
KEYNOTE LECTURE

Internet Research: The Question of Method—A Keynote Address from the YouTube and the 2008 Election Cycle in the United States Conference

Richard Rogers

ABSTRACT. Digital studies on culture may be distinguished from cultural studies of the digital, at least in terms of method. This lecture takes up the question of the distinctiveness of “digital methods” for researching Internet cultures. It asks, initially, should the methods of study change, however slightly or wholesale, given the specificity of the new medium? The larger digital methods project thereby engages with “virtual methods,” the current, dominant “e-science” approach to the study of the Internet, and the consequences for research of importing standard methods from the social sciences in particular. What kinds of contributions are made to digital media studies, and the Internet in particular, when traditional methods are imported from the social sciences and the humanities onto the medium? Which research opportunities are foreclosed? Second, I ask, what kinds of new approaches are worthwhile, given an emphasis on the “natively digital” as opposed to digitization? The goal is also to change the focus of humanities and humanities computing away from the opportunities afforded by transforming ink into bits. The effort is to develop the study of natively digital objects (the link, the tag, etc.) and devices (engines and other recommendation machines) that make use of them. After critically reviewing existing approaches to the study of the digital, which largely import method onto the medium, I subsequently propose research strategies that follow the medium. How can one learn from methods in the medium, and repurpose them for social and cultural research? The lecture launches a novel strand of study: digital methods.

KEYWORDS. Digital methods, hyperlink analysis, Internet archive, blogosphere, social networking sites, Wikipedia

Richard Rogers is the Chair in New Media & Digital Culture at the University of Amsterdam. He is Director of Govcom.org, the group responsible for the Issue Crawler and other info-political tools, and the Digital Methods Initiative, which is reworking method for Internet research. Rogers is author of *Information Politics on the Web* (MIT Press, 2004), which was awarded the 2005 best book of the year by the American Society of Information Science & Technology (ASIS&T). His current research interests include Internet censorship, Googlization, and the Palestinian–Israeli conflict on the Web, as well as the post-demographics implied by recommender systems.

The author would like to thank Stuart Shulman and Michelle Sagan Goncalves. All figures, with exception of the heat map, are reprinted by permission from the Govcom.org Foundation, Amsterdam.

Address correspondence to: Richard Rogers, University of Amsterdam, Turfdraagsterpad 9, 1012 XT Amsterdam, The Netherlands (E-mail: rogers@uva.nl).

I would like to introduce to you a research program that strives, in some sense, to redo the agenda for Internet research.¹ I am coming at this from a number of years of experience in building tools that try to embrace the methods in the media, instead of importing methods from the social sciences and elsewhere. That's my point of departure. I am going to talk about where Internet research has been going, especially the major contributions made by the social sciences, in particular beginning around 1998, and then I'll move to the current period, introducing a general approach that I term "digital methods." Up until the call by Steve Jones in 1999 in his edited volume *Doing Internet Research*, we were in an "Internet as cyberspace period," where cybercultural studies dominated, with the idea of the Internet as a virtual realm apart—something that has an asterisk attached to it, something that is out there in its own world, with its own dynamics (Jones, 1999). Those cybercultural ideas have informed, and continue to inform, both popular as well as intellectual perspectives on how to study the specificity of the Internet. Is it a realm that allows for identity-altering transformation? Is it a realm that allows for different kinds of developments than the off-line? Should the Internet be studied separately? If it is studied as embedded in society, are user studies the only way to perform Internet research? How else to study the Internet for social and cultural research purposes?

It was in the year 2000 when the British ethnographers Slater and Miller came out with what, to me, was an important study (Miller & Slater, 2000). They grounded the Internet. They went to Trinidad and Tobago, and they studied Trinis' use of the Internet in cybercafés. And what they found was not that the Internet or that cyberspace was some kind of separate world apart, with those "inhabiting" it being transformed by it. Rather, what the ethnographers found, which is of course typical of ethnography in general, was that the Internet was a space where Trinis performed their own culture. They appropriated the medium in ways which were Trini-specific. While a case study, the implication of this work was more general: If the Trinis were doing it their way, most likely national or other cultures were embracing the medium in their own ways

as well. In some sense it grounded the Internet both culturally as well as intellectually. But what I want to talk about today is what that sort of work accomplished methodologically.

Arguably it set a methodological agenda: you had to go off-line; you had to go to the off-line, or the ground, in order to study the online. One had to study users. And, indeed, this has been the social scientific project. To me, some of the most significant work has been done by the research program run by the sociologist and science and technology studies scholar Steve Woolgar in what was called the Virtual Society? Program, from 1997 until 2002. The question mark was very important for them. They debunked first of all this idea of cyberspace as a realm apart, but they also subsequently grounded findings in a series of empirical studies. Woolgar formulated what he called "Five Rules of Virtuality." Among these rules is that there is no desolation for people who spend a lot of time online (Woolgar, 2002). Rather, online activity stimulates more off-line activity. They formulated what has come to be called the classic digital divide critique, which is to say that people's skills relating to and understanding of the risks of the Internet are unequal. It has to do with particular demographics, et cetera. In formulating these rules, the program also solidified the dominant methodological program for Internet-related research in social science. The program has been summarized in the notion of virtual methods (Hine, 2005). A series of volumes and handbooks has now appeared where the researchers continue to develop quite a classic social scientific armature, which includes interviews, surveys, observation, and others. What I would like to point out in particular is these could be categorized or conceptualized as digitized methods. That is, taking methods—existing methods—and trying to move them online. How best to do an online survey? Is Survey Monkey the way to go or not? Should one opt for the Pro version? How best to formulate your first contact e-mail with a group, with a community? To which mailing lists should I send my questions? All these sorts of things take into account some small differences that the online environment brings with it, and they make for slight changes to existing methods—digitized methods with small amendments.

What I'd like to try to do—and I think that many of us, in a sense, are doing this already—is introduce a new era in Internet-related research where we no longer need to go off-line, or to digitize method, in order to study the online. Rather, in studying the online, we make and ground findings about society and culture with the Internet. Thus, the Internet is a research site where one can ground findings about reality. With this particular idea, I have introduced the term digital groundedness, or online groundedness, where claims about society are grounded in the online. I want to come directly to an example. One of the seminal cases that has come out recently that I think could be situated under this term is Google Flu Trends (Ginsberg et al., 2009). Google Flu Trends is a different kind of Google project, because it's run by Google.org, the non-profit arm of Google. Google Flu Trends uses the Web as an anticipatory medium, so it reintroduces the discourse that things happen online first, or you can find out what's happening in society first by going online. The online is quicker to the ground than other ways of getting to the ground. Google Flu Trends collates search engine queries, and geolocates where these queries have taken place for flu—the word “flu,” and flu-related symptoms—and arguably makes findings that are about one week ahead of those by the Centers for Disease Control, which bases its findings on emergency room reports and similar formal reporting mechanisms. In this case, the Web becomes an anticipatory medium again.² Of course it remains controversial to use the Web as the site to base claims about where flu is happening, as opposed to on the ground in sites like the emergency rooms.

I want to give you one more example. It comes from an August 2007 article in the *NRC Handelsblad*, a leading Dutch newspaper (Dohmen, 2007). They published a story where the question was: Is right-wing culture becoming more extremist in the Netherlands? This sort of question may be applicable for any number of countries, but what I want to talk about is the method employed. Instead of traditional investigative journalism, embedding a researcher (“going native”), going to a repository of leaflets and other ephemera—instead of using those

standard methods—they used the Internet Archive. They looked up in the Internet Archive about 100 sites, and made a data set—an Excel sheet, which they also published—in another special Internet-related data-sharing practice (NRC Handelsblad, 2007). They read the content of right-wing and right-wing extremist sites over a period of about ten years, and they found that the language on the Web sites over the years has become more and more extreme; the words were harsher and harsher. They thereby concluded that right-wing culture in the Netherlands is hardening. They made these findings on the basis of the Internet. For those people who consider the Internet a virtual realm or who have a sense that although the Internet may tell you something, you ultimately have to go to the ground for your baseline, the investigative reporter's method and means of grounding claims are quite radical. They used the Web as the site to ground the findings about society.

What I'd like to do today is think about what kind of data are available in the medium, first of all. And, second of all, I'd like to also think through the ways in which the Internet offers particular research possibilities. And by research possibilities, I'd like to introduce the question of learning from the methods in media. I would like to talk about what Internet-specific analysis would entail. I've already given you an introduction to this, but now I would like to take you through a series of digital objects as well as devices such as engines and platforms, and think through how they offer method. I would like to think through with you what I call “digital methods,” the repurposing of methods in media for social and cultural research. I'd like to talk about the link, how links are normally studied and how I would propose that they could be studied. What kinds of opportunities are on offer in link analysis if we follow the medium and its methods? I'd like to talk about the Web site, engines, and spheres. The blogosphere is, of course, the well-known one; scholars have coined the term “Web sphere” (Foot & Schneider, 2002). I'd also like to talk about the newssphere. I view the spheres as engine-demarcated spaces. I'd like to talk about the Webs, in the plural. Who's the senator from Alaska with the Internets? Or was that George

W. Bush? I think they were unwittingly onto something. There are Webs in the sense of national Webs, and much of this has to do with particular locative Web technology that has emerged as GeoIP. Finally, I'd like to talk about social networking sites and Wikipedia. How they are normally studied and how else one might study them if one seeks to learn from the methods of the medium and think about how to ground claims about society from those spaces.

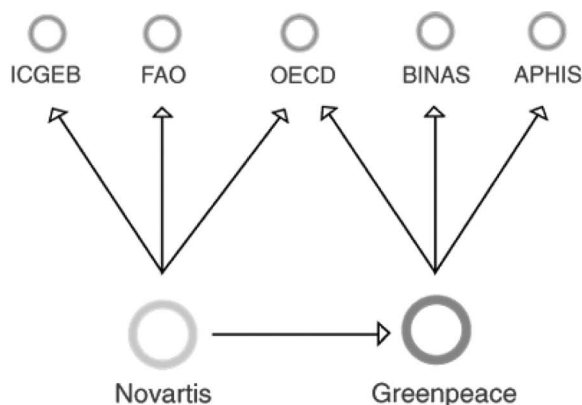
THE LINK

How are links normally studied? There are a couple of traditions here, one of which is from humanities and hypertext theory. This line of thought argues that links author potential stories; the surfer becomes the author going from link to link, authoring a trail. Similarly, in path theory, small-world theory, and social network analysis, what we're concerned about often times are ideas of distance. So how far away are Web sites from one another if we follow links—building on Stanley Milgram's seminal work on six degrees of separation (Milgram, 1967)? And, is there an optimal path between two points? What are the optimal paths to reach someone or something? Social network analysis often times concerns people and their positioning in networks. Are they central or are they in between? Are they brokers or bottlenecks? Et cetera. What would happen if we were to think

through what to do with a link by following methods in the media? One would think immediately of Google. How does Google treat links? Google treats links in some sense as reputation markers, borrowing initially from scientometrics, though also in the tradition of associational sociology. That is to say, sites are ranked on the basis of the number of links they receive from highly influential sites. How can you make use of this particular way of thinking about links? I've created a piece of software called the Issue Crawler, online since July 2001, which builds on the insight that links are reputation markers (Rogers, 2009). The Issue Crawler crawls sites, captures outlinks in any number of degrees of separation, and puts them into a dataset for different sorts of analyses. I present you two types of analyses done with the Issue Crawler. One concerns how links show the politics of association between organizations, and the reputation of sites. What one can learn not from a sort of Google macro analysis treatment of the whole Web, but rather a subset of the Web, a network. One can profile an actor according to the links it gives and receives, either in total or in a particular subject or issue area. The other concerns a method to build out the lists of sites to be checked for blockage, contributing to Internet censorship research.

Figure 1 is a classic piece of work from 1999 and one of the first analyses I conducted. It is an image of the micropolitics of association on display through hyperlinks between Web site types.

FIGURE 1. Normal politics of association displayed through hyperlinking between site types, 1999.



Design by hand. Govcom.org, Design and Media Research Fellowship, Jan van Eyck Academy, Maastricht.

It shows three different kinds of sites: the governmental sites on top; commercial site (Novartis) bottom left; and a nongovernmental organization (NGO; in this instance, Greenpeace) bottom right. You'll notice Novartis links to Greenpeace, but Greenpeace does not link back. Both Greenpeace and Novartis link to government, and government does not link back. These are classic politics of association: governments normally only link to other governmental sites, for instance. Normally NGOs do not necessarily want to endorse other sites that they are critical of by linking to them. This is increasingly the way links are made. In other words, one can begin to gain a grasp of very normal politics of association by showing how sites link to each other, and which links are not reciprocal.

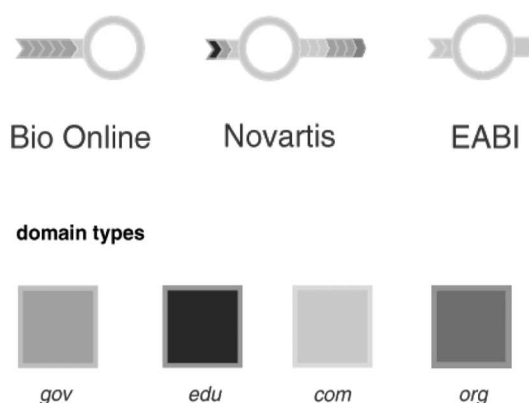
Figure 2 is another example. Here are three separate corporations. They are profiled according to the types of links they receive and the types of links they give. Their reputational status is different depending on the types of links they receive. If a commercial source, a .com, receives links from government, it's a very different status marker than if it received links from only other .com's.

Figure 3 is a picture of the Issue Crawler. It is an Issue Crawler cluster map output, in a particular subject area (e-culture in the Netherlands), with selected profiles of the top actors, and the links they receive and give, to the right.

I would like to talk a little bit about what else one can do with link analysis in the area of Internet censorship research. I occasionally work together with the people at the University of Toronto in the Citizen Lab and the Open Net Initiative (ONI), which was on the cover of *The New York Times* for having discovered a cyber-espionage network, or "ghostnet," allegedly operating out of China (Markoff, 2009). They were contacted by the offices of the Dalai Lama because the Dalai Lama office's computers were acting up. What they discovered and made public was an intriguing information warfare practice: social malware. I'm not going to talk about infowar; rather I'd like say something about my group's contribution to Internet censorship research, particularly methodologically.

The ONI makes a kind of directory of Web sites along 37 different categories including human rights sites, famous bloggers, humor sites, anonymizers, etc. In total, across all categories, at least when I last undertook analysis, they have approximately 2,000 URLs as their sample. They use these 2,000 URLs and query them, or fetch them, in each country in question to see the level of blockage, the level of Internet censorship across some 40 countries. However, I read in the *Cyberdissident Handbook* that came out from the organization Reporters Without Borders, a Paris-based NGO, an article

FIGURE 2. Actor repetitional profiles by inline and outlink types, 1999.



describing the “worst enemies of the Internet” (Pain, 2005). There was a passage where the Saudi Arabia Information Ministry spokesperson boasted that they were blocking something on the order of 400,000 sites. And I thought, ONI is checking 2,000 sites, and the Saudis say they’re blocking 400,000. How do we expand our lists? So I developed a technique called Dynamic URL Sampling whereby we take the initial list the Toronto researchers have drawn up, crawl all of the URLs, and fetch all of the outlinks from these URLs, all these additional pages. We then check them against the original list and the ones that are left over we subsequently check for blockage using, initially, some proxies. Later, we run these through computers that are located on the ground in these various countries because of the lack of transparency of proxies. This “ground check” hardens the findings for ONI. In the example, on Iranian sites, the proxies and the ground check had the same findings.

Figure 4 shows a picture of a network of Iranian social, political, and religious sites, which is an ONI category for Iran. The ones in dark gray are the sites that are blocked, the ones in light gray are the ones that are not blocked, and the ones in dark gray with the little pins on them are sites that we discovered to be blocked but were previously unknown to be blocked by the researchers. I’d like to highlight one site in particular and the difference it makes when one uses hyperlink analysis over building a list of sites in a manual practice. The ONI researchers had BBC.co.uk on their list, and the URL was resolving in Iran. When we ran the network analysis we also ran the page level and we found was that BBC.co.uk/Persian was far more relevant than BBC.co.uk. Indeed the BBC.co.uk newspaper wasn’t blocked in Iran, but the Persian language BBC newspaper was blocked. So we made a contribution not only to methodology in how to expand the lists through dynamic URL sampling, but we also made contributions to the findings. This stuff is double-edged. We would be very good censors because all the blue sites are waiting to be blocked. It’s actually quite difficult to deal with this issue, and the associates from the University of Toronto at the Berkman Center at Harvard sometimes

talk about “data escrows” where they could keep lists away from the prying eyes of censors.

THE WEB SITE

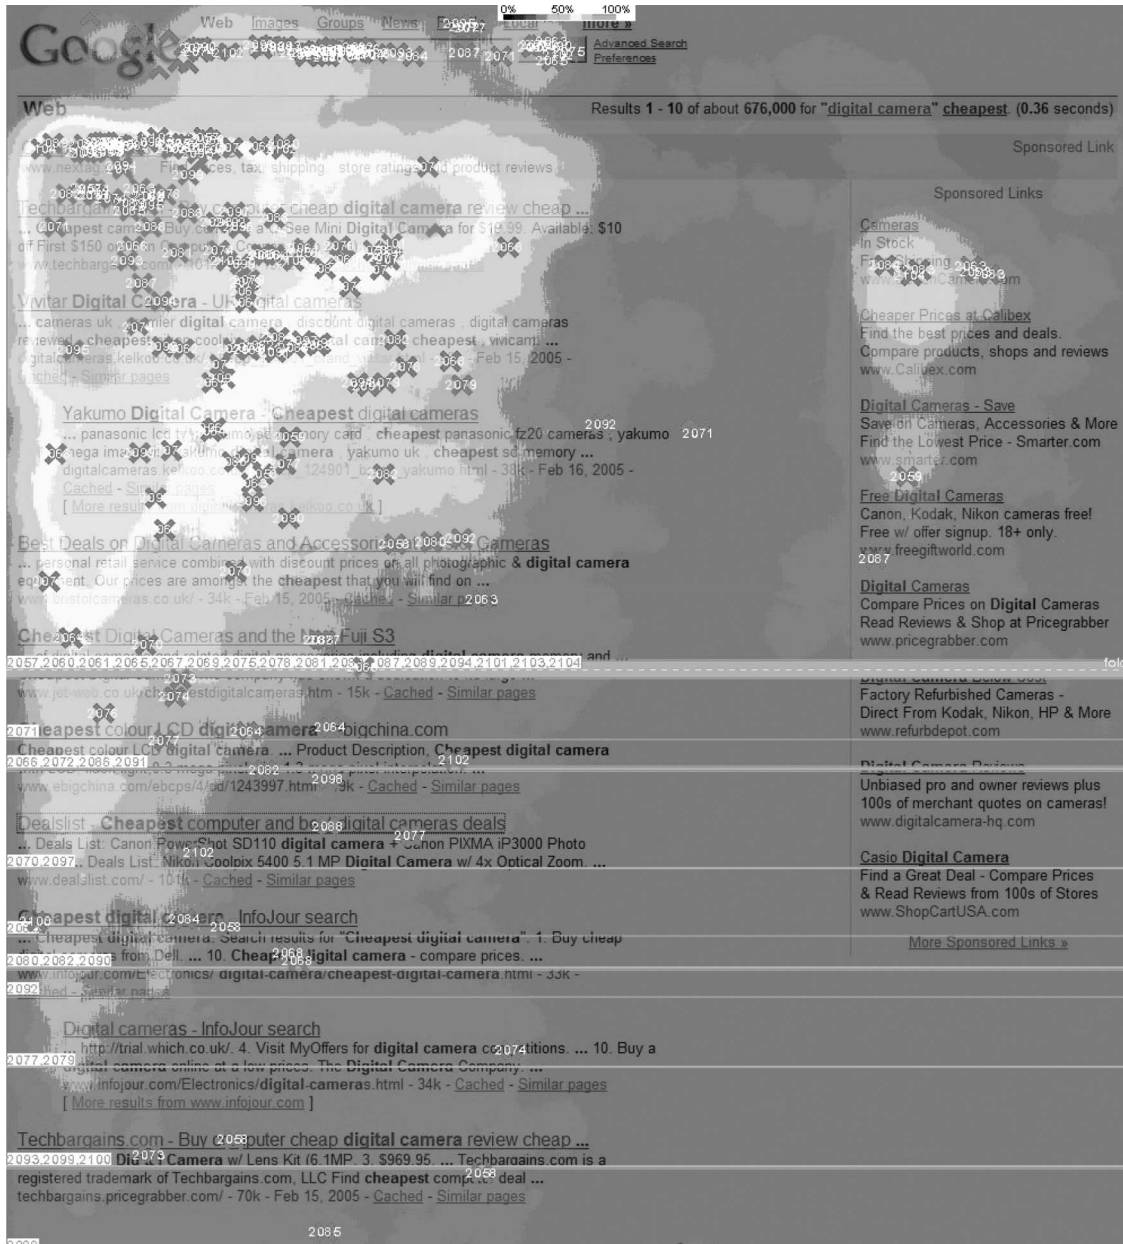
The Web site—how is it normally studied? Well, classically, it’s studied in usability circles. There is a debate between the “Don’t make me think” school and those who are more interested in a poetics of navigation. Web sites are often studied in design circles; for instance, you may or may not know that the majority of the Web is blue. Eye tracking is another classic method for Web site study. The outputs are useful heat maps of eye movement (see Figure 5).

Web sites are also studied as something that needs to be optimized for any number of different reasons, largely because of what I call the drama of search engine space. That is to say, you need to have your site in the top five or ten rankings since search engine user studies have found the number of pages and search engine results pages people look at is in decline. Also, and this is something that I’ve been hearing a tiny bit here, people often study Web sites in terms of their features. Which sites have more features, and is there a correlation or relationship between the features of a Web site and the number of visitors? If there’s more interactivity, is there more participation? Things like this.

Figure 5 is an example of a heat map. The site shown is the Google results page, where eyes are pointed upper left. They call the cluster of heat in the upper left corner the golden triangle of search. And indeed, if you noticed, not too long ago Google moved its menu upper left.

How else can we study Web sites? I have been looking into the Web site as an archived object for some time now. If you run a quick search on Google Scholar, you’ll notice that most of the articles about the Internet archive are about how it works as opposed to how to use it. I’ve also done this in the Web of Science. What I have been trying to do is develop methods or means by which I can use the Internet Archive for research. Similar to how we learn

FIGURE 5. Heat map showing “golden triangle of search,” Enquiro Eye Tracking Report 2005, www.enquiro.com.



from Google for link research, I follow the medium for clues or guidance for archive research. When you look at the Wayback Machine’s interface at the Internet Archive, you notice that it privileges single site histories. You have to type in an individual URL, and a list of dates when that URL has been archived

appears, with an asterisk indicating that the URL content has changed from the page stored at the previous date. What can you do with single site history? It is unlikely that you would like to study the history of the Web through single site histories only, in a kind of biographical approach, but there are things you can do with it.

We've developed a couple of tools that can capture a Web site's history in order to tell a story about it. I made a video documenting this not too long ago. It is a YouTube-style video, and it concerns the Google Directory. You will be familiar with this particular project. The Google Directory sits on top of DMOZ, the Open Directory Project. It was put online by Google in 2000, and it has increasingly been marginalized over the years. Now the Google Directory, and the human-edited Web more generally, are in decline. If you look at the seminal directory, the Yahoo! directory, it's now something that is based on an advertising model—something that you pay for to be listed in quickly. It is also no longer the default search engine on Yahoo!. On Google, the directory is no longer on the front page. What I wanted to do in this particular video is explore what one could learn from the history of a particular page as an example of how one might work with the Internet archive. The video is called *Google and the Politics of Tabs*.³ The video narration says:

This is the history of Google as seen through its interface, from the beginning, sometime in November 1998 all the way up until late 2007. These are screen shots of the Google interface taken from the Wayback Machine of the Internet Archive. The history of Google is important. For some people, Google is the Internet, and for many, it's the first point of access. And Google, as the face of the Internet, has remained virtually the same over the past ten years. But there have been some subtle changes to the interface. So let's go back and look at this in a more detail. You see initially Google with a standard Web search button and its intriguing "I'm feeling lucky" button have been your only options. Then the Directory gets introduced with some front page fanfare. It's the Open Directory Project, DMOZ.org, that Google's built an engine on top of. Then come the Tabs on top of the search box with Web search being privileged at the far left, followed by Images, Groups (that's searching Usenet), and the Directory makes it to the front page. News, the

Google news service, the news aggregator was next. Froogle is introduced; that was that cost-comparison e-commerce service. And that stayed on the front page for a while, then was dropped, followed by Local, which later became Google Maps. You can see that the services are becoming more and more present; there are now five or six at the top bar. Then they add a "More" button. What we're interested in is which services remain on the front page and which get relegated to "More" or "Even More." Let's look at this in some more detail. Let's look at the fate of the Directory over time. It's a story of the demise of the librarian, of the demise of the human editors of the Web, and the rise of the back end, of the algorithm taking over from the editors. Now you see that it's introduced with great fanfare in 2000. The Web is organized by human editors. It remains on the front page. It achieves the Tabs status that we talked about previously. And keeps its place on the front page even as other services are introduced. However, in 2004 something happened: It got placed under the "More" button. You had to click "More" to find the Directory. And in 2006, if you clicked "More," the Directory wasn't there; you had to click "Even More" and there you would find the Directory. As it loses its standing, it also loses recognition. Lots of people don't really remember that there is a Directory just like other services that have left the front page real estate. Also of interest are the services that climb from being "Even More" to "More" and all the way to the front page. But with the Directory, it's a sadder story. As the interface of Google moves upper left, and you click "More," you see that there's no Directory any longer. And you also see that there is no "Even More." So nowadays you have to search Google for its Directory to find the Google Directory.

The "Even More" button is back, by the way. It wasn't there at the end of 2007, and now it's back under the "More" button on the upper left

menu. In the film, I talk about the rise of the back-end, the rise of the algorithm, and the demise of the human. I would like to point out something important in this respect. If you go to the Yahoo! directory, and you type in a query, what you get back is a listing of sources that are now ranked by popularity by default. That is to say the sources are no longer in alphabetical order. So the egalitarian listing of information sources is no longer the default at the original directory. The algorithm is also spreading beyond the Internet to other digital spaces. In a comparative media analysis perspective, you would look into these sorts of new rankings and recommendations taking over from the alphabetical list or chronological list; consider TiVo.

THE ENGINE

How are engines normally studied? There is a body of work on the politics of search engines that looks into search engines as sites or spaces of inclusion or exclusion, where particular sites are buried or they no longer exist in practice if they are way down in the rankings, or certainly out of the top 1,000. Engines are also studied in terms of what could be called the attention deficit disorder. Users are increasingly looking at fewer and fewer returns, fewer and fewer pages. Jansen and Spink in particular have been studying this for a number of years, and not only are people looking at fewer and fewer pages and returns, but they're clicking sites that are closer and closer to the top (Jansen & Spink, 2003, 2004). Engines are studied as a space where placement really matters.⁴

They are also studied in terms of the notion of Googlization. It's a term that has been introduced by library science scholars in particular as a reaction to the Google Books project. The minute Google entered the hallowed halls of the library, library science scholars began critiquing Google quite heavily, but also developing sophisticated ideas about what Googlization would imply for knowledge provision, and knowledge access more generally, if it keeps going like this. Googlization arguably, as a term, has a political economy connotation. Google is creeping into more and more different services.

They're no longer just a search engine; that much is obvious. Google and engines are also studied from surveillance and privacy studies points of view. In particular, search engine results are being personalized on the basis of your query histories. If you are signed into Google, especially, queries are not only logged, but results personalized. And it's interesting in my view that it's becoming more and more difficult to study Google results, because Google results are not necessarily the results increasingly of some universalizing algorithm for all, but they're also partly your results. People increasingly do not receive the same results. So I call Google the "inculpable engine," as it's taking itself off the hook by having the user influence the results. But before it did so, or while it's still doing so a bit, I have been developing is a means to study engines—Google in particular—as an ordering device, as an epistemological machine. And in order to do that, I have captured and stored the engine results, which is not in compliance with Google's Terms of Service, so I put up a notice asking for its forbearance. I look into what is an understudied aspect of engines: the volatility or stability of the actual results. Do results change day by day, or are they relatively stable? Does it matter *when* you search for the kind of results that you receive?

Table 1 shows the results of a query, made daily, over a 30-day period in November 2007. It's a query for RFID. Within a 30-month period, you see that most of the sites are rather stable. But some returns vary somewhat from a top ranking of four to the low ranking of 12, from the top ranking of 11 to 26, from 1 to 17, and from 14 to 31. The major change is one particular site that went from rank number 9 to 213 during the 30-day period. These are important changes, given how users interact with results, and the significance of the top returns.

What does one do with captured search engine results? What I am interested in is in one sense a follow-up on the classic idea of what social research is. According to C. Wright Mills, it is to present no less than conflicting realities themselves (Mills, 1959). What I was looking at was whether or not Google results are increasingly becoming more and more familiar,

TABLE 1. Volatility or Stability in Search Engine Results Space

URL	Highest rank	Lowest rank	Fluctuation
en.wikipedia.org	1	1	0
rfidjournal.com	2	2	0
news.google.com	4	4	0
aimglobal.org	4	5	1
spychips.com	6	13	7
webopedia.com	7	47	40
rfidgazette.org	8	12	4
news.com	9	9	0
rfid.weblogsinc.com	10	12	2
networkworld.com	11	13	2
rfid-weblog.com	11	15	4
epic.org	14	18	4
eff.org	15	26	11
rfidvirus.org	16	17	1
rfidinc.com	16	18	2
ti.com	17	31	14
theregister.co.uk	19	27	8
howstuffworks.com	19	30	11
rfidc.com	20	29	9



Note: Google results rankings for the query of RFID, with highest and lowest ranking per site over 30-day period, September 19–October 19, 2007. Output of the Issue Dramaturg by Govcom.org, Amsterdam.

that is, able or unable to present conflicting realities. The question is whether the results that come out of Google are aligned with the familiar or the mainstream. That is, are Google results becoming quite similar to the sources that you would hear on the evening news? I have been tracking results of the query for 9/11 over about two years, collecting the top 1,000 results for that query because Google serves a

maximum of 1,000 results, and looking in particular at the rankings over time of three important sources for 9/11 accounts generally: *The New York Times*, the New York City government, and a third one which I'll tell you about in a second. In Figure 6, the New York city government, nyc.org, is denoted A, and *The New York Times*, c. The other one, B, is a site that presents a conflicting view of reality, the Movement for 9/11 Truth, 911truth.org. For approximately six months, 911truth.org was in the top five results of Google. Something happened around the 17th of September 2007; it dropped precipitously from result 5 to result 200 and then off the charts to under 1,000. I believe this is the first fully documented case of the “disappearance” if you will, or the apparent removal of a site from Google results, apart from legal cases (Grimmelman, 2008). There are a number of reasons why this may have occurred, for example if Google believes, or has detected, that 911truth.org is in a spammy neighborhood, or if the Web site received too many inlinks too quickly, as if operating a link farm. While I do not know, I believe 911truth.org received many links around the 11th of September from franchise or chapter sites, and it was a false positive for Google.

In the case above, we are studying Google mainly, though one could interpret 911truth.org's high placement as socially significant more generally. How do we use Google to study more specifically what is happening in society? We've built a piece of software called the Google Scraper, also known as the Lippmannian device because Walter Lippmann was always

FIGURE 6. A Web site is gone. The apparent removal of 911truth.org from Google results for the query “9/11” September–October, 2007.



interested in equipment or tools that could provide a “coarse view” of the partisanship of an actor. The Lippmannian Device captures the first 100 engine results from one query, and then for each of these 100 sources, we query them individually for a particular subject matter. We’re interested in seeing whether or not we can provide a coarse sense of the partisanship of a particular organization. Figure 7 shows the top 100 returns for the query “climate change” in Google in July 2007. We queried each of the individual sites—the EPA, BBC News, UN Environmental Program, IPCC, Pew Climate, et cetera, for the names of well-known climate change skeptics. We wanted to look into not only whether we could detect the partisanship of the source, but also we wanted to look whether the Web, like the news, was providing quite a lot of space or voice to the skeptics. As Figure 7 shows, the skeptics are actually are not named too often on very many sites.⁵ With the

Lippmannian Device you gain a sense of not only partisanship, but also issue or position commitment per source. One notes both the partisans as well as the watchdogs.

THE SPHERES

How are blogs often studied? Blogs are often studied as a genre; they’re recognizable because they have particular formats: reverse chronological order and a blogroll, for instance. The blogosphere often times is studied in relationship to the news. What other researchers reported earlier at this conference is counterintuitive to me; in the previous studies that I’ve looked at and the studies that I’ve done, I’ve always found the blogosphere to be parasitic on the news as opposed to the news being parasitic on it. But in any case, the blogosphere, often the political blogosphere is quite obsessed with mainstream media. Blogs

FIGURE 7. Climate change skeptics on the Web (Frederik Seitz). Quantity of mentions of climate change skeptic in top climate change sites on the Web, July 30, 2007.

Climate Change Sceptics on the Web (Frederick Seitz)

Research Question To what extent are climate change 'skeptics' present in the climate change spaces on the Web?

Findings There is distance between the skeptics and the top of the search engine returns.

epa.gov (0) bbc.co.uk (0) defra.gov.uk (0) unep.org (0) bom.gov.au (0) ipcc.ch (0) pewclimate.org (0)
 davidsuzuki.org (0) panda.org (0) mfe.govt.nz (0) ec.gc.ca (0) exploratorium.edu (0) climatechange.com.au (1)
 greenpeace.org (0) climatechallenge.gov.uk (0) guardian.co.uk (0) iisd.org (0) g8.gov.uk (0) campaigncc.org (1)
 foe.co.uk (0) state.gov (0) scidev.net (0) eea.europa.eu (0) whoi.edu (0) cbc.ca (0) energy.gov (0)
 marshall.org (8) climateark.org (4) un.org (0) dar.csiro.au (0) theglobeandmail.com (0)
 acfonline.org.au (0) gcrio.org (0) nature.com (0) grida.no (0) nature.org (0) ecokids.ca (0) royalsoc.ac.uk (0)
 climatechangecentral.com (0) iea.org (0) ecn.ac.uk (0) ecy.wa.gov (0) worldwildlife.org (0)
 realclimate.org (35) faqs.org (0)
 metoffice.gov.uk (0) open2.net (0) scienceagogo.com (0) eldis.org (0) ft.com (0) who.int (0) climatecrisis.net (0)
 itscotland.org.uk (0) abc.net.au (0) climatechange.ca.gov (0) envirolink.org (0) mofa.go.jp (0)
 sourcewatch.org (21) iucn.org (0) dfat.gov.au (0) ncdc.noaa.gov (0)
 climatescience.gov (0) climatechangecollege.org (0) ciel.org (0) ucar.edu (0)

Source_google.com
 Query "Frederick Seitz"
 Method Search for query "Frederick Seitz" in top 100. Organized in order.
 Tools Google Scraper and Tag Cloud Generator
 Date_30 July 2007

Product of the Digital Methods Initiative,
 dmi.medialstudies.nl. Analysis by Bram
 Nijhof, Richard Rogers and Laura van der
 Vlies. Design Anne Helmond.



CC-BY-NC-SA

are also studied as organizers of voice, as voice-giving, or as authentic voices as in the case of the famous Iraqi bloggers and others.

How else to study spheres? I take spheres to be engine-demarcated spaces. That is to say the blogosphere is in some sense authored or at least demarcated by Technorati, and more recently Google blog search. I take the Web as in some sense demarcated by Google Web search; I take the newssphere if you will as demarcated by Google News; I take the social bookmarking sphere, or tagosphere, as demarcated by Delicious. And what I then do is perform “cross-spherical” analysis. That is to say, what are the differences in available or privileged sources between these subspaces, or spheres, on the Web? I also ask questions about the quality of media. Is the blogosphere something that treats issues, subject matters, or substance qualitatively differently, in terms of source composition, than the news or the Web? One brief case study, again having to do with climate change, is called the Issue Animals Project. We queried the various sphere-engines for climate change, and saved the results. We subsequently made a list of animals associated with climate change, manually, from reliable sources. We queried each of the sources per sphere for these animals, that is, on the Web through Google Web search, in the news through Google News, and in the blogosphere through Technorati. We queried all of these individual sources, initially returned from the climate change queries, to see whether or not particular animals are privileged per sphere, in particular to look into whether each of the spheres have tendencies towards creating media icons or not. When looking at the news, we noticed that the polar bear really stands out. As I said, I have found the blogosphere to be quite parasitic on the news, so the effects of the news become amplified in the blogosphere, whereas on the Web, intriguingly, the animals are treated in a more egalitarian fashion, in that they are more evenly distributed. That may say something about the “quality” of the Web over the news and over the blogosphere, or at least its lacking media icons. Perhaps that is the difference between a media space and an information space.

THE WEBS

The Webs are in the plural as I mentioned. How are they often studied? They’re often studied in the singular, as in qualities of the World Wide Web, or Web Studies. They’re often times studied, as I mentioned before, in terms of cyberspace. They are also treated as technical infrastructure, as a technical infrastructure with particular end-to-end principles and with particular engineering—the packet-switching—which conventionally, at least way back when, was supposed to allow us to route around censorship. Anyway, cyberspace, the idea of cyberspace arguably grows out of a particular understanding of the effects of the infrastructure. And the Web is also studied as a realm apart, as something different, certainly not as a potential baseline.

What I started to do two years ago was to see if I could use the Web as an indicator for the state of a country. I was exploring Iraq in particular. Information about Iraq came from news reports, the authentic bloggers from Iraq, and some Presidential candidates and Congress people on fact-finding missions. There were no tourists until very recently. I have tried to develop a means by which I could find out what was going on in Iraq by looking at the state of its Web. How broken is it? Were the university Web sites up? Things like this. On the basis of this thought, I began to develop a series of methods in order to diagnose the condition of a society according to its Web.

The reason for national Web thinking more generally, more conceptually, has to do with GeoIP technology and the idea that the Web has been grounded. I notice it a lot because I’m a baseball fan, and I subscribe to MLB.tv, and any time I’m in the U.S., I’m often blacked out, whereas when I’m in Europe I can watch all the games. You may have noticed this with the Olympics. It means that the content (and advertising) are served according to your location. That is, your IP address is detected. This technology goes back to the famous Yahoo! lawsuit brought by two French NGO’s in 2000 when French users ultimately were blocked from looking at Nazi memorabilia pages in the U.S. The GeoIP detection technology was developed directly as a result

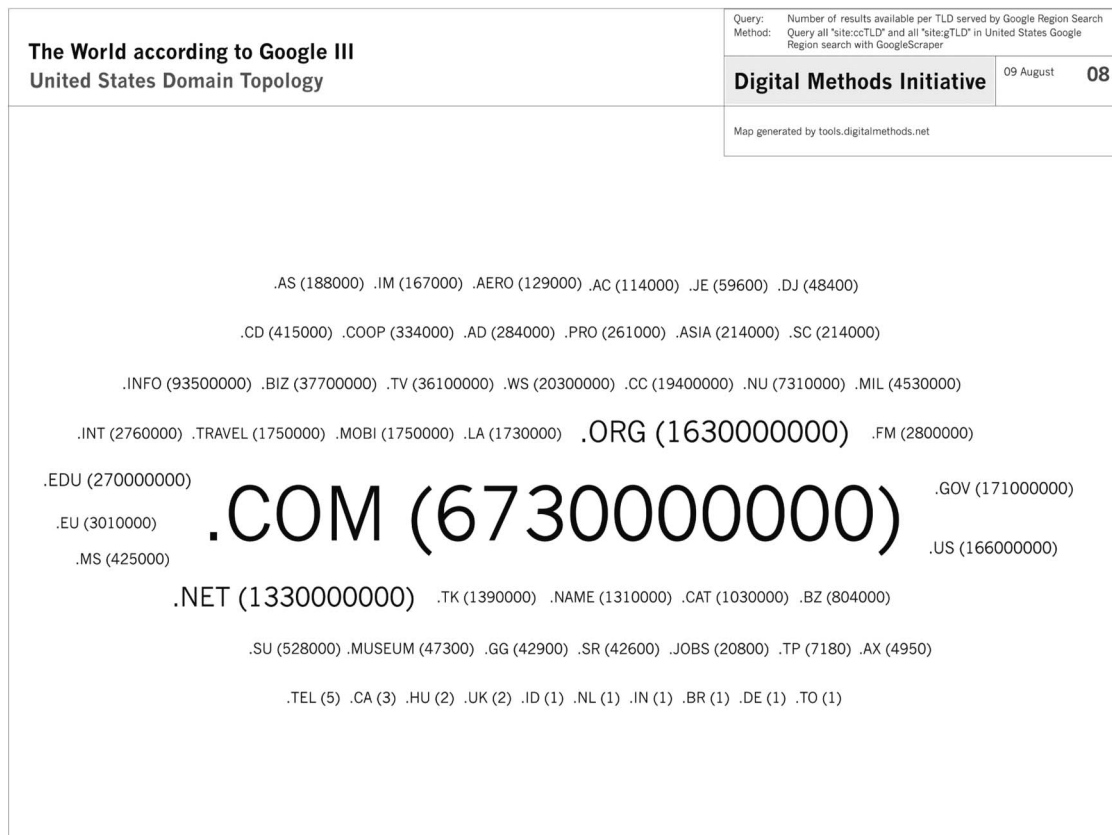
of that particular lawsuit. Goldsmith and Wu have written a book about it (2006).

So what kinds of ideas can one gain about a society if you look into the condition of its Web? I mention briefly five ways to study the condition of the Web: youthfulness, brokenness, cohesiveness, datedness, and dated users. *Youthfulness* asks if pages are fresh. You can determine that through analyzing page date stamps. How *broken* is the Web? Brokenness can be tested using link validators. To find out the *cohesiveness* of a national Web, you can use hyperlink analysis, which asks are sites linking to one another nationally, or more internationally? How *dated* is the Web? What kind of software or server versions are they running? Also, how dated are the users, if you will? That is to say, what kind of browser versions are they using? You can use server logs for that, if you have access.

Here is another way to gain indications of societal compositions according to the Web. Figures 8 and 9 are top- and second-level domain topology of the U.S. and Palestinian territories, respectively, showing rankings of which domains are in use. You can see that “.com” is the most important by far in the U.S. Compare that to Palestinian territories, where “.org” is most significant.

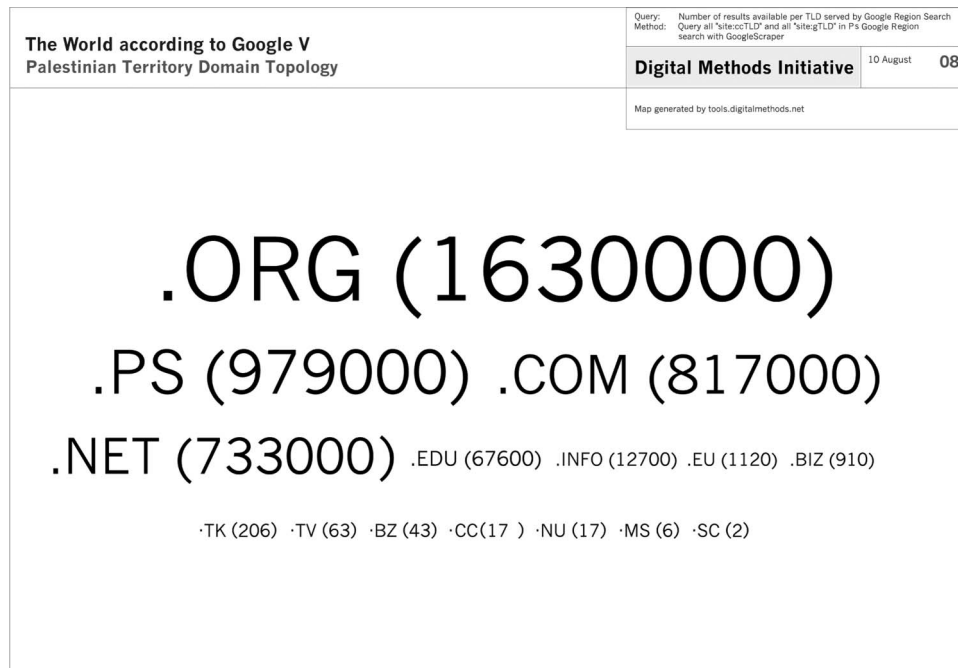
I'll briefly mention the Palestinian Web mapping project, in this context, for it demonstrates how to study social divisions. We undertook this study with Cambridge University and the University of Toronto. It is an effort to see what is going on in the Palestinian territories by looking at the Web. We took all the Fatah-related Web sites and all the Hamas-related Web sites, and we looked at their linking patterns. What we found is that Fatah sites link to the news, to local NGOs, to international NGOs;

FIGURE 8. The world according to Google. Distribution of domains in the U.S. according to Google region search, August 2008. Centered domain cloud.



Output of the Google Scraper by Govcom.org, Amsterdam.

FIGURE 9. The world according to Google. Distribution of domains in the Palestinian Territories according to Google region search, August 2008. Ordered domain cloud.



and Hamas linked to no one. They link only to RSS readers. The work gives you a clear indication of how their networks operate. It provides a particular view of the differences between Fatah's and Hamas's information cultures, one linking with national, international NGO's, and news organizations, and the other off-the-radar, with individuals receiving the information through subscription only. And it also gives you an indication of the splits on the ground with two different kinds of Web presences and cultures.

SOCIAL NETWORKING SITES

How are they normally studied? They're often studied in terms of Goffman's idea of the presentation of self (Goffman, 1959). One of the more intriguing reports was about social class divides performed on social networking sites. MySpace and Facebook are said to have different classes of users associated with it. The extent to which Facebook is a walled garden versus MySpace is another way of looking at

social class. Social networking sites are also studied in terms of the difference or similarity between real-life friends and "friended" friends and how difficult it is to "defriend" online, and how it amplifies the effect. Touchy social relations are not resolved very well by clicking, and having alerts broadcasted.

How else can they be studied? I've been introducing a term recently that I like to call "post-demographics." The term takes into account the kind of information on the profile of social networking sites that is different from the standard demographics. What's on a profile of a social networking site? The ones that are highlighted are favorite media interests—that is to say, movies, TV shows, books, heroes, and things like this. So with post-demographics, I propose to again follow the media and study how profilers already make use of these preferences, of these particular favorites, and then repurpose the way they do it, their method, for social and cultural research.

EiFRIENDO.com is my first attempt at this; it's more of an art project. I made EiFRIENDO

when my team and I were artists in residence at the National Media Arts Institute in the Netherlands, Montevideo. It gives an indication of the work that can be done in post-demographics in a more systematic way. It shows the difference in favorites between all the friends of Obama versus all the friends of McCain. What are their friends' aggregate favorites? One could study Obama's and McCain's supporters, according to demographics, but what about studying Obama's and McCain's friends, according to their favorites? It's interesting that they have quite distinctive profiles, according to the music that the friends listen to, the movies, TV shows, books, and heroes. And TV shows may be of interest to advertisers and political consultants. Obama's friends' favorites, for instance, include *The Daily Show*, *Lost*, and *The Office*; McCain's friends: *Family Guy*, *Project Runway*, *Top Chef*, *America's Top Model*, and *Desperate Housewives*. The larger point of post-demographics is that relationship between candidates and friends' media favorites may be distinguished from the relationship between candidates and supporters' demographics (e.g., gender, income, education level).

WIKIPEDIA

This is the last topic. How is Wikipedia often studied? It's often studied in terms of its accuracy. You will have seen the studies in *Nature* about Wikipedia vis-a-vis *Encyclopedia Britannica*. It's also studied in terms of its "encyclopedianness," if you will. Indeed, it is remarkable that Wikipedia is encyclopedia-like. It's also often studied or often used and studied as a kind of scandal machine. This practice in particular has picked up since Virgil Griffith at Caltech made the Wikiscanner, which de-anonymizes anonymous edits. And it also studied or thought about at least in terms of the highly "vigilant" community. How can Wikipedians be so vigilant and also accurate given the fact that they are (a) amateurs and (b) free labor. To put them to the test, a number of scholars created false Wikipedia entries or changed things in a Wikipedia entry

and then waited for something to happen. Some have later expressed regret for doing this. What happened was that many of these changes were corrected quickly, which came as a surprise. Wikipedians are highly vigilant. How?

I just want to tell you first that the Wikiscanner rocked the Netherlands. One of the Princes, Prince Friso, and his Princess apparently were caught editing a Wikipedia entry. On the Princess Mabel entry, it was written that she had given "false and incomplete" information to the government prior to wedding the Prince. And this was in quotation marks in the Wikipedia entry. It was found that the Royal Family, or at least someone at their IP address, had removed one of the words, changing, "false and incomplete" to "incomplete." It was front-page news, and created a scandal. What wasn't reported was that the edit was reverted (changed back) within minutes, because one of the vigilant Wikipedians probably received a software alert saying that the entry had been edited.

One of the things that I'd like to point out is that most Wikipedia research today has forgotten the bots. In fact, if you go to the statistics of Wikipedia, the top Wikipedians are bots; they're not humans. The bots are working in tandem with the humans. Why are Wikipedians vigilant? They have bots. And they have software alerts that tell them when something is changed, something has been reverted, or that something has been edited.

The initial question is, "How dependent is Wikipedia actually on bots?" And how would one begin to think that through? It turns out that in total, the number of human edits is far greater than the number of bot edits. However, when you look into different languages, you see that particular languages are more reliant on bots than others. In particular, the languages most endangered have the most bot activity compared to human activity. What are the bots doing? They're looking for vandalism, they are interlinking pages, et cetera. The most active bots are referred to internally as maintenance bots, a term which is disarming. However, the question is, where does the "maintenance" end and content co-authorship between humans and bots actually begin?

NOTES

1. This article is an edited transcription of the keynote address given at the YouTube and the 2008 Election Cycle in the United States Conference at the University of Massachusetts–Amherst on April 16, 2009.
2. In referring to the Internet as becoming an anticipatory medium “again,” the author is referring to Rogers, 2003.
3. The complete video may be seen as part of the author’s keynote address at <http://youtubeandthe2008election.jitp2.net/conferencevideo>.
4. The author edited a volume on the subject of preferred placement and paid inclusion in search engine returns. See Rogers, 2000.
5. Additional sample output from the Lippmannian Device was presented in Rogers’s keynote address and is available at <http://www.digitalmethods.net/> (the Digital Methods Initiative site), <http://www.mappingcontroversies.net> (the Macospol EU project site) and <http://youtubeandthe2008election.jitp2.net/conferencevideo> (the conference Web site).

REFERENCES

- Dohmen, J. (2007, August 25). Opkomst en ondergang van extreemrechtse sites [The rise and fall of right-wing extremist websites]. *NRC Handelsblad*. Retrieved January 15, 2010, from http://www.nrc.nl/binnenland/article1831689.ece/Opkomst_en_ongegang_van_extreemrechtse_sites
- Foot, K., & Schneider, S. (2002). Online action in Campaign 2000: An exploratory analysis of the U.S. political Web sphere. *Journal of Broadcasting & Electronic Media*, 46(2), 222–244.
- Ginsberg, J., Mohebbi, M. H., Patel, K. S., Brammer, L., Smolinski, M. S., & Brilliant, L. (2009). Detecting influenza epidemics using search engine query data. *Nature*, 457(7232), 1012–1014.
- Goffman, E. (1959). *The presentation of self in everyday life*. Garden City, NY: Anchor Doubleday.
- Goldsmith, J., & Wu, T. (2006). *Who controls the Internet? Illusions of a borderless world*. New York, NY: OUP.
- Google.org. (2010). *Google flu trends*. Retrieved January 15, 2010, from <http://www.google.org/flutrends/>
- Grimmelmann, J. (2008). The Google dilemma. *New York Law School Law Review*, 53, 939–950.
- Hine, C. (Ed.) (2005). *Virtual methods: Issues in social research on the Internet*. Oxford, England: Berg.
- Jansen, B. J., & Spink, A. (2003). An analysis of Web information seeking and use: Documents retrieved versus documents viewed. In *Proceedings of the 4th International Conference on Internet Computing* (pp. 65–69).
- Jansen, B. J., & Spink, A. (2004). An analysis of Web searching by European Alltheweb.com users. *Information Processing and Management*, 41(6), 361–381.
- Jones, S. (1999). Studying the Net: Intricacies and issues. In S. Jones (Ed.), *Doing Internet research: Critical issues and methods for examining the Net* (pp. 1–28). London, England: Sage.
- Markoff, J. (2009, March 29). Vast spy system loots computers in 103 countries. *The New York Times*, p. A1.
- Milgram, S. (1967). The small-world problem. *Psychology Today*, 1(1), 61–67.
- Miller, D., & Slater, D. (2000). *The Internet: An ethnographic approach*. Oxford, England: Berg.
- Mills, C. W. (1959). *The sociological imagination*. New York, NY: Oxford University Press.
- NRC Handelsblad. (2007, 25 August). Overzicht van vrijspreekende, rechtse en extreemrechtse websites. *NRC Handelsblad*. Retrieved January 15, 2010, from http://nrc.nl/redactie/binnenland/overzicht_websites.htm
- Pain, J. (2005). Internet-censor world championship. In S. Devilette (Ed.), *Handbook for bloggers and cyberdissidents*. Paris, France: Reporters without Borders. Retrieved August 16, 2006, from http://www.rsf.org/IMG/pdf/Bloggers_Handbook2.pdf
- Rogers, R. (ed.) (2000). *Preferred placement: Knowledge politics on the Web*. Maastricht, The Netherlands: Jan van Eyck Editions.
- Rogers, R. (2003). The Viagra files: The Web as anticipatory medium. *Prometheus*, 21(2), 195–212.
- Rogers, R. (2009). Mapping public Web space with the Issuecrawler. In C. Brossard and B. Reber (Eds.), *Digital cognitive technologies: Epistemology and knowledge society* (pp. 115–126). London, England: Wiley.
- Woolgar, S. (2002). Five rules of virtuality. In S. Woolgar (Ed.), *Virtual society? Technology, cyberbole, reality* (pp. 1–22). Oxford, England: Oxford University Press.

APPENDIX

Question and Answer Period

Moderator: All right, thank you, Richard, as always for telling us what everybody’s doing and has done in an authoritative way and for telling us all the things that we might consider doing, which is one of the reasons that we’re here. So it’s question time for Richard. I’ve got a hand in the back, Steve?

Attendee: My team, the research group that I’m a member of at Cornell, studies very similar things to the things you talked about today. One of the things we do is take advantage of problems

in the algorithms to make the algorithms better and let me give you an example. If you flip back to your slides to the section on where you were pulling down the hundred search results and then querying the next level of detail to look for under-cited authors. We've actually built systems that do essentially the same thing but in order to improve rankings based upon user preferences. So what you think about the fact if we do that, what are the implications of it?

Rogers: I think it's excellent. Let me contextualize a larger point, and then I'll come back more specifically to your question. Recently one of the leading new media theorists, Lev Manovich, has called for a program called Cultural Analytics. The term is borrowed from Google Analytics. It would build quite large-scale data collection facilities to take advantage of all the digital traces online and analyze them to think about culture production, state of culture, et cetera. That's a particularly large-scale model; it's kind of big science type of idea. And what I'm interested in are far more modest research undertakings. That is, instead of thinking about the models of these large companies and their large datasets and getting negotiated access to them, I'd rather think about ways that we can use the methods and computing techniques that are being implemented online and then think through what kinds of other sorts of research can be done with them, how these sorts of techniques can be repurposed. So indeed when you create techniques to better the rankings, or improve algorithms for ranking, I'm interested in using those algorithms—repurposing them—in order to query different sites to tell us whether or not these sites are in league with a particular position or friendly with particular kind of funders, et cetera. So it's a different kind of purpose, but we build on the very important work that you're doing.

Questioner: There's a huge level of funding that's being pumped into this area, it's called, we call it "learning to rank" in the computer science world. And so thoughts that you have about this area would be very useful to any team that's working on that research agenda.

Attendee: I'm interested in your discussion of national Web diagnostics. And the focus on

Iraq is fascinating. I mean Iraq is certainly an extreme case. To what extent have you gone beyond Iraq, I mean to what extent have you developed metrics to try to and how much data have you gathered on all 178. It's embarrassing as an IR person to say I don't know the actual number of jurisdictions, but do you know the number of (ccTLD's)?

Rogers: 245. I'm not exactly sure how many. I mean, we're now developing metrics. So we've done a very brief case study on Iraq. We've done a more extensive study on Palestine. And of course these are very specific.

Attendee: My question is, aren't they outlier cases? I mean, you're going to get interesting stuff from that, but are you trying to develop a metric that's broader?

Rogers: Let me just address a larger point and then I'll come to the specifics. I make situated software. So what I normally do is develop techniques, software applications, for specific kinds of research questions. And then later I see whether or not they could be made into something more generic. So this is a particular kind of research practice and I just want to make this clear. It's very different from the standard social scientific instrument-making, whereby you build an all-purpose tool that you then install on your machine, then you go and look for datasets, and plug them in. So what I do is I normally make things that are for quite specific situations. But in most of the work that I've shown you apart from the Iraq case, many tools later have been developed in to something more generic, that is, for more than just the one research project. What I'm trying to develop at least in the first instance are a set of metrics that will aid in diagnosing the condition of particular kinds of countries when we don't have good diagnoses from media. I mean, someone wrote me very recently and wanted to do something on Zimbabwe. But, apart from these specific cases, what can we learn from more country data, for example, from the Alexa data? Have you looked at the top 100 sites per country? It's interesting in the sense that you can profile a country according to what kind of sites are in that top 100. Which kinds of countries are relying on the mega-upload sites? Just to give you one short example. So one can think about different

sorts of Web indicators for ideas about the societal condition.

Attendee: I like your discussion of links. Actually I like the way you organized it from how it's been studied and what might be done differently. But on links, I was interested in the contrast you were drawing and I wasn't sure I understood the contrast between social network analysis which was in the "how it's been done before" versus this sort of reputational understanding, which is the "how might we do it." How is looking at sorts of links as reputational different from just in degree centrality as a prestige measure as in social network analysis?

Rogers: That's a very good question. What I tried to do is go one step further and talk about the micro-politics of association. So that would be the answer in short. Normally when one would study hyperlink relationships in a qualitative way, one would try to think about whether or not they have an off-line relationship, whether or not they're partners, whether or not they're allied—these sorts of things. And that's how one would explain why it is that they're linked. That's how one would do it in a social networking sense, if you will. But what I'm proposing is you can find a politics of association where you don't need that kind of baseline of off-line relationships in order to come up with a reputational marker.

Attendee: I'm so grateful that you're doing things like this. One of the things that happens to me is I immediately start thinking: What kinds of interesting questions can we ask and then use these tools to answer? And I would like to say, what are two or three interesting research questions that you think these would answer? And a follow-up on Dan's question about Iraq and Palestine—I just think the next one is China and I'll just give you a quick answer. We have a group that comes regularly and does some executive ed. from China and I give them some lectures on technology and technology-policy and we have this conversation about "censorship" of the Internet in China.

And they acknowledge that these Web sites are blocked but they are grateful to their government because it prevents terrorism. And so I'd be really interested in having some of this evidence and say, "Okay so this is blocked but what kind of things does it block?" Just throw that out, I mean my brain's going a little faster than I think I can articulate, but what interesting broader research questions would you like to apply these things to?

Rogers: One of the things that I'm interested in, in the context of Internet censorship research, is this: To what extent does a circulationist Web pre-empt or forestall state Internet censorship? I didn't show it, but I did a case study on the Balochi or Baluchi Web, some people pronounce it different ways—sites in Pakistan. Those sites are routinely blocked by the state. However, the question is: Does one see the Web in a kind of old media style, as a set of single Web sites that are blocked or unblocked? Or do you see the Web as a content circulation space? So what I've tried to do is test the idea that the Web is a content circulation space by taking all of the sites that are blocked in Pakistan—all the Baluchi sites—skimming the content off of them, and then querying engines for that content to see whether or not, literally, that content has been repackaged, or moved to other parts of the Web. The answer is: very little. In other words, Pakistan is doing a very good job censoring Baluchi sites, and Baluchi content is not circulating well by other means. I did a similar project on China. I looked into the women's rights case. China routinely blocks at least three or four women's rights sites, according to the Open Net Initiative. I looked at what issues are discussed on those sites. The one child policy. The high suicide rate amongst women in China. The issue of sexual diseases. All kinds of sensitive subjects for the government. Some nine of them. I wanted to see whether or not that kind of content was available on women's rights sites that aren't censored by China. Answer: yes.