

**Invited intro for “Knowledge Graphs in Use - The What, How and Why”, by Dieter Fensel, Kevin Angele, Elwin Huaman, Elias Kärle, Oleksandra Panasiuk, Umut Şimşek, Ioan Toma, Jürgen Umbrich, and Alexander Wahler, Springer, 2020**

In 2005, Manuel Sahli became an active contributor to Wikipedia. He was particularly interested in his hometown, the beautiful Swiss city of Winterthur. Over the years, he edited and meticulously maintained the article on Winterthur. One of the sections he worked on was the list of sister cities.

Around the same time, a contributor added Winterthur to the list of sister cities on the article on Ontario, California, a city about an hour east of Los Angeles. Even though sister cities are supposed to be reciprocal, the information never made it to the article on Winterthur.

Manuel, although interested in all things Winterthur, did not learn about this addition for almost a decade.

2012 saw the launch of Wikidata, a sister project to Wikipedia. It was a knowledge graph that anyone could edit and use. As time went by, a contributor added the list of sister cities of Ontario to Wikidata, using the Ontario article on Wikipedia.

Manuel Sahli had since been elected to the cantonal parliament of Winterthur. Using Wikidata, he wanted to make a map of Winterthur's sister cities - and Wikidata has a powerful query interface that allows to create ad-hoc visualizations of query results over its knowledge graph. To Manuel's big surprise, the map contained one place he was not familiar with yet: Ontario.

At first, Manuel thought that this was an act of vandalism that has remained undiscovered for a while, but the fact came with a source: the official government Website of Ontario. And when he went to that site, there he found Winterthur! As a member of the cantonal parliament he asked the Winterthur city administration and archives if they knew anything about that claim. Both denied.

Manuel wrote an official letter to the city of Ontario to clarify the situation - and lo and behold, they delivered documents, signed by both sides, establishing the sister city relation in 1982. Given these documents and the precise dates, the Winterthur city archives were now able to find the relevant documents. Thanks to Wikidata, Winterthur discovered its long lost fifth sister city, and the two cities have since rekindled their relationship.

\*

No one was trying to hide that fact from Manuel, or the world. The city archives had the documents (but couldn't find them at first), Wikipedia was publishing the fact for almost a decade, it was on the official Ontario Website - and yet, Manuel, although deeply invested in his home town and Wikipedia, did not learn about it. It was only when it was added to a knowledge graph, one that Manuel had access to and knew how to query, that the fact was surfaced.

It is far too easy to bury knowledge in documents and in heaps of natural language, and very hard to surface it at the right time. A knowledge graph makes facts easier to index, process, and find.

But that is only half the story. The other half is that it was sheer luck that Manuel was not only interested enough to discover this fact, but also capable of querying Wikidata using the SPARQL query language. Most people do not know how to do that.

\*

We are witnessing a paradigm shift in how we use computers, as momentous as the introduction of time-sharing systems, graphical user interfaces, the Web, and smartphones before: the nascent ability to interact with computers through natural language. Conversational interfaces have been around for a long time - just as any novel technology has been preceded by various adoptions in niche areas. But so-called intelligent assistants are quickly spreading and evolving, and more and more devices, from wristwatches to cars, from televisions to earbuds, are being equipped with the capability to listen to your commands and questions, and to answer them.

The number of people who will gain access to knowledge through intelligent assistants for their daily tasks is growing remarkably fast. In order to find answers to burning questions, whether trivial or life-changing, we will not need the traditional computer in the office or school or library anymore. No one will need to learn how to use the newest smartphone features, how to use a mouse, or even how to read and type, thus widely opening the door to a much more inclusive world.

This new paradigm needs new ways to allow people to join the new data space, to give access to their data and their services to their prospective clients and consumers, to allow people to discover and enjoy what they have to offer. In a world, where more and more people do not use apps and websites to book a hotel room, rent a car, buy a flight ticket, order their lunch, and ask for information, we need to make sure to understand the users' intents and offer our data and services in ways that allow their integration and composition.

This is not a vision of a future to come, but the world we already built. Tens of thousands of capabilities and services are already integrated into the Google Assistant, Apple Siri, Amazon Alexa, the open source Mycroft, Microsoft Cortana, Samsung Bixby, and others.

Bringing together the natural language technology required to understand users' queries and the power of structuring facts in large knowledge graphs will allow to bring knowledge to the end user when they need it. People will become empowered to explore and ask about the world, to follow their curiosity, to learn. Businesses will be able to expand the reach of their services and products to a large number of new customers.

\*

This short booklet offers a hands-on overview of the architecture and the steps to join the world of intelligent assistants. It guides us through building, growing, and maintaining our own knowledge graph. It discusses how knowledge graphs represent the facts that are important in your organization, and how to make them available to other providers.

Dieter Fensel was one of the people who have very early recognized the potential of semantic technologies and the need to allow for the composition of services into novel interfaces. He and his research groups were leading the field in the semantic description of services, in linked data applications, and in ontology engineering and learning methods and tools. This book by him and the researchers of his team is distilling the experience of many years.

Enjoy going through the book! It will help you to better understand the novel field of knowledge graphs and how to use them. It will allow you and your organization to move to a world where the knowledge you already have in your archives and systems does not remain so stubbornly hidden, and where you and your customers both will be enabled to work with the knowledge you already have, so that we can empower users, discover new sources of revenue, and never lose a sister city again.

Denny Vrandečić, San Francisco  
June 2019