

The .GOV Internet Archive: A Big Data Resource for Political Science

Emily Gade¹
John Wilkerson
University of Washington

Abstract

We introduce the Internet Archive's .GOV database, which includes more than 1 billion .gov webpage captures in a format that supports large-n systematic analyses. Understanding how to work with .GOV provides an introduction into the ever expanding world of big data. We offer instructions and tips for using .GOV, and illustrate its potential by exploring attention to several recent issues across U.S. federal government websites.²

¹Corresponding Author: ekgade@uw.edu

²We thank Vinay Goel of the Internet Archive, Altiscale, the University of Washington Data Science Incubator and Nitin Borwanker for their assistance. This work was supported in part by the National Science Foundation under Grant No. 1243917 (Division of Social and Economic Sciences, Directorate for Social, Behavioral Economic Sciences). Any opinions, findings, and conclusions or recommendations expressed are those of the authors and do not necessarily reflect the views of the National Science Foundation.

1 Introduction

“Big data” will transform social science research. By big data, we primarily mean datasets that are so large that they cannot be analyzed using traditional data processing techniques. However, big data is further distinguished by diverse types of information and the rapid accumulation of that information.³ We introduce one recently released big data resource, and discuss its promise along with potential pitfalls. For nearly 20 years, governments have used the web to share information and communicate with citizens and the world. .GOV is an archive of nearly two decades of content from .gov domains (US federal, state, local) organized into a database format that is nine times larger than the entire print content of the Library of Congress (90 terabytes, or 90,000 gigabytes).⁴ Big data resources like .GOV pose novel analytic challenges in terms how to access and analyze so much data. In addition to the difficulty posed by its size, big data is often messy. Additionally, .GOV is neither a complete nor a representative sample of government presence on the web across time.

1.1 The Internet Archive

In 1963, J.C.R. Licklider of the Advanced Research Projects Agency (ARPA) drafted a “Memorandum For Members and Affiliates of the *Intergalactic Computer Network*” (emphasis added). Subsequent discussions ultimately led to the creation of ARPANET in 1968. Soon after, major government departments and agencies were constructing their own “nets” (DOE and MFENet/HEPNet, NASA and SPAN). In 1989, Tim Berners-Lee proposed (among other things) using hypertext links to enable users to post and search for information on the internet, creating the World Wide Web. The first commercial contracts for managing network addresses were awarded in the early 1990s. In 1995, the internet was officially recognized by the Federal Networking Council, and Netscape Navigator, “the web browser for everyone,” went public.

In 1996, a non-profit organization, the Internet Archive (IA) assumed the ambitious task of documenting the public web. The current collection contains more than 450 billion webpage “captures” (downloads of URL linked pages and metadata) dating back to 1995. The best way to quickly appreciate what’s in the IA holdings is to visit the WayBack Machine website (archive.org/web), where specific historical website captures (e.g. the White House home page from Dec. 27, 1996) can be viewed.

2 .GOV: Government on the Internet

The Internet Archive also curates sub-collection: .GOV.⁵ .GOV contains approximately 1.1 billion page captures of URLs with a .gov suffix (from 1996 through Sept 30, 2013). At the federal

³For example, see this article on understanding Big Data: Sagioglu, S., Sinanc, D. (2013, May). Big data: A review. In Collaboration Technologies and Systems (CTS), 2013 International Conference on (pp. 42-47). IEEE. <https://xa.yimg.com/kq/groups/72986399/1585974627/name/06567202.pdf>

⁴In thinking about using the volume of the Library of Congress as a unit of measure, see: “A “Library of Congress” Worth of Data” by Leslie Johnston, April 25, 2012. <http://blogs.loc.gov/digitalpreservation/2012/04/a-library-of-congress-worth-of-data-its-all-in-how-you-define-it/>

⁵See the Internet Archive’s description of their sub-collections here: https://archive.org/details/additional_collections

level, this includes the official websites of elected officials, departments, agencies, consulates, embassies, USAID missions and much more.⁶ Whereas the Wayback Machine makes it possible to view date-specific individual websites, the .GOV collection can be used to investigate patterns across websites and over time.

.GOV offers four types of data from each webpage capture: the link data (the page URL and every other url/hyperlink found on the page); the parsed text of the page; the full content of the page (the text including html markup language; images; video files etc); and the CDX index file that is used to access the page via the Wayback Machine.

2.0.1 Messy Data

There is no way to download the entire content of the internet, or even a representative sample. The IA (as well as major search firms such as Google) capture content by “crawling” from one page to another. Starting from a limited number of “seed” URLs (web page addresses) a “bot” (software program) collects content from all of the URLs found on the originating page, then all of the URLs on those pages (etc.) until it encounters no more unique pages, or a user defined search constraint tells it to stop. This sequential process inevitably offers an incomplete snapshot of a constantly evolving World Wide Web. In 2008, the official Google blog reported that developers had collected 1 trillion unique URLs in a single concerted effort but also noted that “the number of pages out there is infinite.”⁷ Crawl results are also incomplete because webpages are sometimes located behind firewalls (the “dark web”), or include scripts that discourage bots from collecting content. The Internet Archive will also delete a website at the owner’s request. We have also discovered other limitations of the .GOV data that users should be aware of in designing projects.⁸

The quality of the Internet Archive also improves over time, both because of changes in the way the Web is used and because of changes in the way the Internet Archive conducted its crawls. Figure 1 displays how often the White House website was captured across four different years starting in 1997.⁹ The Wayback Machine indicates that `whitehouse.gov` was crawled just 3 times in 1997. In 2001, it was not crawled at all in the month of August and then hundreds of times in the three months following the terrorist attacks on September 11. In 2007, it was captured much more often - at least once a week. And in 2014, `whitehouse.gov` was captured at least once a day.


⁶.GOV also includes state and local websites that use the .gov suffix.

⁷See Google’s Official Blog (July 25, 2008) for discussion at “We knew the web was big...” <http://googleblog.blogspot.com/2008/07/we-knew-web-was-big.html>

⁸These are listed in the on-line Appendix.

⁹The graphs are copied from Wayback Machine search results for `whitehouse.gov`.


Figure 1: Frequency of whitehouse.gov crawls (selected years)



The most complete .GOV crawls occurred during three month time periods (Nov-Jan) of election years starting in 2004.¹⁰ Using congressional websites URLs as the seeds, the IA captured more government web presence than before. Figure 2 indicates spikes in unique .GOV URLs captured during election years. For example, the number triples from about 500 million to 1.5 billion between 2003 and 2004.

¹⁰According to Vinay Goel, senior data engineer at the IA, the Library of Congress contracted with the IA to systematically capture congressional websites during these time periods.

Figure 2: Total .GOV Unique URLs




Although .GOV is less than ideal as a data resource from a conventional social science perspective, there is no other option for investigating two decades of White House website content, or the content of millions of other pages of government website content. Importantly, because these crawls contain snapshots of each page, researchers could hypothetically examine language that agencies or individuals chose to remove - something scraping those same pages now could not provide. The challenge researchers face is finding the hidden gems in a resource that cannot be easily explored.

3 Big Data and Distributed Computing

.GOV is an excellent platform for learning about “big data” analysis techniques. The basic challenge is that the dataset is too large (90,000 gigabytes) to download and explore. Big data is stored and managed differently. Traditional databases (aka “structured data” are organized into neat rows and columns.¹¹ Big data projects rely on more flexible data storage processes where portions of the data are distributed across a cluster of computers. Each computer in the cluster is a node, and portions of the data are stored in “buckets” or “bins” within each node (see Figure 3). To access the data, researchers use special software to send simultaneous requests to the different nodes. The piecemeal results of these multiple queries are then recombined into a much smaller, single working file.

¹¹See Sagioglu Sinanc: <https://xa.yimg.com/kq/groups/72986399/1585974627/name/06567202.pdf>

Figure 3: Hadoop System for .GOV



3.1 Querying .GOV

The .GOV database is currently hosted on a Hadoop computing cluster operated by a commercial datacloud service, Altiscale (www.altiscale.com). Within the cluster housing .GOV, the data are distributed across nine separate “buckets.” Each bucket contains thousands of large (100mb) WARC (Web Archive Container) files (or ‘ARC’ files for earlier records). Each of these WARC files then contain thousands of individual webpage capture records. As mentioned, each capture record includes the parsed text, the URLs found on the page, the full content of the capture (including images and video files); and the CDX index file. The CDX file includes useful metadata about specific records that can be used to find and exclude particular records, such as the URL, timestamp, Content Digest, MIME type, HTTP Status Code, and the WARC file where it is located.

The data are accessed using Apache software programs. Apache Pig and Hive are SQL-based languages that can be used for basic data processing such as joining or merging files, searching for specific URLs, and more generally retrieving data of interest. Many Apache commands will be familiar to users with working knowledge of SQL, R or Python. To search all of the capture records in the .GOV database, one must write a query to search thousands WARC files across each of the nine buckets.

3.2 Obtaining a key and creating a workbench

Here we describe the big picture process of querying .GOV. In the next section, we present some preliminary findings using the parsed text data. The specific annotated scripts used to accomplish the latter can be found in the on-line Appendix.

Users must first gain access to the Altiscale computing cluster by requesting an “ssh” key (detailed on Altiscale’s website: <https://documentation.altiscale.com/getting-started> - you have to email accounts@altiscale.com to request a key). Each key owner is granted a local workbench (an Apache Work Station (AWS)) on the cluster that is similar to the “desktop” of a personal computer and contains the Apache software programs needed to query the database. About 20 gb of storage is also provided (the .GOV database is about 4500 times larger).

3.3 Writing scripts to extract information

1. Specifying what is to be collected

Apache Hive and Pig are used execute SQL queries. This can be done on the command line directly (there is no GUI option), but it is easier to write and store scripts on the workbench, and then write a command to execute them across the buckets of interest. For example, one can write an Apache Pig script that requests each parsed text file from a specified URL (e.g. whitehouse.gov), separates the parsed text fields (“date” “URL” “content” “title” etc.), searches each field for each record for a keyword or regular expression, and then counts how many times a match occurs. The full parsed text could also be downloaded in order to explore the content in more detail later. But with so much data (1.1 billion pages), such a collection can quickly become too large to export.

The functionality of Pig and Hive (like SQL) is limited. For example, Pig will return and count the webpage captures that contain a keyword (true/false) for a date range (e.g., per month/year), but it can’t compute the frequency of keyword mentions. To do more detailed or custom analysis, researchers can write user defined functions (UDFs) in Python. The Python script is stored on the workbench as a .py file and then called by the Pig script.¹²

Processing time is a major consideration. Even simple jobs such as keyword counts can take hours or even days to run over so much data. More computationally intensive methods, such as topic models may be impractical. The best way to discover whether a script is going to work and how long it will take is to test it on a subset of the data, such as on just one WARC file in one of the buckets. Running a complete job without testing it is likely to lead to many hours or days of waiting only to discover that it did not work. Linux “Screen” (already installed on the cluster) can then be used to run the script across the cluster remotely (so that your own computer can be used for other things).¹³

2. Providing instructions about where to search on the cluster

The CDX files provide guidance that makes it possible to limit queries to particular URLs, date ranges, WARC files etc.¹⁴ For queries than cannot be restricted in advance (e.g. the research objective is to identify all parsed text files that contain a particular keyword), breaking a job into steps can be more efficient. For example, the first query might identify and produce a list of all of the URLs that contain the keyword. The next query would focus on extracting

¹²See <https://wiki.apache.org/Pig/UDFManual> and <https://Pig.apache.org/docs/r0.11.0/udf.html>

¹³For instructions about how to use Screen see: <https://www.rackaid.com/blog/linux-screen-tutorial-and-how-to/#detach>

¹⁴Instructions for querying the CDX file can be found at <https://webarchive.jira.com/wiki/display/Iresearch/IA+-+GOV+dataset+-+Altiscale>

the relevant information (such as keyword and total word counts) from that more limited set of URLs.

3. Concatenating and exporting the results

Query results for each WARC or ARC file (containing thousands of captures) are stored separately on the cluster. Additional scripts must be written to concatenate them. Whether the results can be exported can also be calculated at this point.¹⁵ If they cannot, Apache Giraffe (<http://giraph.apache.org/>) is designed to facilitate analyses and graphing on the cluster.

4 Application: Government Attention to the Financial Crisis, Terrorism, and Climate Change

As a starting point to discovering what's in .GOV, we investigate keyword frequencies for three recent issues in American politics. We hope to observe patterns consistent with what is generally known about the issues, and perhaps more novel patterns that begin to illustrate the potential of this new data source.

4.1 Collecting the data

We first created a limited list of top level URLs (departments, agencies and political institutions) relevant to the three issues. For example, the regular expression “.house.gov” theoretically captures every webpage of every branch of the official website of the U.S. House of Representatives that the IA collected. This includes, among other things, every Representative's official website (e.g., pelosi.house.gov) and every House committee's website (e.g., agriculture.house.gov). Aggregating results during the collection process in this way means that we cannot pull out just the results for a particular member or committee's website. That would require a different query using more specific URLs.

We then counted keyword mentions on every subpage of that root URL. We first developed broad lists of keywords related to the three issues. After obtaining results, we created more refined lists by dropping terms that seemed problematic or were used less often (see online appendix). For example, we dropped “security” from the terrorism keyword list because it was too general (e.g., financial security). Running the query over all WARC files took about five days of processing time. All together, the results reported below are based on 8.3 billion keyword hits generated by searching about 600 billion words found on the parsed text pages of the specified URLs.

Focusing on raw counts of keywords gives more weight to larger domains. Any changes in attention to terrorism at the much larger State Department will swamp changes at the Bureau of Alcohol, Tobacco and Firearms (ATF). We focus on the proportion of attention given to the issue within an agency or political institution, by dividing the number of keyword hits by total website words. A proportion-based approach also does a better job of controlling for the

¹⁵Instructions for exporting documents from the Altiscale cluster can be found here: <http://documentation.altiscale.com/exporting-data-from-your-altiscale-cluster>

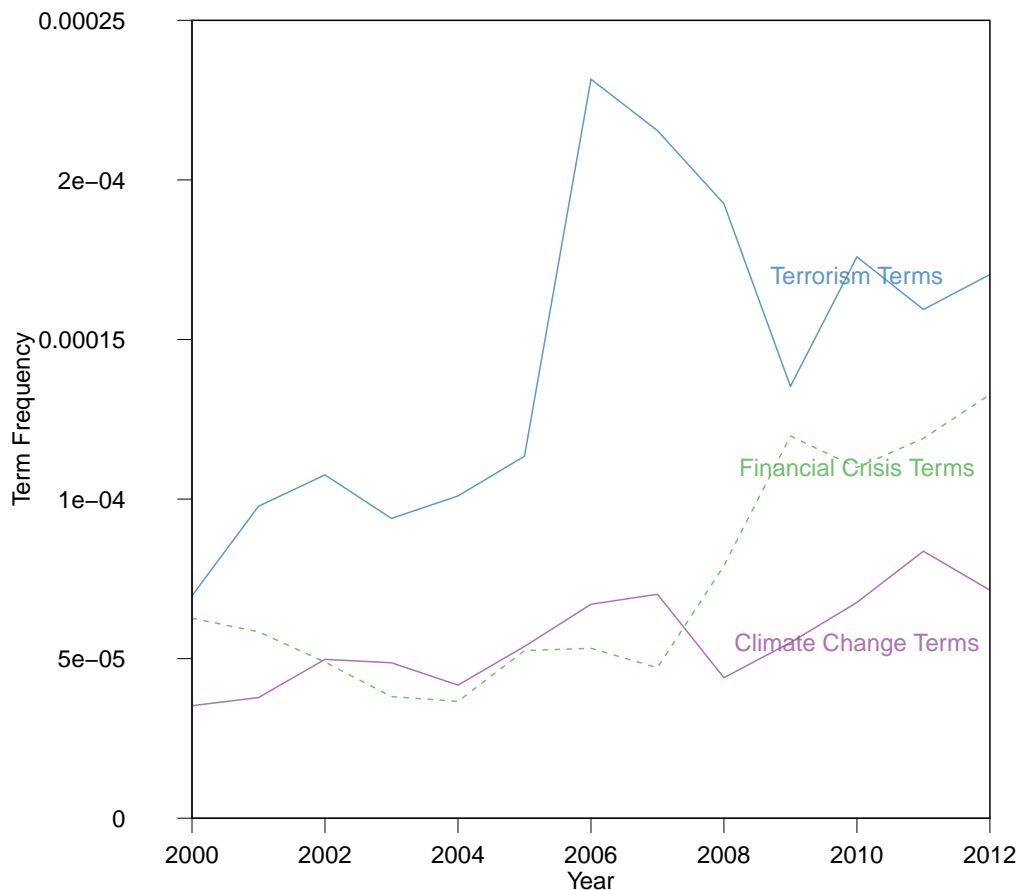
expanding size of government web presence.

4.1.1 Overall Trends

One way to begin to assess the validity of using website content to study political attention is to ask whether changes in content correlate with known events. For each issue in Figure 4, we identified the URLs (federal government organizations) thought to play a role on the issue (see Appendix II for these lists and URLs). The graphs then report average proportions of attention across these URLs.¹⁶ Financial crisis term usage (as a proportion of all terms) spikes upward in 2007-08 as expected (and also in 2001 when there was another stock market decline). Attention to terrorism similarly increases after 9/11/2001, but government-wide attention to terrorism increases most dramatically from 2005 to 2006. Institutionalization is almost certainly part of the explanation. We are capturing attention to terrorism (relative to other issues) on the websites of government agencies. The Department of Homeland Security was not created until 2003 and one of the purposes of its creation was to re-orient the missions of existing agencies (such as FEMA) towards preventing and responding to terrorism. In addition, while 9/11 was an important focusing event for the US, terrorism worldwide continued to increase post 9/11. As we might expect, there is no evidence of equivalent shocks for climate change.

¹⁶The different proportions for the different issues are not comparable because they are dependent on the keyword lists.

Figure 4: Issue Attention Across .GOV




Diffusion of Attention to Terrorism

Political scientists have long been interested in how “focusing events” impact political attention (Birkland 1998). Many studies have examined the impact of 9/11 on the organization and activities of specific government agencies and departments. Here, we ask how attention to terrorism spread across government departments and agencies. Entropy is a measure disorder that is frequently used to study the dispersion of political attention (Boydston et al. 2014). Figure 5 confirms that attention to terrorism in the federal government became more dispersed post 9/11.¹⁷

¹⁷Our measure is based on the proportion of domain content for 23 departments and agencies, where entropy is based on each domain’s proportion of the sum of all agencies’ proportions.

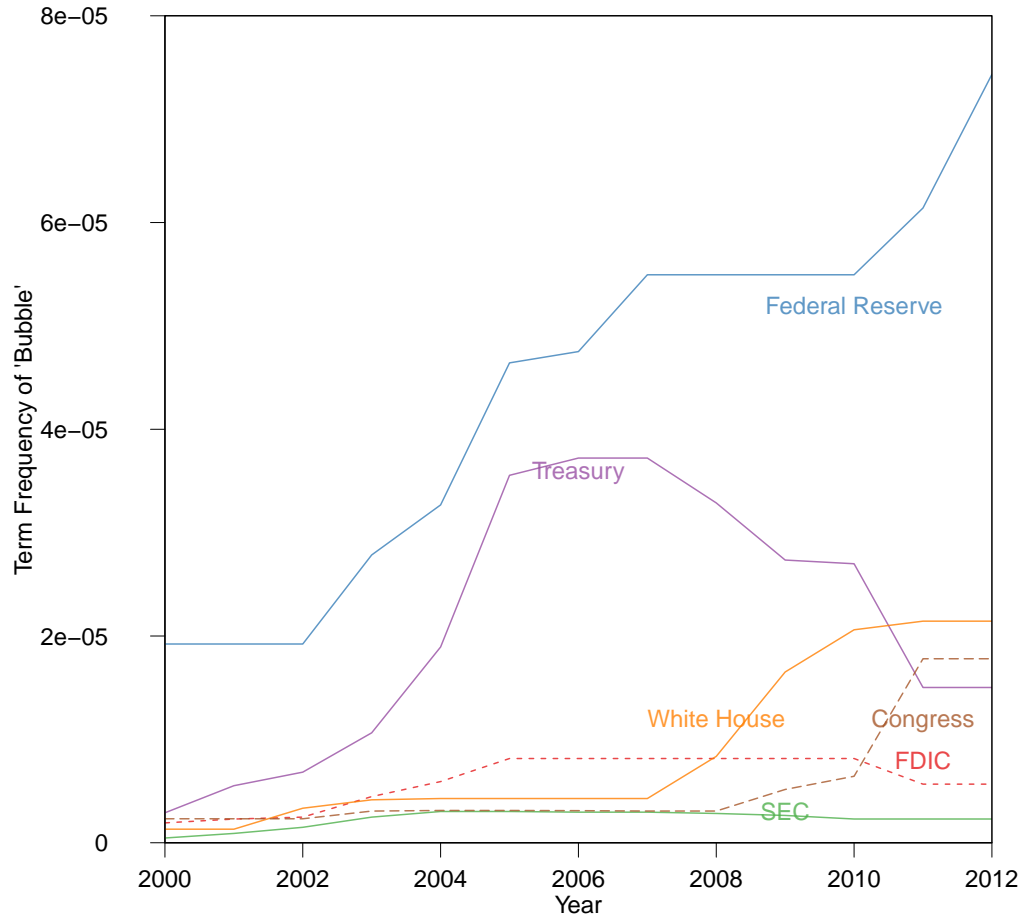
Figure 5: Diffusion in Attention Across .GOV



A Financial “Bubble”?

One of the questions raised in congressional hearings after the 2007-08 financial crisis was whether it could have been anticipated and averted. A related question is whether government agencies saw it coming. As an historical archive, .GOV may provide some clues. Here we simply examine “bubble” mentions across organizations (as a proportion of total website words). Figure 6 indicates that references to bubbles spike in the elected branches after the meltdown, whereas bubble mentions at the four agencies most responsible for the economy increase 2-3 years ahead of the crisis. Bubbles also see increased attention at the Federal Reserve before the stock market sell-off in 2001.

Figure 6: Attention to Financial Crisis Across .GOV

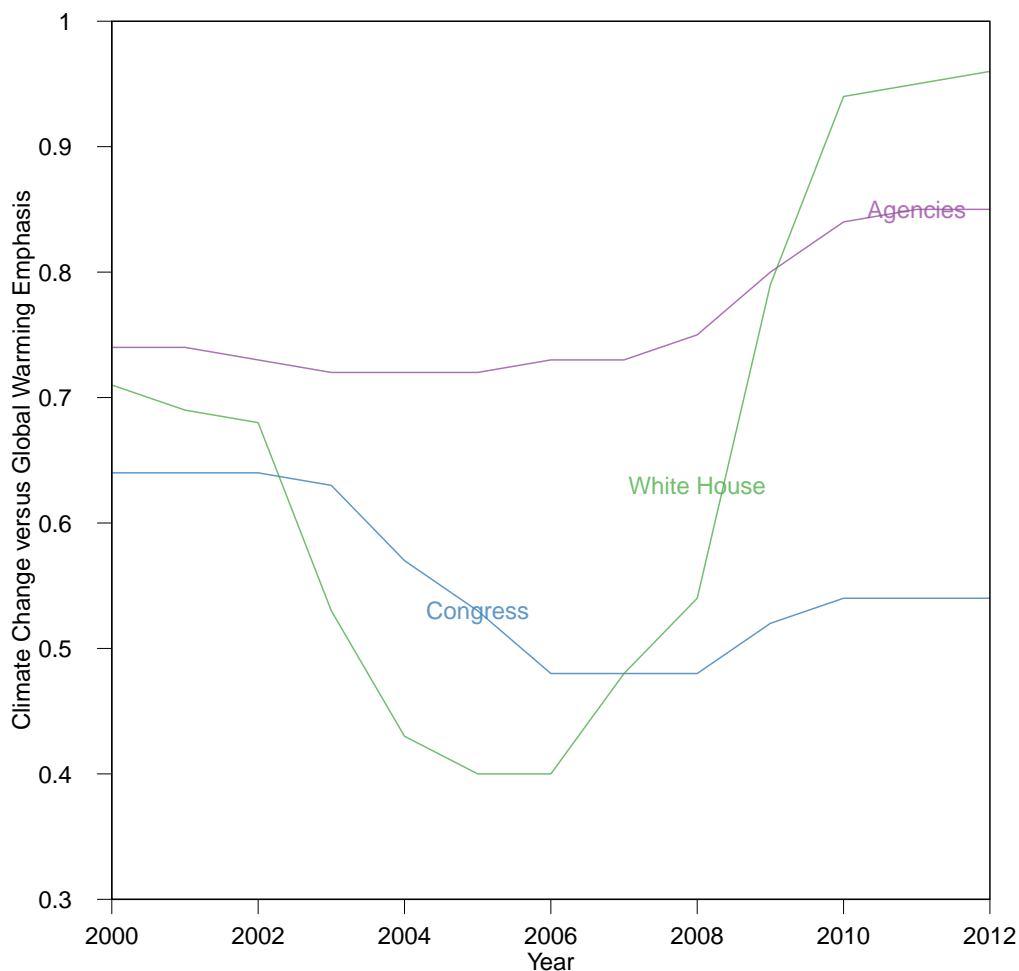


Framing Climate Change

Early in President G.W. Bush's first term of office, pollster Frank Luntz advised Republicans to talk about "climate change" rather than "global warming" because focus groups saw the latter as more of a threat (Leiserowitz et al. 2014, 7). Subsequent academic research also found that the public is somewhat more likely to support action to address global warming. However, it seems as though conservatives also spend much of their time ridiculing global warming. Recently Senator James Inhofe (R-OK) brought a snowball to the Senate floor to question scientists' claims that 2016 was one of the warmest years on record (Bump 2015). If conservatives have discredited global warming in the eyes of the public, then proponents of climate action may have less incentive to use that frame.

In Figure 7, values above .50 indicate that climate change mentions are more common than global warming mentions. "Agencies" refers to the average emphasis on climate change for four agencies with central roles (the EPA, NSF, NOAA, and NASA). According to Figure 7, scientific agencies have always emphasized climate change over global warming, with climate change increasingly favored in recent years. For the elected branches, the patterns are more variable and seem to support the notion that conservatives control the global warming frame. In Congress, global warming has been a more popular frame during periods of Republican control (2001-2008; 2011-2013) and has been used more often over time. The patterns for the White House do not support what Luntz advised. Global warming receives more attention than climate change for most of the years of the Bush administration. The Obama administration, in contrast, has gone all in for climate change. Although preliminary, these results do suggest that conservatives have defanged what was once the most effective frame for winning public support for climate action.

Figure 7: Attention To Climate Change Across .GOV



5 Conclusion

Accessing this big data resource requires new skills and a new mindset. In terms of skills, we hope that our description of the process and working scripts lower the bar. In terms of mindset, political scientists working with statistical methods are used to immediate results. Exploring .GOV in this way is not an option (the Wayback Machine is probably the best way to get a sense of what's in .GOV, and it can take days or even weeks to run a query). On the other hand, .GOV contains insights available nowhere else. Although the current database has important limitations, the Internet Archive recently embarked on a collaboration with many partners to scrape all federal government agencies as completely as possible prior the end of the Obama administration. If this effort is successful and if similar efforts follow in subsequent years, .GOV will be an even more valuable resource for investigating a wide range of questions about the federal bureaucracy and federal programs.

References

- “Bash Shell Basic Commands.” *GNU Software*. <http://www.gnu.org/software/bash/manual/bash.pdf>
- Boydston, A. E., Bevan, S. and Thomas, H. F. (2014), “The Importance of Attention Diversity and How to Measure It.” *Policy Studies Journal*, 42: 173–196. doi: 10.1111/psj.12055
- Bump, P. “Jim Inhofe’s Snowball Has Disproven Climate Change Once and for All.” Washington Post. The Washington Post, 26 Feb. 2015. Web. 28 June 2016.
- Birkland, T. A. (1998). “Focusing events, mobilization, and agenda setting. *Journal of public policy*.” 18(01), 53-74.
- Edwards, J., McCurley, K. S., and Tomlin, J. A. (2001). “An adaptive model for optimizing performance of an incremental web crawler”. *Tenth Conference on World Wide Web* (Hong Kong: Elsevier Science): 106–113.
- “The History of the Internet.” *The Internet Society*. <http://www.internetsociety.org/internet/what-internet/history-internet/brief-history-internet>
- “The Internet Archive.” *Internet Archive*. <https://archive.org/>
- Kahn, R. (1972). “Communications Principles for Operating Systems.” *Internal BBN memorandum*.
- Leiner et al. “Brief History of the Internet.” http://www.internetsociety.org/sites/default/files/Brief_History_of_the_Internet.pdf
- Leiserowitz, A. WHAT’S IN A NAME? GLOBAL WARMING VERSUS CLIMATE CHANGE. Rep. Yale Project on Climate Change Communication, May 2014. Web. 28 June 2016.
- Licklider, J. C. (1963). “Memorandum for members and affiliates of the intergalactic computer network.” M. a. A. ot IC Network (Ed.). Washington DC: KurzweilAI. ne.
- Najork, M and J. L. Wiener. (2001). “Breadth-first crawling yields high-quality pages.” *Tenth Conference on World Wide Web*, (Hong Kong: Elsevier Science): 114–118.
- “Pig Manual.” *Apache Systems* https://pig.apache.org/docs/r0.7.0/piglating_ref1.html
- “The Rise of 3G.” *THE WORLD IN 2010*. International Telecommunication Union (ITU). <www.itu.int/ITU-D/ict/material/FactsFigures2010.pdf>.
- Sagiroglu, S., Sinanc, D. (2013, May). “Big data: A review.” In *Collaboration Technologies and Systems (CTS)*, 2013 International Conference on (pp. 42-47). IEEE. <https://x.a.ying.com/kq/groups/72986399/1585974627/name/06567202.pdf>

“A “ssh” key (Secure Shell)” (2006). <http://tools.ietf.org/html/rfc4252>

Vance, A. (2009). “Hadoop, a Free Software Program, Finds Uses Beyond Search”. *The New York Times*.