

With the rise of natural language processing and artificial intelligence, computers can now speak our language. In addition to impacting our lives as consumers (when was the last time you consulted a paper map or tried a restaurant without reading reviews?), these new technologies have fundamentally changed how marketers must approach their craft.

In a world where consumers are asking more and more questions on their path to purchase, brands need to be able to understand those questions and answer them accordingly. But that's not as easy as it sounds. Spreadsheets, relational databases, and even traditional content management systems can't answer complex questions, because they don't understand how multiple data points relate to each other.

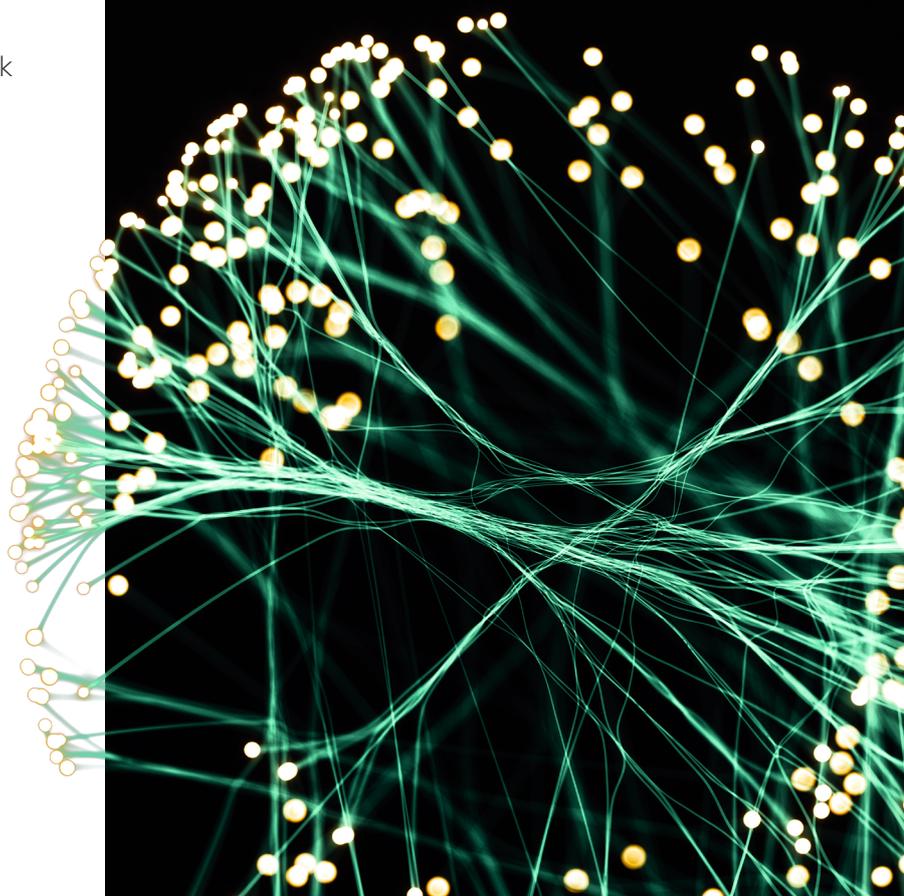
How are you, the brand, going to market, sell, and support your customers in this new world? To understand the complex questions your customers ask and deliver them an experience that exceeds expectations, you'll need an entirely new computing model. You'll need technology that's **semantic**, that understands language, at the center of your marketing stack.

Say hello to the Knowledge Graph.

What is a knowledge graph?

A knowledge graph is a brain-like database that's optimised for answering questions.

It organises data into a graph, which means it stores both data points themselves, and the multiple relationships between data points. With this method of storage, you can look up how data points are linked or related to each other. This mirrors how the human brain works – our brains store and retrieve "data" that is linked together, so you instantly understand the context and connections behind different bits of information.



Why is a knowledge graph important to my marketing?

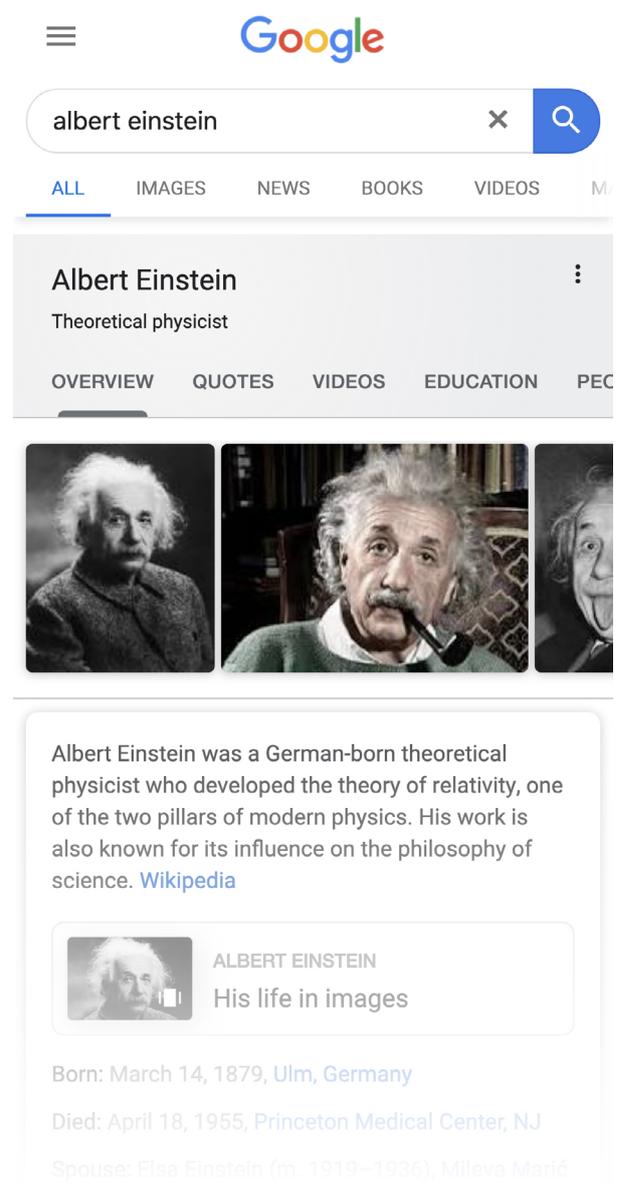
Today, the customer journey starts with a question. And every time you search on Google or ask Alexa a question, you interact with a knowledge graph.

Google first unveiled its Knowledge Graph in 2012. By 2016, the search giant returned structured answers in the form of Knowledge Cards to roughly 1/3 of all queries monthly, all powered by a knowledge graph that contained 70 billion connected facts.

Google built its knowledge graph because people had begun asking questions as they searched, not just typing in simple keywords. In turn, as they received direct answers powered by the knowledge graph, more people asked *more complex* questions. So began the positive feedback cycle that continues today: the better answers people receive, the more complex questions they ask. And the more complex questions they ask, the better answers they expect!

But we all know that consumers don't only have expectations for their experience on Google – they also have the same expectations for *your* brand. When they ask a question about your brand, they want accurate information, no matter where they ask it. And, why can't they ask *you* a complex question and receive just as good of a response as they're used to on the world's biggest search engines?

To win in this new world, just like Google, Facebook, Apple, and Amazon Alexa all use a Knowledge Graph, every company must deploy **their own Knowledge Graph** to answer all the questions anyone could ask about them.



You – and your customers – interact with knowledge graphs every day. Here's an example of a Knowledge Card powered by data from Google's Knowledge Graph.

How is a knowledge graph different from what I have today?

There are **five qualities** that make knowledge graphs special and differentiate them from other types of databases. Their graph data structure is the first.

1 Graph Data Structure

Say a consumer is looking for which products are on sale at a specific retail location. If that retailer were to store its locations, its products, its special offers, and its in-store events in different tabs in a spreadsheet, it would be difficult for the retailer to find the answer to that question without having a human page through the tabs. But with a knowledge graph, the brand can relate relevant data points to each other: special offers would be *applied* to specific products, which would be *sold* in certain stores, which would *host* certain events. The relationships add context and meaning to the data points themselves, and the brand's technology could then answer the customer's question.

2 Semantic

The second thing that makes knowledge graphs special is that they're semantic.

When your customers search for things like stores, products, offers, or events, they're not searching for those words – they're searching for the real *things* those words refer to.

With a traditional database, words like "store", "provider", or "offer" are simply strings of letters. Knowledge graphs, however, understand what those words *mean*. Because the meaning of data is inscribed inside the data they store, knowledge graphs can understand what your customers are actually looking for in the real world, so they can get them exactly what they need.

Take this example: let's say your customer is looking for the best financial advisor near them, open now, that covers family trusts. With a Knowledge Graph, you can understand that your customer wants a financial advisor, which is a *professional*, who offers the *service* of covering family trusts. You can also tell that your customer wants an *office location* near them, that's *open now*.

Without semantic understanding, a spreadsheet or other database doesn't understand that an advisor is a person, or that a family trust is a service, because it simply stores strings of data. It doesn't interpret or relate data points to each other. But, a knowledge graph can – and since its graph structure understands the relationships among your professionals, locations, products, and hours, it can serve up the best answer to this question without human intervention.

3 Smart

Third, knowledge graphs are smart, because they understand not only that different data points are related, but specifically *how* they relate to each other.

Any database table can contain a *taxonomy*, which is a hierarchical series of definitions and attributes. An auto manufacturer may manage a taxonomy for the cars it produces. That taxonomy would include each “parent” car (“Model X”) and the “child” attributes it contains (Body, Color Options, Transmission Type, and Tires). This type of information is sufficient to answer questions about the attributes of each car (“What color options come standard with Model X?”).

But, what happens if a customer were to ask how an attribute relates to *any* car, like “Which cars work with ABC Tires”? With a relational database, each “child” attribute (e.g., a tire) only exists in relation to its “parent” (the car). Since a relational database’s taxonomy would only consider a tire to be a “child” of a certain car, rather than a first-class citizen in its own right, the database would struggle to answer a question about how that tire relates to *any* parent. **(See Fig. 1 on next page)**

Unlike the parent-child hierarchy of relational databases, knowledge graphs contain an *ontology*, which is a network of interconnected things and how each of them relates to another. The nature of these relationships can change depending on which direction they go – for example, tires *go on* a car, but cars *contain* tires; tires are *sold* at stores, but stores *carry* tires. This is a far more accurate representation of the real world: depending on who’s asking, the same person can be someone’s parent, someone’s child, someone’s employee, or someone’s teacher. **(See Fig. 2 on next page)**

Knowledge graphs are smart because they understand that relationships between things change depending on which way they’re viewed. This allows knowledge graphs to answer questions that are asked from any starting point. In addition, they also understand the *strength* of each relationship between entities. For example, if a product is sold in every store, that relationship is different from the one between a product and the *only* store it’s offered in.



Relational Database

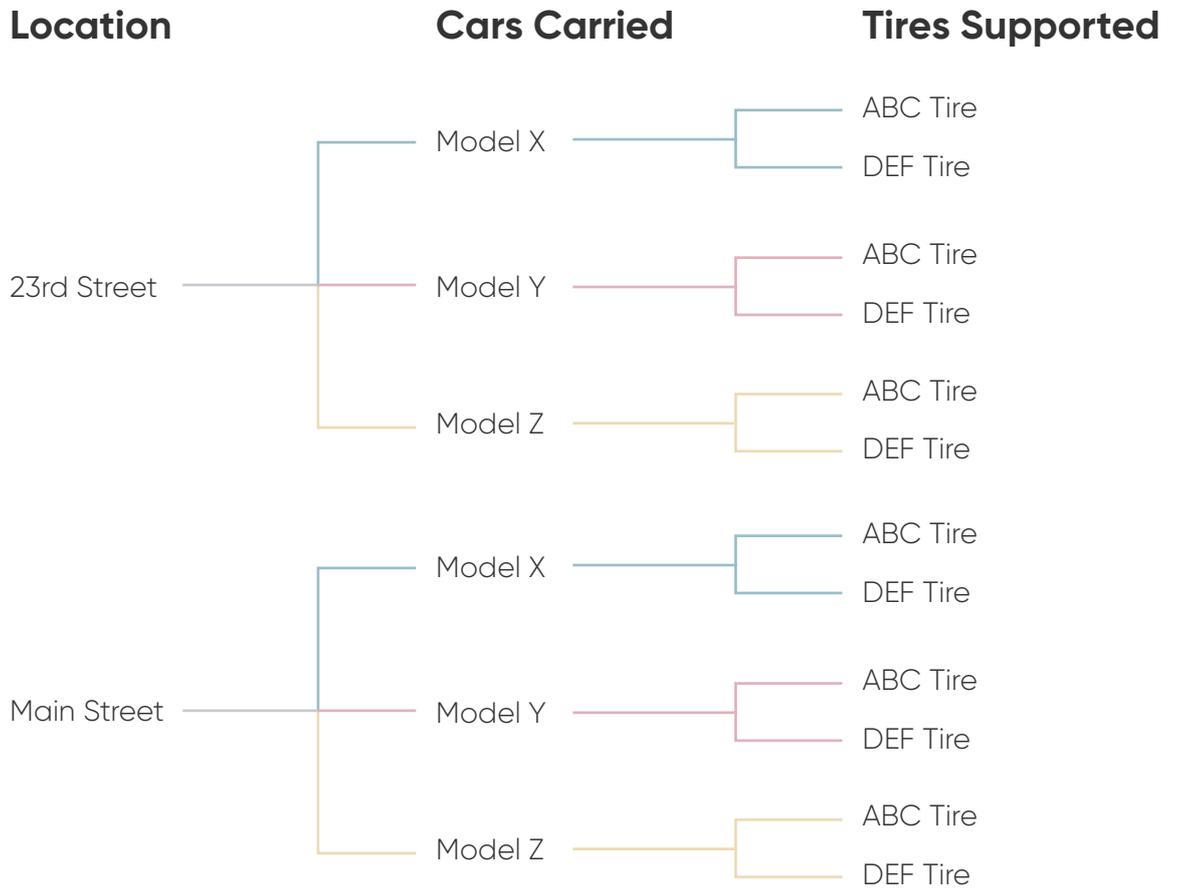


Figure 1

Knowledge Graph

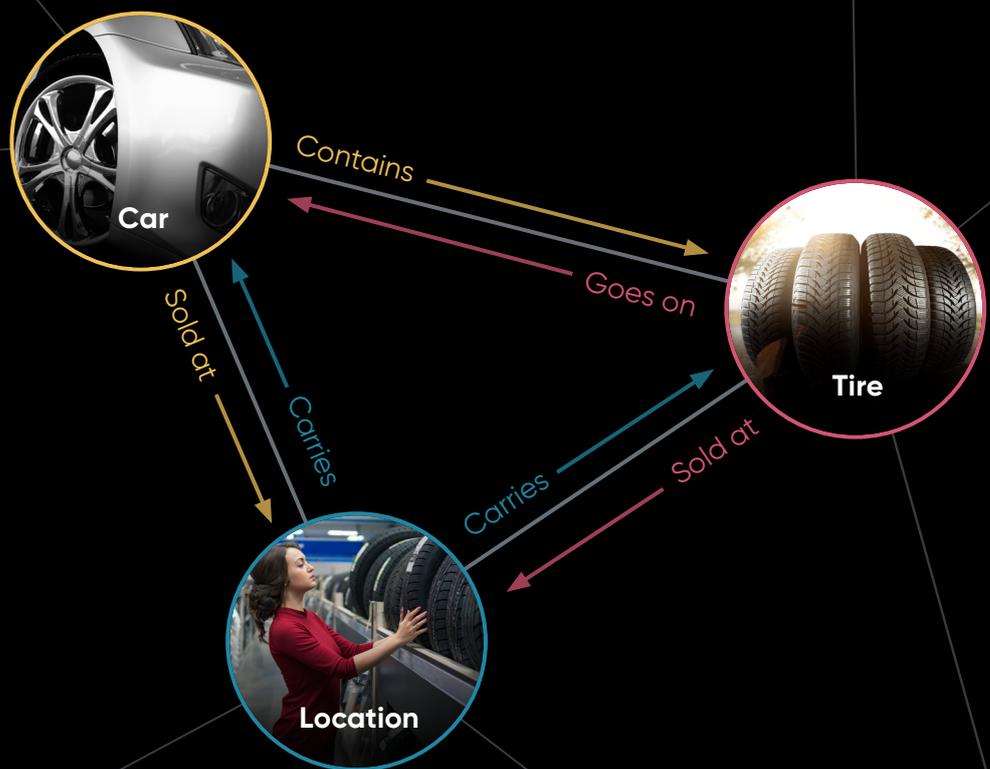


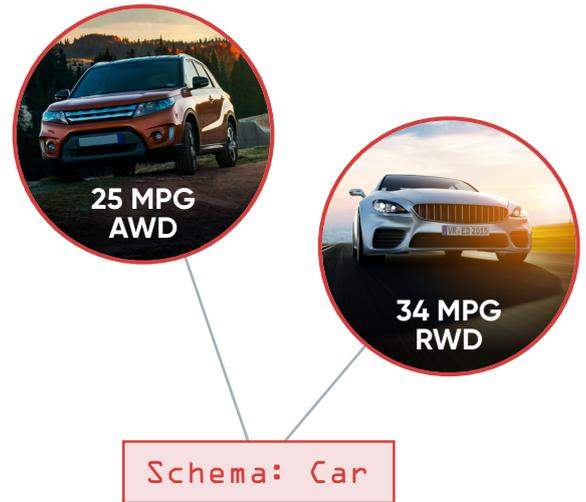
Figure 2

4 Flexible

Next, knowledge graphs are flexible: they're capable of learning and changing their data structure as they grow.

To continue with our example, say the auto manufacturer introduced a new line of electric vehicles (EVs). Since its prior cars were not electric, these new EVs would have a whole new set of attributes the manufacturer had not previously stored, like charge time, outlet type, and battery energy. Additionally, the EVs would *not* include some attributes of non-electric cars, such as fuel type, engine torque, and fuel capacity.

In order to serve up the correct cars in response to consumer questions like "What is the charge time for your cars?", the manufacturer could simply add each new EV to its graph as a new node. It would add schema to that node specifically for Electric Vehicles. It would also go back and apply schema for Non-Electric Vehicles to the nodes for its older cars. Since the schema within a knowledge graph can be changed easily without creating new tables or having to move data between them, adding a new type of data to a knowledge graph is far less time-intensive than doing the same in a relational database.



↓ Add new line of Electric Vehicles ↓



5 **Alive**

Finally, knowledge graphs are alive: they can accommodate new, streaming data points and the new relationships they bring with them. Why does this matter to your business? Think of a piece of information that may interest your customers, like your stock price. It constantly changes based on new bid asks and new sales. Would you expect an employee to manually update that stock price in a relational database every time it changes? Unlikely. So how could you answer a question about it?

Because knowledge graphs can ingest real-time streams of data, they can accommodate constant changes without requiring heavy write transactions each time. As you accumulate more and more data from disparate sources, you can correlate it effectively and refine the connections among different parts without rebuilding the entire structure. In short, knowledge graphs are better equipped to keep up with the pace of your business.

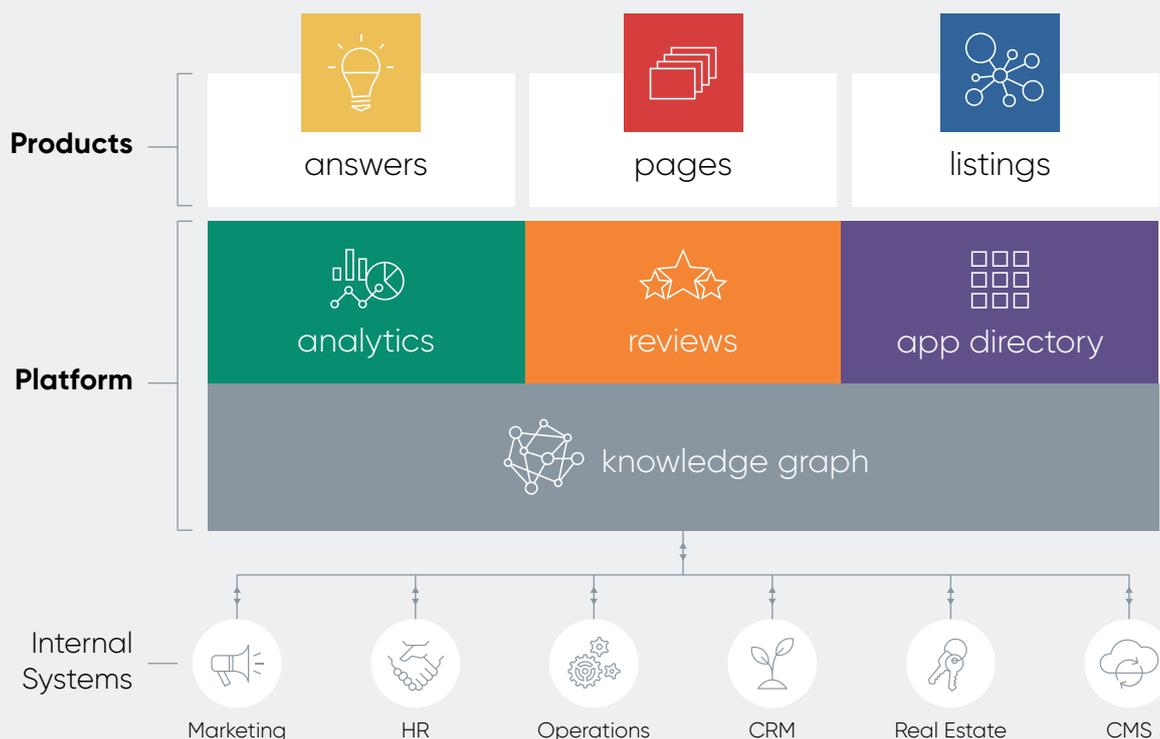
The Yext Search Experience Cloud

With a mission to provide perfect answers everywhere customers search, Yext's Search Experience Cloud puts businesses in control of their facts online. By providing brand-verified answers on their own websites – then across search engines, maps, apps, voice assistants and chatbots – Yext helps brands take back control of the customer journey.

At the foundation of our offering is the Yext Knowledge Graph, the single source of truth for the many public facts about your brand. With the Knowledge Graph, brands can create, connect, and optimise all the facts about themselves to answer their customers' questions. Brands around the world use the Yext Knowledge Graph to manage information from store locations and open jobs, to menu items and their nutritional information, to doctors and their specialties and educational backgrounds.

When they open new stores, host new events, or add new jobs, they simply add those to their Knowledge Graph.

Then, the services on top of the Knowledge Graph let you power answers anywhere a customer could ask for them. Yext Answers delivers a natural-language search experience on your website and other owned properties, where customers can search your Knowledge Graph and get direct answers. Yext Pages helps you create dynamic landing pages for your stores, events, doctors, ATMs, or any other facts inside your Knowledge Graph to power rich experiences in search results. And, Yext Listings integrates your Knowledge Graph directly with more than 150 voice assistants, mapping apps, and other third-party services, like Google Maps, Amazon Alexa, Bing, Facebook, and many more.





The customer journey starts with a question, and every day consumers search for answers about brands. However, they are increasingly served false or misleading information from sources other than the brand. Yext, the Search Experience Cloud company, exists to help brands regain and maintain a direct relationship with their customers. With a mission to provide perfect answers everywhere, Yext puts businesses in control of their facts online by delivering brand verified answers straight from the source – wherever their customers are searching. Companies like Tesco, Marriott, Jaguar-Land Rover, and businesses around the world use Yext to take back control of the customer journey, starting on their own website.

Yext has been named a Best Place to Work by Fortune and Great Place to Work[®], as well as a Best Workplace for Women. Yext is headquartered in New York City with offices in Amsterdam, Berlin, Chicago, Dallas, Geneva, London, Miami, Milan, Paris, San Francisco, Shanghai, Tokyo, and the Washington, D.C. area.