

BLOG BOX

Google just got a whole lot smarter, launches its knowledge graph

Today, Google is launching one of its most ambitious and interesting updates to its search engine in recent months. Starting in a few days, you will start to see large panels with additional factual information about the topic you were searching for take over the right side of Google's search result pages. The panels are powered by what Google calls its new "Knowledge Graph" and they will serve two different functions. Google will use this space to show you a summary of relevant information about your queries (think biographical data about celebrities and historical figures, tour dates for artists, information about books, works of art, buildings, animals etc.), as well as a list of related topics. In addition, Google will now allow you to clarify what exactly you are looking for and will use these boxes for disambiguation. Thanks to this, you will soon be able to tell Google you were looking for the L.A. Kings ice hockey team and not the Sacramento Kings when you searched for 'kings.'

The company has actually been working on the semantic technology that drives this knowledge graph for quite a few years. This specific project, Google told me earlier this week, has been in the works for about the last two years. During this time, the company has been working hard on creating the vast database of structured knowledge that powers the features it is launching today (through Google's acquisition of Freebase). Today, the knowledge graph database holds information about 500 million people, places and things. More importantly, though, it also indexes over 3.5 billion defining attributes and connections between these items.

As Google Fellow Ben Gomes told me yesterday, the company really wants to move beyond just understanding the characters you are typing into its search engine to getting a better understanding of what it is you are really looking for ("strings to thing" is what Gomes likes to call it). To do this, Google is using both its own and other freely available sources like Wikipedia, the World CIA Factbook, its own Freebase product, Google Books, online event listings and other data it crawls, but it is also using some commercial datasets (though Google wouldn't reveal which companies specifically it is working with here).

Looking ahead

Google, of course, has been adding bits and pieces of semantic search smarts to its search engine over the last few years (and so has Microsoft after its acquisition of Powerset). With Google Squared, one of its recently shelved experiments, it also once launched a pretty ambitious project to understand information on the web and then display it in a table (some of this technology likely lives on in the Knowledge Graph now). Today's launch, however, represents Google's most ambitious move in this direction.

As Gomes has told me, now that Google's algorithms have access to this structured data and can understand it better, the next step will be to understand more complex questions like "Where can I attend a Lady Gaga concert in warm outdoor weather?" For now, though, it is worth noting that this update isn't about natural language processing and answering questions so much as about displaying relevant data.

It'll be interesting to see how this new feature will influence how people search and what links they click on. I wouldn't be surprised if this had quite a negative influence on traffic to Wikipedia, for example. At the same time, though, the disambiguation feature may just help drive more relevant traffic to the sites Google links to as well.

Source: F. Lardinois, 'Google just got a whole lot smarter, launches its knowledge graph', TechCrunch.com, 16 May 2012.