INQUIRY · INQUIRY IN PRACTICE

Search Queries as a Record of Public Inquiry

Billions of search queries form a running record of what people want to know. This article explains how researchers have used search data to track illness, unemployment, markets and public attention, what went wrong with some celebrated predictions, and the methodological and ethical rules for using search data well.

ModQu research article · 2,281 words · 24 references (APA 7th edition)

Every day, people type billions of questions into search engines. They ask about symptoms, jobs, products, news, prices, laws and words they have just heard. They ask in private, in their own words, at the moment the question occurs to them. Taken together, these queries form the largest record of human curiosity ever assembled. Researchers in public health, economics, political science and marketing have learned to read it, with some remarkable successes and some instructive failures. This article explains what search queries reveal about public inquiry, how they have been used as evidence and what care their use requires.

A query as a question

Robert Taylor (1968) described how a person's information need passes through stages, from an unexpressed sense that something is missing to the compromised form in which the question is finally put to an information system. A search query is the latest form of that compromised question. It is short and often ungrammatical, but it records a genuine need at a specific time and place. Nicholas Belkin (1980) emphasised that people search because they are aware that their knowledge is incomplete, often without being able to specify exactly what they lack. Search logs therefore capture not only what people know they want but their attempts to find out what they need.

Early studies of search logs established basic patterns. Silverstein et al. (1999) analysed nearly a billion queries submitted to one search engine and found that most were very short, most search sessions consisted of a single query, and most users looked only at the first page of results. Jansen et al. (2000) reported similar patterns. Andrei Broder (2002) classified queries by their purpose: navigational, to reach a particular site; informational, to learn something; and transactional, to do something. Only informational queries are questions in the ordinary sense, but all three record intentions. Even a navigational query tells us that someone wanted to reach a particular organisation at a particular moment, and the volume of such queries over time is a rough measure of public interest in that organisation.

The language of questions

Search queries also show how people phrase their questions, which is often quite different from how experts and organisations phrase the same topics. A medical guideline may refer to "gastro-oesophageal reflux"; the public searches for "burning in chest after eating". A bank may describe a "balance transfer facility"; its customers search for "how to

move credit card debt". The gap matters, because information written in the vocabulary of experts may never reach the people who need it, and research questions written in that vocabulary may not be understood by the people answering them.

Queries also reveal the *form* of public questions. Some are framed as "how to" questions, seeking a procedure. Some ask "why", seeking an explanation. Some ask "is it safe", "is it worth it" or "is it normal", seeking reassurance or a judgment. Others compare: "X or Y", "X versus Y". The mix of question forms for a topic indicates what stage of understanding people have reached and what kind of answer they want. A topic dominated by "what is" queries is still unfamiliar to most people; one dominated by comparisons is one in which people are choosing.

Search data as evidence

Figure 1 traces some milestones in the use of search data for research.

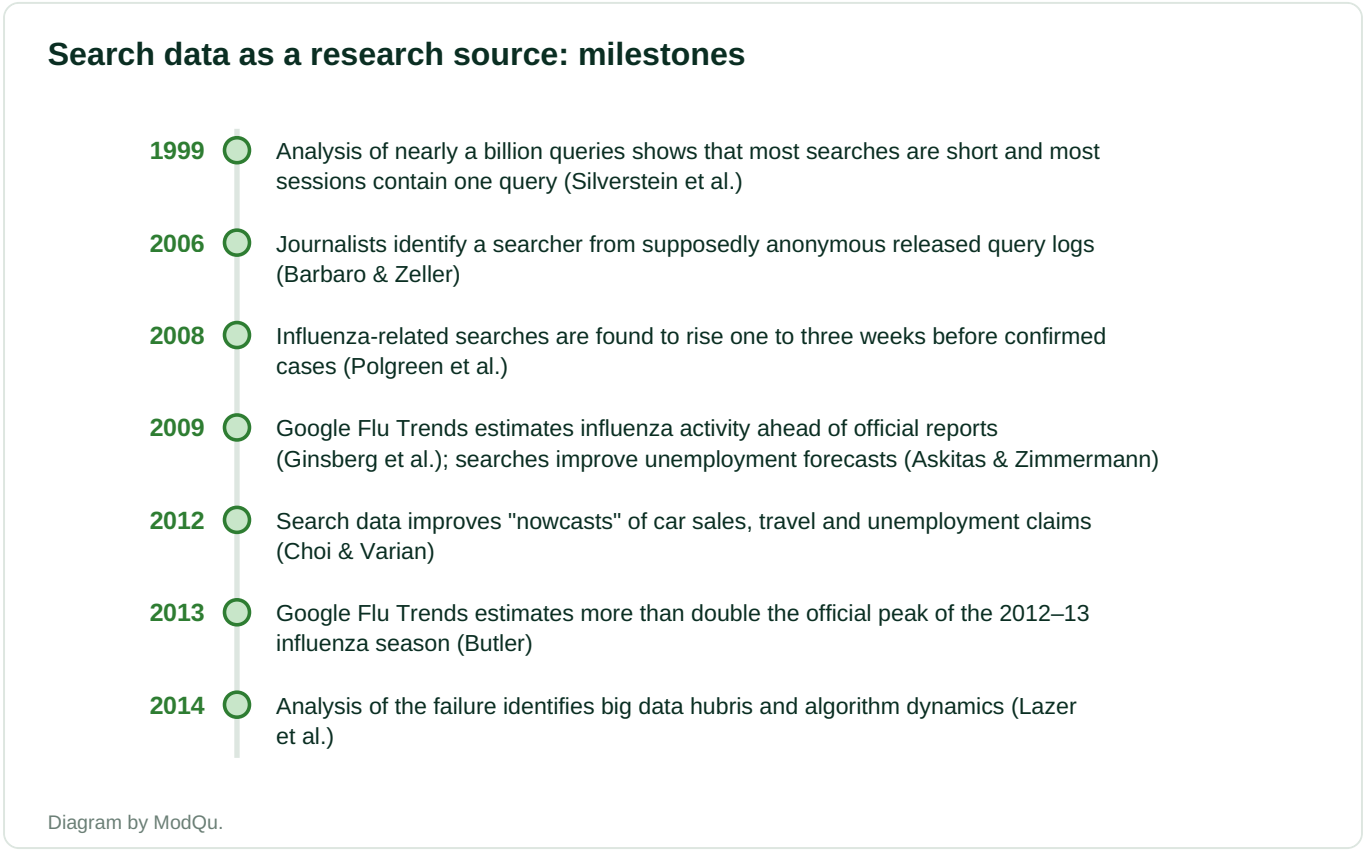


Figure 1. Milestones in research using search data. Diagram by ModQu.

Health

Public health researchers were among the first to see the potential. Polgreen et al. (2008) found that searches for influenza-related terms on one search engine rose one to three weeks before increases in laboratory-confirmed cases. Ginsberg et al. (2009) built a model from queries on Google and reported that it could estimate levels of influenza-like illness in regions of the United States about one to two weeks ahead of official surveillance reports. The resulting service, Google Flu Trends, attracted wide attention as an example of how big data could transform public health.

Search data also shows how people inquire about their own health. Fox and Duggan (2013) reported from a national survey that about 72% of American internet users had looked online for health information in the previous year, and that about a third of adults had used the internet to try to work out what medical condition they or someone else might have. Ryen White and Eric Horvitz (2009) studied search logs and found that searches about common symptoms frequently escalated to queries about serious but rare conditions, a pattern they called *cyberchondria*. Paparrizos et al. (2016) went further, showing that patterns in the earlier searches of people who later searched for information about a diagnosis of pancreatic cancer could, in some cases, be distinguished from those of other users before the diagnosis was made, raising the possibility, and the ethical difficulties, of screening through search behaviour.

Economics

Economists found that searches often precede economic activity. Nikos Askitas and Klaus Zimmermann (2009) showed that searches for terms related to unemployment and job seeking improved forecasts of unemployment in Germany. Hyunyoung Choi and Hal Varian (2012) demonstrated that search data could improve short-term "nowcasts" of indicators such as car sales, unemployment claims, travel and consumer confidence, because people search for information about these activities as they prepare to undertake them. Preis et al. (2013) found that increases in searches for financial terms tended to precede falls in stock market prices, suggesting that search data captured investor concern before it showed up in trading.

Goel et al. (2010) offered an important corrective. They found that search volumes predicted the opening box-office revenue of films, first-month sales of video games and the chart rankings of songs, but that in many cases the improvement over simple models using other publicly available data was modest. Search data adds information, but not always much.

Politics and society

Because searches are private, they may reveal attitudes that people would not express to a survey interviewer. Seth Stephens-Davidowitz (2014) used the regional frequency of searches containing racial slurs as a measure of racial animus and found that areas with more such searches gave Barack Obama a smaller share of the vote than would otherwise have been expected, estimating a substantial cost to his vote share in 2008. In a later book, Stephens-Davidowitz (2017) argued that search data provides a uniquely candid view of human concerns, because people tell search engines things they tell no one else.

Political scientists have also used search data to measure public attention. Joseph Ripberger (2011) found that search trends tracked other measures of attention to public issues and proposed them as a measure of public attentiveness. Jonathan Mellon (2013) tested search data against survey measures of the issues people consider most important and

found that it was a valid measure for some issues but not others, depending on whether searching is a natural response to concern about the issue. Christopher Scheitle (2011) reached a similarly cautious conclusion: search data can complement other social research data, but researchers must validate it for each use.

When search data misleads

The most famous failure in the field involved its most famous success. In the 2012–13 influenza season, Google Flu Trends estimated a peak level of influenza-like illness more than double that reported by official surveillance (Butler, 2013). Lazer et al. (2014) analysed the failure and found that the model had overestimated influenza prevalence in most weeks over the preceding two years. They identified two problems. The first was *big data hubris*: the assumption that large volumes of data could replace, rather than supplement, traditional measurement. The model had been built by selecting, from millions of search terms, those that happened to correlate with past influenza data, and some reflected seasonal patterns unrelated to illness. The second was *algorithm dynamics*: changes in the search engine itself, such as suggested searches, altered what people searched for, so that the relationship between searches and illness shifted over time.

The episode illustrates general problems with search data, summarised in Figure 2. Danah Boyd and Kate Crawford (2012) warned that "bigger data are not always better data", that data taken out of context loses its meaning, and that data being accessible does not make its use ethical. Matthew Salganik (2018) noted that digital trace data are typically big, always on and non-reactive, but also incomplete, unrepresentative, drifting over time and *algorithmically confounded*: shaped by the design of the systems that record them.

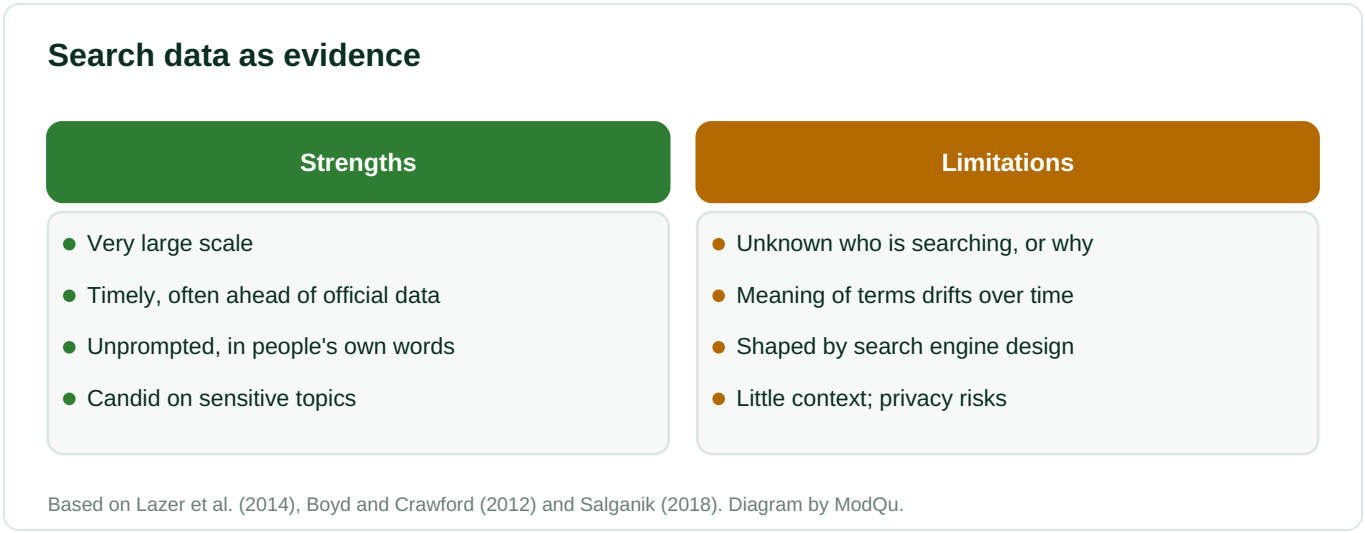


Figure 2. Strengths and limitations of search data as evidence. Based on Lazer et al. (2014), Boyd and Crawford (2012) and Salganik (2018). Diagram by ModQu.

Search data also raises a basic question of representation. Not everyone searches, and those who do are not a random sample of the population. Older people, people on low incomes and people in areas with poor internet access may be under-represented, and people differ in how readily they turn to a search engine rather than to a friend, a professional or

a printed source. A rise in searches may therefore reflect a change in who is searching as much as a change in what people want to know. Without independent information about the searchers, which search data does not provide, such effects are hard to detect.

Privacy and ethics

Search queries are intimate. In 2006, an internet company released a large set of search queries for research, with user names replaced by numbers. Journalists at *The New York Times* were able to identify one of the users from the content of her searches alone (Barbaro & Zeller, 2006). The incident showed that search histories can identify the people who made them even when names have been removed. Researchers using search data now generally work with aggregated figures, such as the relative volume of searches for a term in a region, rather than individual records, and they must consider whether their findings could harm the groups they describe.

Using search data well

The research suggests several rules for anyone using search data as evidence of public inquiry:

- **Validate against other measures.** Check that search trends for a topic correspond to an independent measure before relying on them, and recheck periodically.
- **Use search data to complement, not replace.** Combine it with surveys, sales data or official statistics, which can show who is asking and why.
- **Choose terms with care.** Select search terms on the basis of meaning, not only correlation, to avoid spurious matches.
- **Watch for drift.** Changes in search engines, language and behaviour can alter what a term measures over time.
- **Protect privacy.** Work with aggregated data and avoid analyses that could identify individuals or stigmatise groups.

Search data in organisations

Businesses and public bodies increasingly use aggregated search data alongside their own research. Used carefully, it can show when interest in a topic rises and falls through the year, which regions show the most interest, which related questions people ask and which terms they use. It can provide early warning of an emerging concern, such as a rise in questions about a product fault or a change in regulations, before the concern appears in complaints or sales.

Its limits in organisational use are the same as in academic research. Search data cannot say who the searchers were, whether they were customers, what they concluded or what they did next. A rise in searches for a competitor's product may reflect growing interest, a news story or a problem with the product. For these reasons, organisations get most value from search data when they treat it as a way of finding questions worth investigating and then investigate them with methods that can answer them, such as surveys of defined populations, interviews and analysis of their own customer records. The search data shows what the public is asking; the follow-up research shows who is asking, how many and why.

What search queries reveal

Used carefully, search data offers something no other source does: a continuous, unprompted record of what people want to know, in their own words, at the moment they want to know it. It shows when a question first appears, how quickly it spreads and when it fades. It reveals the vocabulary people use for their concerns, which often differs from the language of experts and organisations. And it captures questions people are reluctant to ask aloud.

Consider a hypothetical example. A public health agency notices a sharp rise in searches about a particular side effect of a common medicine in one region. Search data alone cannot say whether the side effect has become more common, whether a news story has raised concern or whether a local rumour is circulating. But it tells the agency where to look and what question to ask. A rapid survey of the region, combined with a review of reports from pharmacists, shows that a social media post has caused alarm. The agency responds with clear information, phrased in the words people were searching for. Within days, searches on the topic begin to fall, and a follow-up survey finds that most people in the region who had heard the rumour now know it was unfounded. The search data did not answer the agency's question, but it showed the agency which question to ask, where and when.

Search queries do not answer questions; they record them. Their value lies in showing which questions the public is asking, so that research, communication and services can be directed towards answering them.

References

Askitas, N., & Zimmermann, K. F. (2009). Google econometrics and unemployment forecasting. *Applied Economics Quarterly, 55*(2), 107–120. https://doi.org/10.3790/aeq.55.2.107

Barbaro, M., & Zeller, T., Jr. (2006, August 9). A face is exposed for AOL searcher No. 4417749. *The New York Times*.

Belkin, N. J. (1980). Anomalous states of knowledge as a basis for information retrieval. *Canadian Journal of Information Science, 5*, 133–143.

Boyd, D., & Crawford, K. (2012). Critical questions for big data: Provocations for a cultural, technological, and scholarly phenomenon. *Information, Communication & Society, 15*(5), 662–679. https://doi.org/10.1080/1369118X.2012.678878

Broder, A. (2002). A taxonomy of web search. *ACM SIGIR Forum, 36*(2), 3–10. https://doi.org/10.1145/792550.792552

Butler, D. (2013). When Google got flu wrong. *Nature, 494*(7436), 155–156. https://doi.org/10.1038/494155a

Choi, H., & Varian, H. (2012). Predicting the present with Google Trends. *Economic Record, 88*(s1), 2–9. https://doi.org/10.1111/j.1475-4932.2012.00809.x

Fox, S., & Duggan, M. (2013). *Health online 2013*. Pew Research Center.

Ginsberg, J., Mohebbi, M. H., Patel, R. S., Brammer, L., Smolinski, M. S., & Brilliant, L. (2009). Detecting influenza epidemics using search engine query data. *Nature*, *457*(7232), 1012–1014. https://doi.org/10.1038/nature07634

Goel, S., Hofman, J. M., Lahaie, S., Pennock, D. M., & Watts, D. J. (2010). Predicting consumer behavior with Web search. *Proceedings of the National Academy of Sciences*, *107*(41), 17486–17490. https://doi.org/10.1073/pnas.1005962107

Jansen, B. J., Spink, A., & Saracevic, T. (2000). Real life, real users, and real needs: A study and analysis of user queries on the web. *Information Processing & Management*, *36*(2), 207–227. https://doi.org/10.1016/S0306-4573(99)00056-4

Lazer, D., Kennedy, R., King, G., & Vespignani, A. (2014). The parable of Google Flu: Traps in big data analysis. *Science*, *343*(6176), 1203–1205. https://doi.org/10.1126/science.1248506

Mellon, J. (2013). Where and when can we use Google Trends to measure issue salience? *PS: Political Science & Politics*, *46*(2), 280–290. https://doi.org/10.1017/S1049096513000279

Paparrizos, J., White, R. W., & Horvitz, E. (2016). Screening for pancreatic adenocarcinoma using signals from web search logs: Feasibility study and results. *Journal of Oncology Practice*, *12*(8), 737–744. https://doi.org/10.1200/JOP.2015.010504

Polgreen, P. M., Chen, Y., Pennock, D. M., & Nelson, F. D. (2008). Using internet searches for influenza surveillance. *Clinical Infectious Diseases*, *47*(11), 1443–1448. https://doi.org/10.1086/593098

Preis, T., Moat, H. S., & Stanley, H. E. (2013). Quantifying trading behavior in financial markets using Google Trends. *Scientific Reports*, *3*, Article 1684. https://doi.org/10.1038/srep01684

Ripberger, J. T. (2011). Capturing curiosity: Using internet search trends to measure public attentiveness. *Policy Studies Journal*, *39*(2), 239–259. https://doi.org/10.1111/j.1541-0072.2011.00406.x

Salganik, M. J. (2018). *Bit by bit: Social research in the digital age*. Princeton University Press.

Scheitle, C. P. (2011). Google's Insights for Search: A note evaluating the use of search engine data in social research. *Social Science Quarterly*, *92*(1), 285–295. https://doi.org/10.1111/j.1540-6237.2011.00768.x

Silverstein, C., Marais, H., Henzinger, M., & Moricz, M. (1999). Analysis of a very large web search engine query log. *ACM SIGIR Forum*, *33*(1), 6–12. https://doi.org/10.1145/331403.331405

Stephens-Davidowitz, S. (2014). The cost of racial animus on a black candidate: Evidence using Google search data. *Journal of Public Economics*, *118*, 26–40. https://doi.org/10.1016/j.jpubeco.2014.04.010

Stephens-Davidowitz, S. (2017). *Everybody lies: Big data, new data, and what the internet can tell us about who we really are*. Dey Street Books.

Taylor, R. S. (1968). Question-negotiation and information seeking in libraries. *College & Research Libraries*, *29*(3), 178–194. https://doi.org/10.5860/crl_29_03_178

White, R. W., & Horvitz, E. (2009). Cyberchondria: Studies of the escalation of medical concerns in web search. *ACM Transactions on Information Systems*, *27*(4), Article 23. https://doi.org/10.1145/1629096.1629101