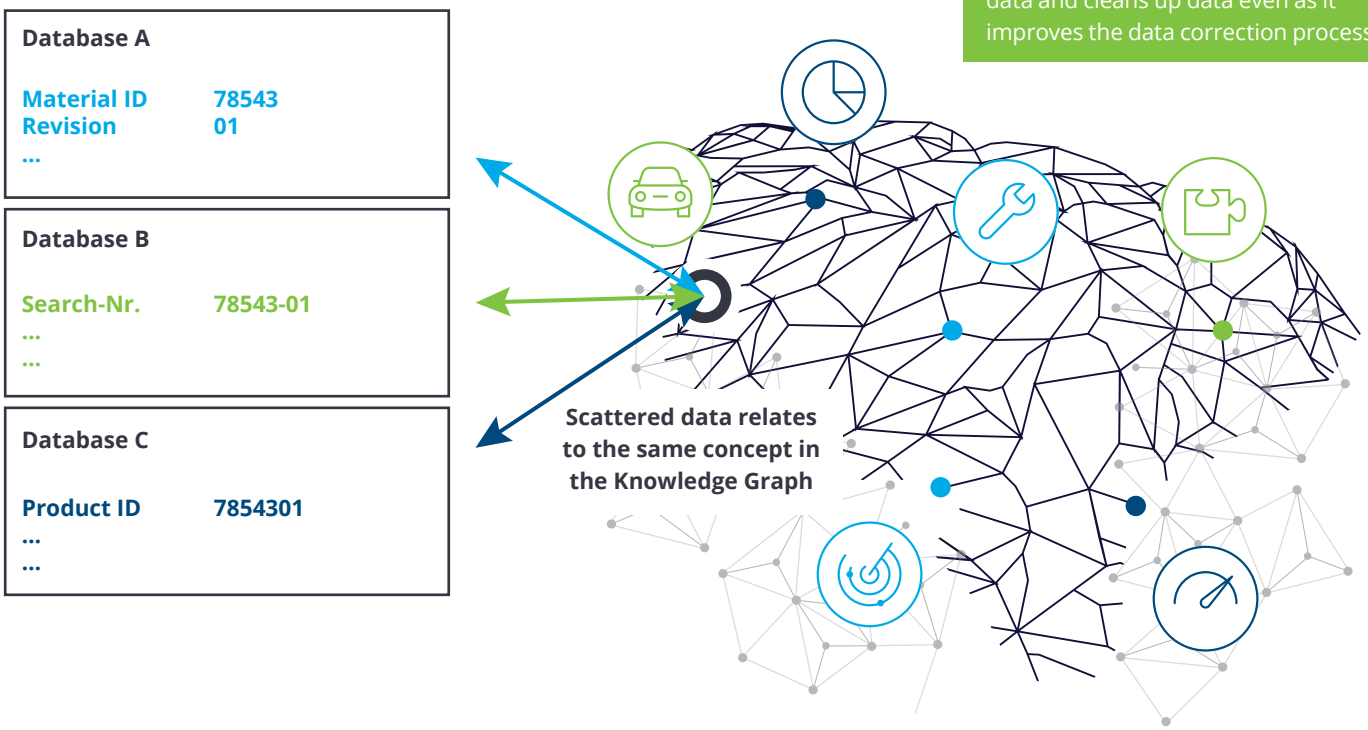

Use Cases

Context-based knowledge transformation

Knowledge Graphs automatically transform and represent knowledge in the required context for any user.

Information can only evolve into knowledge by adding context to it. The act of simply retrieving information on its own is basically useless without a deep understanding of the context, which may vary strongly depending on what the information will ultimately be used for. A mechanical part is associated with, among other things, financial data, engineering data, manufacturing data and procurement data. Depending on their tasks, users access the same data with a different point of view that serves their respective purpose. Storing product structure data and transformation information within a Knowledge Graph allows the user to explore the database within the required context.

Fig. 6 – Context based views on data within the Knowledge Graph.

The Knowledge Graph can be regarded as a cloud of data points watched form different perspectives und thus presenting the same object in different contexts.

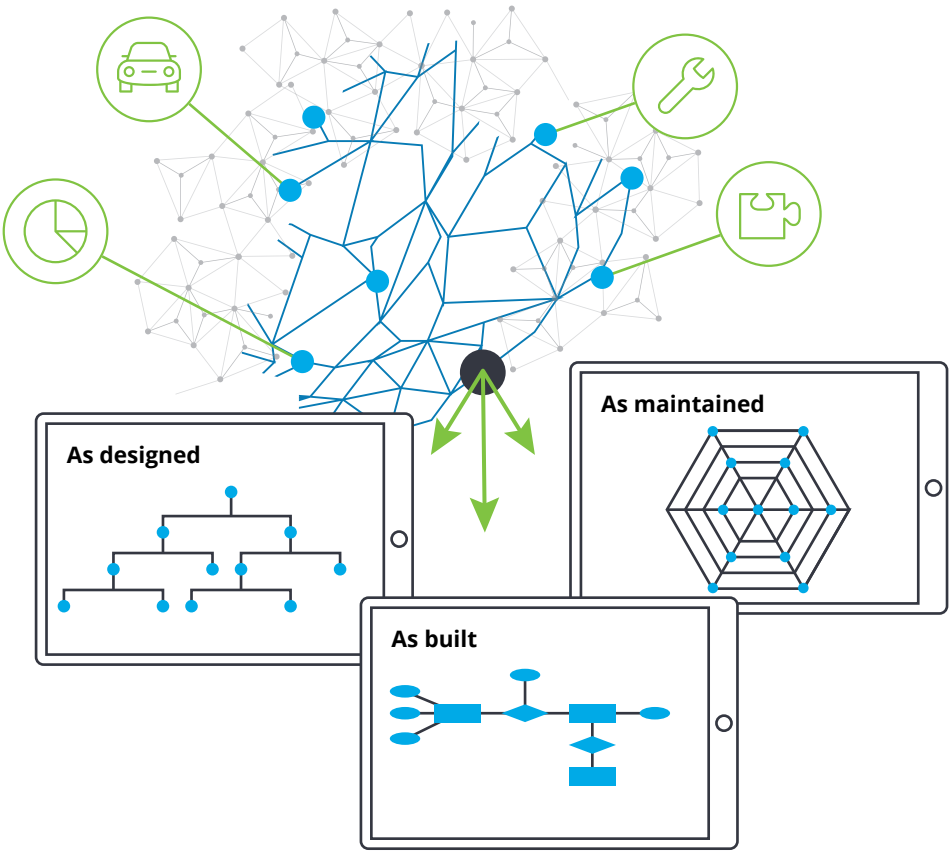

Example
A very common example is the difference between how engineers and manufacturing teams perceive the product structure. A Bill-Of-Materials (BOM) is a structured list of product assemblies, subassemblies, parts, etc. included in the product. The BOM can have different levels of maturity depending on when it is used during the product lifecycle (development, production, after-sales/maintenance). The engineering BOM (E-BOM) contains the hierarchical structure of a product as designed by the engineer. However, when drafting the E-BOM, engineers don't factor in how and where the product will be manufactured. Depending on the component interfaces specified in the CAD model, we can create a site-specific manufacturing BOM (M-BOM) for specific manufacturing processes at the production site or the suppliers. This may entail replacing certain materials, restructuring the production process based on physical limitations or adding non-design items like oil or glue that are not mentioned in the E-BOM.

Use Cases

Leverage knowledge from external sources

Knowledge Graphs easily integrate data from external databases, deep web, open sources and geographic information systems and save costs caused by regulatory and contract conflicts as well as unforeseen risks in companies' ecosystem.

Managing data generated within the company becomes a lot more difficult when you have to consider external factors e.g. legal regulations. The semantic core of Knowledge Graphs makes it easy to integrate external ontologies, vocabularies and databases.

Knowledge Graphs are continuously enriched with information from external regulatory databases, news articles, material databases, supplier information, geospatial data etc. The data is automatically processed with intelligent NLP algorithms and integrated into the graph in a structured and understandable way.

Fig. 7 – Automated integration of unstructured information relating to external regulations or internal contracts and detection of potential compliance issues.

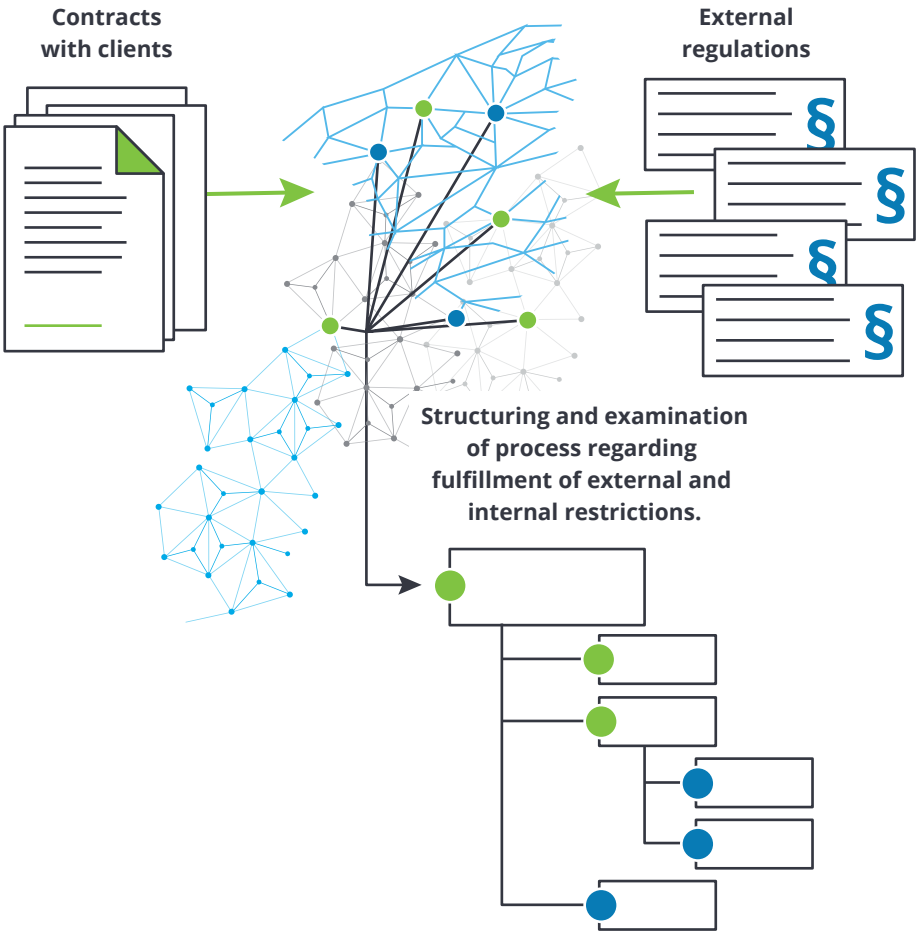

Example
Whether or not a product complies with prevailing regulations and contractual terms will depend on a range of different factors including product functions, production process parameters, storage and transport as well as the respective geographic regions. This information may vary over time and may be stored in part in an unstructured or heterogeneous form, which will require manual analysis.

Product and process complexity can be represented in a Knowledge Graph, enabling users to apply AI-based technologies for analytics and data management such as NLP, text clustering and semantic classification. The system can automate regulatory compliance checks and offer recommendations to resolve compliance conflicts. It can also evaluate the similarities of different regulations and contracts. This will allow risk management to easily simulate different scenarios based on potential regulatory changes and analyze how they impact the company in order to validate business decisions. Similarly, Knowledge Graphs can be used for homologation purposes in the automotive sector.

What are Knowledge Graphs?

A Knowledge Graph is a close-to-reality model that represents a company's business logic and serves as its central knowledge platform. You could say it is the brain of a company, as its architecture is very similar to that of our brains – unlike traditional databases and data lakes. Both connect objects directly with each other rather than storing data in different tables and then connecting them via JOIN-tables.

Experiment:
Think of something blue!
You probably thought of the sky, a car, a shirt or some other object with the color blue that you have recently seen. You brain directly associates all of these things with the named color. It behaves like a graph, a network of entities where the entity "blue" classified as a color is directly connected to cars and shirts that are blue. If our brain worked the same way a traditional relational database does, it would have to look at every item in the entire closet and check the value of the attribute "color". Obviously, that would be a lot more time consuming – both for our brains and for an enterprise database.

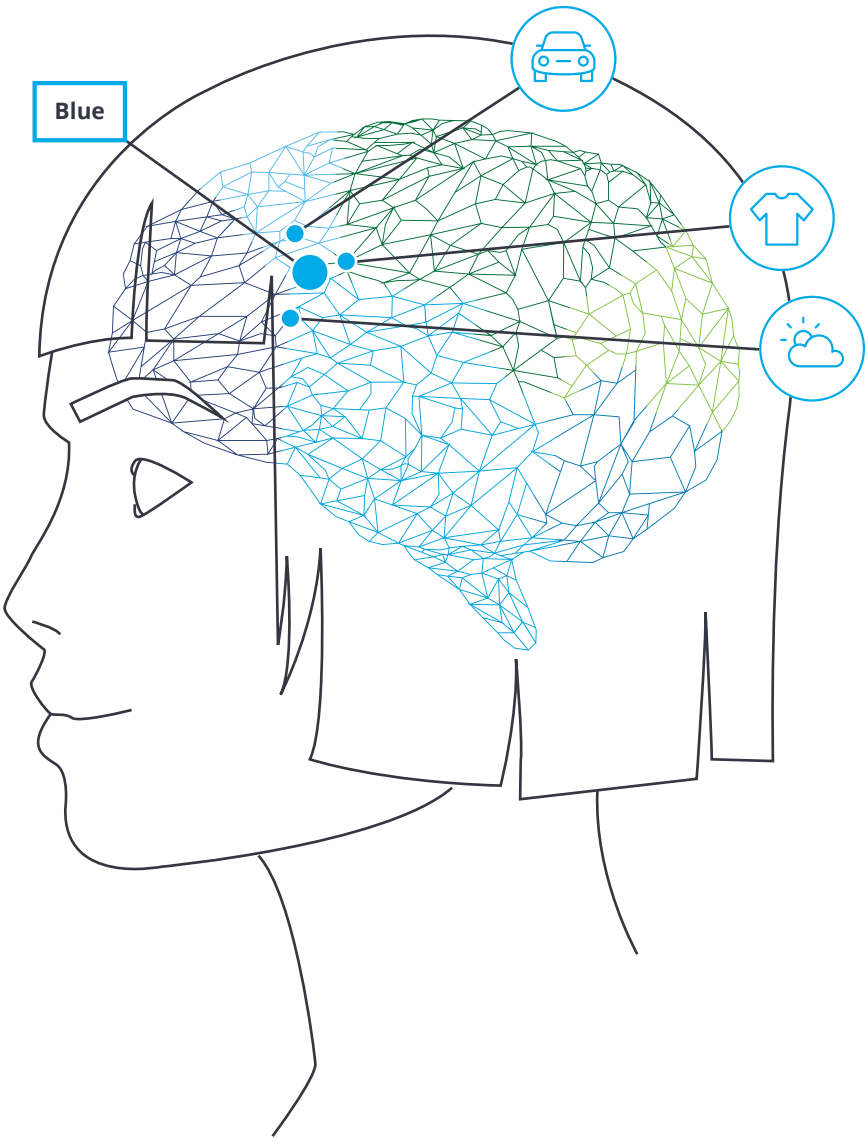

A Knowledge Graph combines sophisticated AI and graph technologies. It is a complex network that accesses data from databases, silos and unstructured sources, bringing it into context. A Knowledge Graph connects different domain knowledge models without altering them and evolves over time with no need to resstructure the models. A Knowledge Graph is able to store complex models and interdependencies that exist in the real world (process instructions of a chemical productions process, configuration rules of a complex mechatronic product, etc.).

Knowledge Graphs connect, structure and simplify knowledge within organizations, creating transparency by bridging data silos.

The foundation for Knowledge Graphs

The creation and management of Knowledge Graphs requires 3 components: capturing the business logic and intelligence of domain experts, connection of domain-specific business logics and their representation in a machine-readable form and storage of the knowledge graph in a semantic graph database.

Component 1: Semantics
capturing Intelligence by connecting information.
Semantic description is the process by which we enrich data with machine-readable meaning by creating the context for an object. Rules and relationships between data define its context.

Cpmponent 2: Ontology
Methodology of representing intelligence in a machine-readable way.
An ontology is a formal representation of knowledge within a particular domain that is created by defining its terms and concepts and enriching them semantically – a specification of a conceptualization.

Compnent 3: Graph Database
Technology to store and manage the intelligence.
Graph databases enable storage, management and analysis of highly connected data made up of objects and the relationships between them (nodes and edges). Graph databases store Information in directly connected objects instead of tables and thus do not require resource-intensive JOIN-operations.

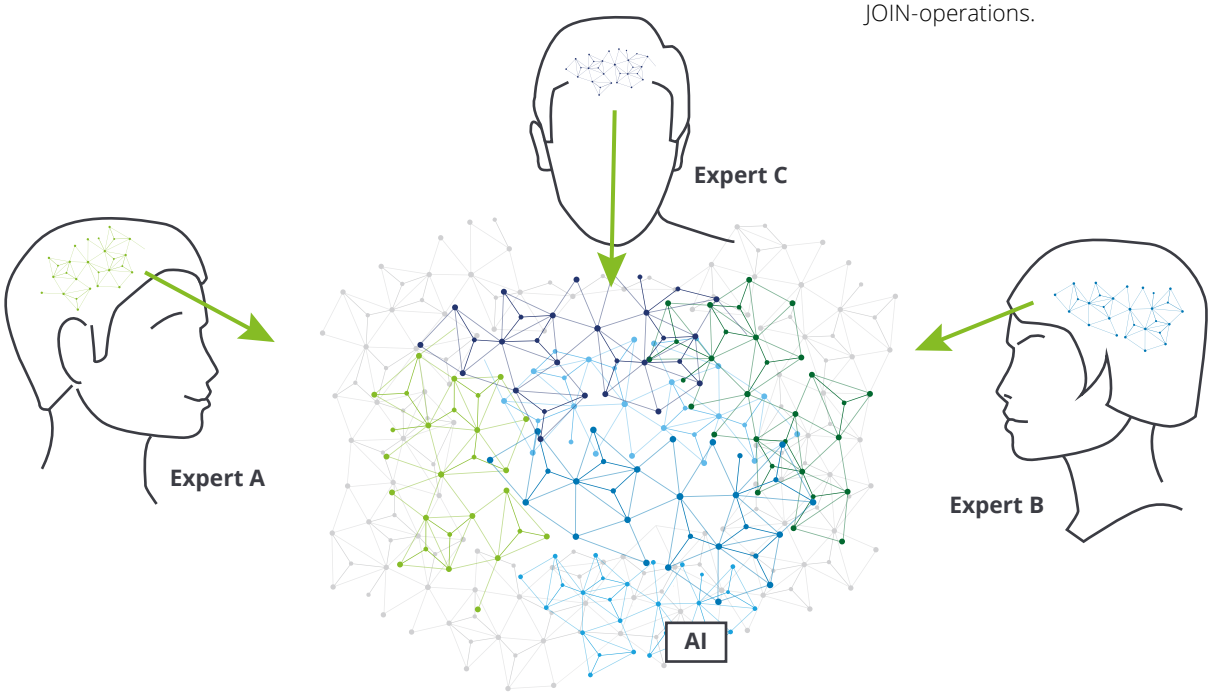

Conclusion

Knowledge is every company's most valuable asset, but it remains hard to grasp at the same time, because it is scattered across different systems and human minds. The key to integrating knowledge efficiently among systems and human users is to provide representation in machine-readable form.

Creating a Knowledge Graph with semantic description of information context allows users to access a machine-readable representation of complex interdependencies that form a real-world model of the knowledge domain.

Knowledge Graphs integrate data across the whole enterprise, support complex decisions and reveal the origin of every correlation within milliseconds. It forms the foundation for AI technologies that are more intelligent than artificial.

The Enterprise Knowledge Graph is a close-to-reality data model that contains the business logic of an enterprise and is stored in a graph database.

Contact

Philipp Obenland
Director | Strategy & Operations
Phone: +49 (0)89 29036 7822
pobenland@deloitte.de

Ulrich Schoof
Senior Manager | Strategy & Operations
Phone: +49 (0)151 5807 0972
uschoof@deloitte.de

Boris Shalumov
Senior Consultant | Strategy & Operations
Phone: +49 (0)151 5807 4474
bshalumov@deloitte.de

Live Enterprise – a paradigm change for companies

Dynamic leadership based on real time data: this approach realizes the potential of new digital technologies.

Welcome in the Here and Now! Digitalization makes this possible. Today more and especially more actual data are available than ever before. Thus, companies are gaining more precise fundamentals for their actions.

They become a Live Enterprise: The Internet of Things is the new data supplier and allows the generation of consistent data pools, which enable continuous evaluations in near real time with innovative analytical methods and Machine Learning. The time lag between analyzing the current state and implementing it into concrete action shrinks potentially towards zero. The result: Relevant data, adaptive processes, agile decision making – a new form of leadership. Live Enterprise is more than just an IT project.

This publication is one of our building blocks how to structure and connect data on the journey to a live enterprise.

Deloitte.

Issue 07/2019