

Predicting consumer behavior with Web search

Sharad Goel¹, Jake M. Hofman¹, Sébastien Lahaie¹, David M. Pennock¹, and Duncan J. Watts¹

Microeconomics and Social Systems, Yahoo! Research, 111 West 40th Street, New York, NY 10018

Edited* by Simon A. Levin, Princeton University, Princeton, NJ, and approved August 10, 2010 (received for review April 29, 2010)

Recent work has demonstrated that Web search volume can “predict the present,” meaning that it can be used to accurately track outcomes such as unemployment levels, auto and home sales, and disease prevalence in near real time. Here we show that what consumers are searching for online can also predict their collective future behavior days or even weeks in advance. Specifically we use search query volume to forecast the opening weekend box-office revenue for feature films, first-month sales of video games, and the rank of songs on the Billboard Hot 100 chart, finding in all cases that search counts are highly predictive of future outcomes. We also find that search counts generally boost the performance of baseline models fit on other publicly available data, where the boost varies from modest to dramatic, depending on the application in question. Finally, we reexamine previous work on tracking flu trends and show that, perhaps surprisingly, the utility of search data relative to a simple autoregressive model is modest. We conclude that in the absence of other data sources, or where small improvements in predictive performance are material, search queries provide a useful guide to the near future.

culture | predictions

As people increasingly turn to the Internet for news, information, and research purposes, it is tempting to view online activity at any moment in time as a snapshot of the collective consciousness, reflecting the instantaneous interests, concerns, and intentions of the global population (1, 2). From this perspective, it is a short step to conclude that what people are searching for today is predictive of what they will do in the near future. Consumers contemplating buying a new camera may search to compare models; moviegoers may search to determine the opening date of a new film, or to locate cinemas showing it; and individuals planning a vacation may search for places of interest, to find airline tickets, or to price hotel rooms. If so, it follows that by appropriately aggregating counts of search queries related to retail activity, moviegoing, or travel, one might be able to predict collective behavior of economic, cultural, or political interest. Determining the nature of behavior that can be predicted using search, the accuracy of such predictions, and the time scale over which predictions can be usefully made are therefore all questions of interest.

Although previous work has considered the relation between search volume and offline outcomes, researchers have focused on the observation that search “predicts the present” (3, 4), meaning that search volume correlates with contemporaneous events. For example, Ettredge et al. (5) found that counts of the top 300 search terms during 2001–2003 were correlated with US Bureau of Labor Statistics unemployment figures; Cooper et al. (6) found that search activity for specific cancers during 2001–2003 correlated with their estimated incidence; and Eysenbach (7) found a high correlation between clicks on sponsored search results for flu-related keywords and epidemiological data from the 2004–2005 Canadian flu season. More recently, Polgreen et al. (8) showed that search volume for handpicked influenza-related queries was correlated with subsequently reported caseloads over the period 2004–2008, and Hulth et al. (9) found similar results in a study of search queries submitted on a Swedish medical Web site. An automated procedure for identifying informative queries is described in Ginsberg et al. (10), and based on that methodol-

ogy, Google Flu Trends (<http://www.google.org/flutrends>) provides real-time estimates of flu incidence in several countries. Finally, Choi and Varian (3, 4) have compared search volume to economic activity, including auto and home sales, international visitor statistics, and US unemployment claims; and similar work has been reported for German unemployment claims (11). In this paper, we further this work by considering the ability of search to predict events days or weeks in advance of their actual occurrence.

In so doing we also emphasize an often overlooked aspect of prediction—namely, that performance is relative. To illustrate, consider predicting the weather in Santa Fe, New Mexico, where it is sunny 300 days a year. A prediction of sunshine every day would be correct 82% of the time, yet hardly impressive; nor could a model that fails to outperform the simple, autoregressive rule that tomorrow’s outcome will be like today’s be said to be predictive in any interesting way. Correspondingly, the predictive power of search should be judged in relation to statistical models fit with traditional data sources, prediction markets, or expert opinions—a point that, with some exceptions (3, 4), has been overlooked in previous related work.

Finally, we extend the domain of past studies from epidemiological and macroeconomic time series to include consumer activity such as that associated with movies, music, and video games. These are a natural class of events to consider as they represent activities (e.g., attending a movie) for which it is plausible that individuals might (i) harbor the intention to perform the corresponding action sometime in advance of actually fulfilling it and (ii) signal that intention through a related Web search. In this sense, our paper is related to other work that uses online chatter and Twitter posts to predict rankings of books on Amazon.com (1) and movie box-office revenues (2), respectively. However, by studying three classes of outcomes—the opening weekend box-office ticket sales for feature films, the first-month revenues of video games, and the weekly ranks of songs on the Billboard Hot 100 chart—we reveal some additional insights regarding variation across domains in how search-based predictions perform, both in an absolute sense and relative to alternative forecasting methods.

Results

The potential predictive power of search activity is illustrated in Fig. 1, which shows the volume of people searching for the movie *Transformers 2* (Fig. 1A) and the video game *Tom Clancy’s H.A.W.X.* (Fig. 1B) around their release dates. In both cases searches peak either on or slightly after the release date, but precursors to the spikes appear several days or even weeks in advance. Correspondingly for music, Fig. 1C shows that the rank of the song “Right Round” in terms of search volume closely tracks its rank on the Billboard Hot 100 chart.

Author contributions: S.G., J.M.H., S.L., D.M.P., and D.J.W. designed research; S.G., J.M.H., and S.L. performed research; S.G., J.M.H., and S.L. analyzed data; and S.G., J.M.H., S.L., D.M.P., and D.J.W. wrote the paper.

The authors declare no conflict of interest.

*This Direct Submission article had a prearranged editor.

¹To whom correspondence may be addressed. E-mail: goel@yahoo-inc.com, hofman@yahoo-inc.com, lahaies@yahoo-inc.com, pennockd@yahoo-inc.com, or djw@yahoo-inc.com.

Given the attention that search-based predictions have received recently, it may seem surprising that search data are, at least in some cases, no more informative than traditional data sources. To further examine this point, we reassess the utility of search data in monitoring influenza caseloads, arguably the best-known example of search queries correlating with real-world outcomes. To place that work in context, we note that public reports of flu caseloads by the Centers for Disease Control and Prevention (CDC) are typically delayed one to two weeks; thus the primary aim of search-based flu estimation is real-time monitoring, as opposed to predicting future activity. In particular, flu trend estimates for a given week rely on query volume during that same week: At the end of week t , query volume from that week is used to estimate the yet-to-be-reported flu level for week t . To better gauge the value of search data for flu surveillance, we compare search-based estimates to those from a simple autoregressive model:

$$\text{flu}_t = \beta_0 + \beta_1 \text{flu}_{t-2} + \beta_2 \text{flu}_{t-3} + \epsilon.$$

We find a correlation of 0.86 between current flu levels and recently reported levels where, because of a two-week reporting delay, the estimated flu level during week t is based only on reported levels from weeks $t - 2$ and $t - 3$ (Fig. 4). By comparison, Ginsberg et al. (10) showed that search activity in a given week for terms such as “flu” and “cold” has a correlation of 0.94 with the number of influenza caseloads that week. Search, therefore, does indeed outperform the autoregressive model; however, the difference in performance is relatively small. Moreover, we note that the CDC has recently deployed a system to reduce reporting delays. To illustrate how such an improvement would impact the relative performance of search-based prediction, we repeat the comparison under the assumption that flu reports are delayed by only one week instead of two. In this case, the corresponding autoregressive model has a fit of 0.95, nearly identical to that of the search-based estimates (0.94) of Ginsberg et al. (10). We conclude, therefore, that in monitoring the flu—as in some of the other domains we have considered—search data are comparable in utility to alternative information sources, but not necessarily superior.

Discussion

Across the empirical cases that we have considered, we observe that search counts are generally predictive of consumer activities, such as attending movies and purchasing music or video games, that will take place days or even weeks in the future. As Fig. 5 shows, however, there is also wide variability in the predictive power of search among the different domains. The reasons for these intriguing differences are ultimately unclear and suggest the need for further work; however, we offer the following three possible factors.

First, the size of the relevant population varies greatly across domains. A major feature film, for example, may draw tens of millions of viewers within the United States, whereas popular, platinum albums sell on the order of 1 million copies. One might




Fig. 4. Actual and estimated flu levels in the United States, where flu level is the percentage of physician visits that involve patients with influenza-like illnesses. Search-based estimates are from Google Flu Trends.




Fig. 5. The correlation between predicted and actual outcomes for movies, video game sequels and nonsequels, music, and flu.

expect, therefore, some variation in signal strength to be accounted for by variation in overall market size across domains. Second, in certain domains searching may be more closely tied to the measured outcomes than in others. For example, it may be the case that movie searches are primarily intended to locate theaters or to purchase tickets online—activities directly related to box-office revenue—whereas music searches may simply reflect interest in a song’s lyrics, as opposed to any intent to purchase it. Finally, the difficulty in identifying relevant queries varies considerably among domains. With movies or video games, for example, one would expect users to search by the name of the object in question, whereas with music it is frequently the case that the song name is less well known than the name of the artist. Because artist searches may arise for many reasons, and may also be conflated with names of other individuals, it is arguably harder to classify these queries. It was for this reason, in fact, that we chose to study song-related search on Yahoo! Music rather than by using general Web search data, in part mitigating the problem of identifying relevant queries.

In addition to variation in the predictive power of search alone, Fig. 5 also indicates substantial differences in the relative value of search data compared to alternative sources. For movies, search is clearly outperformed by the baseline and offers no improvement to it in the combined model. For video games, meanwhile, search performs well relative to the baseline for nonsequels and poorly for sequels, but in each case boosts the baseline in the combined model. And for music, search performs worse than the baseline, but offers a substantial boost in the full model. These findings raise interesting questions regarding the circumstances under which search-based predictions might be useful. Most obviously, the importance of the baseline models suggests that search-based predictions can yield the greatest performance boost when key information—such as marketing or production budgets—is difficult to acquire. In the case of nonsequel video games, for example, it is likely that insiders (e.g., employees at video game companies) would have access to better baseline information and would therefore gain less from search counts than outsiders who lack such information.

Even in domains where baseline models generally perform well, however, search data may still help in some circumstances. In the case of music, for example, although the autoregressive model in general outperforms search, the model performs particularly poorly for highly ranked songs, which have a tendency to drop off rapidly, as can be seen in Fig. 3C. Because search counts

