

# Media Cloud: Massive Open Source Collection of Global News on the Open Web

Hal Roberts,<sup>1,2</sup> Rahul Bhargava,<sup>3,2</sup> Linas Valiukas,<sup>2</sup> Dennis Jen,<sup>2</sup> Momin M. Malik,<sup>1</sup>  
Cindy Sherman Bishop,<sup>2</sup> Emily B. Ndulue,<sup>2</sup> Aashka Dave,<sup>4</sup> Justin Clark,<sup>1</sup> Bruce Etling,<sup>1,2</sup>  
Robert Faris,<sup>5,1,2</sup> Anushka Shah,<sup>6,7</sup> Jasmin Rubinovitz,<sup>7\*</sup> Alexis Hope,<sup>7</sup> Catherine D’Ignazio,<sup>8</sup>  
Fernando Bermejo,<sup>2,1</sup> Yochai Benkler,<sup>9,1,2</sup> Ethan Zuckerman<sup>10,1,2</sup>

<sup>1</sup> Berkman Klein Center for Internet & Society, Harvard University

<sup>2</sup> Media Cloud

<sup>3</sup> College of Arts, Media and Design, Northeastern University

<sup>4</sup> School of Information & Library Science, University of North Carolina at Chapel Hill

<sup>5</sup> Shorenstein Center on Media, Politics and Public Policy, Harvard University

<sup>6</sup> Civic Studios

<sup>7</sup> Media Lab, Massachusetts Institute of Technology

<sup>8</sup> Department of Urban Studies and Planning, Massachusetts Institute of Technology

<sup>9</sup> Harvard Law School, Harvard University

<sup>10</sup> School of Public Policy, University of Massachusetts Amherst

hroberts@cyber.harvard.edu, r.bhargava@northeastern.edu, {linas, dennis}@mediacloud.org, momalik@cyber.harvard.edu, {cindy, ebndulue}@mediacloud.org, aashkad@email.unc.edu, {jclark, betling, rfaris}@cyber.harvard.edu, {anushka18, jasrub}@gmail.com, {ahope, dignazio}@mit.edu, fernando@mediacloud.org, ybenkler@law.harvard.edu, ethanz@umass.edu

## Abstract

We present the first full description of Media Cloud, an open source platform based on crawling hyperlink structure in operation for over 10 years, that for many uses will be the best way to collect data for studying the media ecosystem on the open web. We document the key choices behind what data Media Cloud collects and stores, how it processes and organizes these data, and its open API access as well as user-facing tools. We also highlight the strengths and limitations of the Media Cloud collection strategy compared to relevant alternatives. We give an overview of two sample datasets generated using Media Cloud and discuss how researchers can use the platform to create their own datasets.

## Introduction

Researchers have long studied the production and dissemination of news to understand its impact on democracy, beliefs and behaviours. These efforts have included large scale projects to manually review news content (Kayser 1953), theories of how “themes” of news move subjects into and out-of public favor (Fishman 1978), sociological accounts of newsroom decision making (Gans 2004/1979), analysis of news agendas to help identify areas ripe for public health messaging (Dorfman 2003), and more. While the digitization of news has broadly made many sorts of research newly possible, the volume of content and newer trends toward proprietary platforms have significantly increased the barriers for researchers hoping to build on the scholarly traditions of studying attention, representation, influence, and language in online news.

\*Now at Google.

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

In addition, the democratization of content authoring has significantly expanded the number of different news sources reporting about any topic or geographic area. The process of discovering which online media sources to study has become more involved; how does one identify the set of media outlets that exist and decide which are influential? Furthermore, once a set of media sources has been created, collecting their content introduces many technological difficulties.

Many researchers rely on commercial data stores to solve these problems of media source discovery and content collection, but those stores are often limited by cost, timeliness, geographic focus, linguistic diversity, and content restrictions. Analysis falling outside of those stores requires significant technical and substantive effort.

To address these concerns we created the Media Cloud project, an open source data corpus and suite of web-based analytic tools that support research into open web media ecosystems across the globe. This paper documents the architecture and datasets underlying Media Cloud. We provide some history of the project origins, document the data collection and metadata extraction processes, introduce the datasets Media Cloud allows its users to create, review methods of accessing those datasets, and summarize research the datasets allow. Researchers can create datasets via the API or web-based research support tools to study attention, representation, influence, and language about subjects of their own choosing. They may also download and run the Media Cloud code to create their own instances of the platform.

## About Media Cloud

Since 2008, Media Cloud has collected a total of over 1.7 billion stories. As described below, these data are avail-

able to the public via an open source code base available on GitHub (<https://github.com/mediacloud>), a suite of free web tools (<https://tools.mediacloud.org>), and an extensive open API (<https://mediacloud.org/support>). Our project has a growing number of open-source contributors and collaborators outside of our core team. Media Cloud is limited by copyright restrictions from sharing the full text of news articles with external partners; instead, we surface metadata, data about text content of the documents, and URLs of documents indexed. After reviewing the history of Media Cloud, this section documents the data available in the underlying platform and how those data are generated.

## History

Work on Media Cloud began at Harvard's Berkman Klein Center for Internet & Society in 2007 as a project to examine the influence of blogging on mainstream media coverage. Berkman researchers were using the blog search engine Technorati to identify popular topics in blogs and sought an up-to-date, open source alternative to LexisNexis to identify topical coverage in news media.

Earlier software had been developed at Berkman to track a subset of topics in news media through keywords—the names of individual countries—to support the “Global Attention Profiles” (GAP) research project (Zuckerman 2003). While the GAP tool tracked hundreds of specific keywords by scraping the search engines of select news websites, Media Cloud was designed as a general solution, allowing researchers to track keywords across websites with RSS feeds. GAP maintained only a count of how many stories mentioned given keywords on a specific day, but Media Cloud provided a search engine, fully indexing the content of tracked pages.

An internal debate within the Berkman Klein Center pushed the initial development of Media Cloud. Yochai Benkler's work on commons-based peer production suggested that permissionless online spaces for publication—blogs, online-native publications and wikis—could be more inclusive and diverse than traditional news media (Benkler 2006). Ethan Zuckerman argued that digital media was subject to strong homophily effects and would likely be less diverse than traditional news media (Zuckerman 2013). Media Cloud illuminated this debate with data: would ‘digitally native’ media show greater topical and geographic diversity than traditional media? A 2012 Ford Foundation grant boosted the project, focusing on politically-motivated misinformation and the network architecture of agenda-setting in the public sphere (Benkler et al. 2013).

Media Cloud became useful not only for the analysis of the architecture of large scale media ecosystems, but for analysis of the fine-grained flow of specific terms and ideas online over time. This affordance was facilitated by the system's ability to track story publication dates, as seen in work to understand the spread of news around Trayvon Martin's death in 2012 (Graeff, Stempeck, and Zuckerman 2014).

Open and free access to its code, data, and tools has been a foundational principle for the project from the start. To make the platform more widely accessible, a team at the MIT Media Lab built interfaces to the Media Cloud datasets, usable

by activists and journalists as well as researchers comfortable writing their own code to interact with the Media Cloud API. This expanded access led to calls to broaden the subject matter covered by Media Cloud, including the establishment of collections of news media in multiple countries outside the English-speaking world, beginning with collection of data about Brazil (Gonçalves 2014).

Since those beginnings, our global work has expanded significantly, producing a dataset of geographic collections of media sources that cover the large majority of countries in the world. Much of this diversification effort has been driven by funding partnerships with global grant-making non-profits. The Bill and Melinda Gates Foundation, the Ford Foundation, the MacArthur Foundation, the John H. and James L. Knight Foundation, the Robert Wood Johnson Foundation, the Open Society Foundations and others are heavily invested in assessing and influencing global media narratives on the topics they focus on. They provide significant funding to nonprofits specifically to impact local media narratives around issues such as contraceptive and family planning in Nigeria and sanitation in India. These foundations' support, and our own global research partnerships, have helped Media Cloud grow our international media coverage. The breadth of our collection differentiates the Media Cloud datasets from many commercial offerings.

In addition to increasing the usability of the tools and the breadth of content coverage, the project has greatly expanded the kinds of data collected and the analysis performed on that data, in particular within the Topic Mapper tool. Led by a team at the Berkman Klein Center for Internet & Society, the project grew from its initial scope of spidering and analyzing topics that were keyword-based and web-only, across a limited scope of about 10 thousand stories, to now managing web- and social media-seeded topics spanning the U.S. national media ecosystem and consisting of tens of millions of stories.

## Data Collection

Media Cloud discovers and collects data through two systems: a syndicated feed crawler that ingests roughly a million stories per day from about 60 thousand media sources, and a topic-specific engine that recursively spiders hyperlinks to discover additional topic-specific content, capable of creating topic-specific datasets of more than 10 million stories. An overview of the system is given in figure (1).

**Regular Data Ingest** The syndicated feed crawler ingests RSS, RDF, and Atom feeds associated with specific media sources. A media source is any publisher of online content, such as The New York Times or Gateway Pundit. Syndicated feeds for each media source are discovered through code that identifies feeds by: spidering two levels deep from the home page of the media source, finding any links that match a pattern indicating likely feeds (including terms like ‘RSS’ or ‘XML’), downloading each of those URLs, and attempting to parse the resulting content as a syndicated feed to determine whether it has valid content.<sup>1</sup> All such validated feeds

<sup>1</sup>[https://github.com/mitmedialab/feed\\_seeker](https://github.com/mitmedialab/feed_seeker)




Figure 1: A schematic diagram of the Media Cloud topic-specific dataflow.

are added to the given media source. Also, a human “collections curator” regularly reviews collections and sources and manually adds missing feeds.

Each feed is downloaded following a progressive backoff strategy. After a new story appears in a given feed, that feed is downloaded again within five minutes. If no new story is available, that feed is downloaded again in ten minutes, then twenty minutes, and so on, backing off on the frequency of update until the feed is only downloaded once a week. Once a new story is discovered in a given feed, its interval is reset to five minutes.

The system parses the items from each downloaded feed and either: 1) matches each item to an existing story in the database by URL, GUID, or title, or 2) adds the unmatched item as a new story with the standard metadata provided by the feed (URL, GUID, title, publication date, description, and media source). To match by URL or GUID, the system uses a lossy URL normalization routine.<sup>2</sup> To match by title, the system uses a normalization scheme based on the longest part of a title as separated by characters such as ‘:’, ‘|’, or ‘\_’.

The system downloads the HTML content for each story, extracts the substantive text using the publicly available python-readability package,<sup>3</sup> parses the extracted text into individual sentences, removes duplicate sentences within the same media source and calendar day, and then stores the remaining concatenated sentences as the story text. The system then guesses the language of the story using the cld2 library.<sup>4</sup> That story text is stored along with the language and other story metadata in a database to make it available for search through the platform API and various web tools. The system can also download and generate story text for podcasts (transcribed via the Google Speech-to-Text API). Finally, for each English story, the system also produces geotagging information using disambiguated CLIFF-CAVIN (D’Ignazio et al. 2014), U.S. news subject topics based on New York Times topics modeled using Google News word2vec (Rubinovitz 2017), and entities generated by Stanford CoreNLP (Finkel, Grenager, and Manning 2005).

<sup>2</sup>More details available at <https://mediacloud.org/news/2020/7/14/tech-brief-how-we-deduplicate-content>

<sup>3</sup><https://github.com/buriy/python-readability>

<sup>4</sup><https://github.com/CLD2Owners/cld2>

As a principle, Media Cloud only downloads content that is freely available on the web. The vast majority of media sources make their content freely available to bot systems like Media Cloud, including even paywalled sites like those of the New York Times and the Washington Post. For the handful of sites that do not allow our bot access to full content, we index the content available. For example, the Wall Street Journal website provides the first paragraph of content, which is what we index. Likewise, many academic journals only provide open access to an abstract, which is what we index.

**Topic-specific Data Ingest and Spidering** Media Cloud’s Topic Mapper (<https://topics.mediacloud.org/>) operates independently from the feed collection system. A topic initially consists of a set of stories defined by a seed query, a set of media sources, and a date range. Topic Mapper collects stories for specific topics by searching the Media Cloud index for this initial batch of stories, and spiders out on the open web from hyperlinks in those stories to find other relevant stories. The resulting aggregate set of stories is added to the topic and acts as the seed set of URLs for the further spidering process. The topic spider uses the same normalized URL and title matching as the syndicated feed crawler to match spidered URLs to existing stories within the database and downloads unmatched URLs. This is shown in figure (1).

Whether a matching story is found in the database or the URL is novel and then downloaded, the content for each discovered story is matched against the seed query. The story is added to the topic only if it matches the seed query. For new stories, the system guesses a date using the story HTML,<sup>5</sup> a media source based on the URL hostname, and a title based on the HTML. The system then processes the story through the same pipeline as feed-discovered stories. For each story added to the topic, the system parses hyperlinks from the extracted HTML of the story and adds them to the spidering queue. This spidering process is repeated for 15 rounds.

The Topic Mapper can also seed topics with URLs through searches of social media platforms. For each supported platform, the system searches for all posts matching a given query, parses the hyperlinks out of the posts, and adds those URLs to the topic seed URLs. The system currently

<sup>5</sup>[https://github.com/mediacloud/date\\_guesser](https://github.com/mediacloud/date_guesser)

Field	Search for	Example
story_id	a specific story by its unique Media Cloud id	story_id:12345
media_id	stories from a specific media source	media_id:12345
publish_date	stories published on a specific date and time	publish_date:2018-04-17 13:24:12
publish_day	stories published on a specific date	publish_day:2018-04-17 00:00:00
publish_week	stories published during a specific week	
publish_month	stories published during a specific month	
publish_year	stories published during a specific year	
tags_id_stories	stories with a specific metadata tag	tags_id_stories:9362290 (“Vladimir Putin” tag)
tags_id_media	stories from a source with a specific metadata tag	tags_id_media:32018496 (“digitally native source” tag)
timespans_id	stories from a specific timespan from a Topic	timespans_id:74322
language	text in a language, by ISO-639-1 two-letter code	language:en

Table 1: Fields included as indices in the Media Cloud database. See the online appendix in <https://arxiv.org/abs/2104.03702> for a description of data schema.

uses Brandwatch to search Twitter, Pushshift to search Reddit (Baumgartner et al. 2020) and verified Twitter accounts, and CSV uploads to support user-supplied platform searches (such as content downloaded from CrowdTangle).

The Topic Mapper uses information from this discovery process to create influence metrics for stories and media. During the spidering process, all hyperlinks encountered between stories are made available in topic datasets. Those hyperlinks are also used to generate a media in-link count for each story and for each media source, where the media in-link count for a story is defined as the number of unique media sources linking to that story. Similarly for media sources, the media in-link count (for a source) is the number of unique media sources linking to any story within that media source. Those media in-link counts allow researchers to focus on the stories that have received the most attention in the link economy. The hyperlink data are also used to generate network visualizations of the media sources involved in a given topic, with the media sources as nodes and the hyperlinks as edges. The system also retrieves a Facebook share count for each story from the Facebook API and then sets the Facebook share count of each media source to be the sum of the source’s story Facebook share count.

Any given topic can be analyzed as a subset of stories by specifying a query (as described below) that can restrict stories by text, language, or publication. For instance, an “immigration” subtopic of an “election topic” would include only stories mentioning immigration within the larger election topic. Likewise, every topic can be analyzed within distinct weekly and monthly timespans, each of which consists of all stories published within a given date range or linked to by a story within the date range. And for topics that include a social media seed query, the system generates a subtopic that uses URL co-sharing by users as a replacement for hyperlinks for the network analysis (allowing, for instance, a network visualization of media sources that clusters together sources shared by the same users). The system also generates post, author, and channel counts for each of these social media seed queries. So a topic with a Reddit seed query will include a subtopic that allows analysis of co-sharing of URLs by Reddit users and also counts of the unique posts, channels, and authors that included each URL returned by the seed query. All of these modes of analysis are included

within each topic’s downloadable dataset.

## Data Access

Media Cloud runs on top of a scalable, containerized architecture currently consisting of around 80 docker containers running roughly 550 workers across a pool of servers. The core database is a PostgreSQL database that is responsible for storing nearly all of the system metadata. All story content is also indexed in a cluster of Apache Solr shards, which provides rapid text searching for the system. Below we summarize the multiple points of data access that this architecture allows for.

**Query Language** Queries are authored in a customized version of the basic Boolean query language supported by the Solr search engine. This allows for off-the-shelf inclusion of operators such as AND and OR, negation, proximity, word stemming (in multiple languages), word frequency thresholds, quoted strings, and more. Queries that are the length of long paragraphs are common for big research projects. Researchers can also filter by the fields included as indices in the database, shown in table (1).

We have experimented with more advanced query tools such as query expansion and combining Boolean queries with trained models. These approaches are still deployed as one-offs, and are not yet available as general purpose tools within Media Cloud.

**API** All of the data described above are available through an open API, which includes over 65 distinct end points. The public web tools, which are the primary research tools used by the Media Cloud team for its own research, are implemented entirely on top of this API, so all functionality in the front end tools is available via the API. The API is available both directly as a REST API with extensive documentation<sup>6</sup> and as a Python client we maintain and support.<sup>7</sup>

To support and ease adoption of API-based access for data science researchers, we have created a set of Jupyter Notebooks in the Python programming language. These notebooks provide examples of common operations such as list-

<sup>6</sup>[https://github.com/mediacloud/backend/blob/release/doc/api\\_2.0\\_spec](https://github.com/mediacloud/backend/blob/release/doc/api_2.0_spec)

<sup>7</sup><https://pypi.org/project/mediacloud/>

ing top words used in stories matching a query, or accessing a topic dataset (Bhargava, Dave, and Papakyriakopoulos 2020).<sup>8</sup>

**Web-based Research Tools** For those without programming experience, or those who simply want the convenience of a web-based research platform, we provide access to data through two analysis-focused web apps (Explorer and Topic Mapper), in addition to a website that functions as a library and management console for the media sources included in Media Cloud (Source Manager).<sup>9</sup> For many researchers, these three apps are the only necessary access points to the Media Cloud dataset and its capabilities.

Both Explorer and Topic Mapper enable researchers to investigate a sub-corpus of news content defined through a Boolean search including query text, media sources or collections, and a date range. Both apps present tools for analyzing attention (amount of coverage given to a keyword or topic), language (most frequently used terms), and representation (frequently mentioned people, organizations, and countries). Topic Mapper includes a deeper set of influence and linking metrics and research support features, as described previously.

Explorer (<https://explorer.mediacloud.org/>) queries return in real time, allowing for rapid iteration on query construction. Researchers can use Explorer by itself to glean quick insights, and can use Topic Mapper (<https://topics.mediacloud.org/>) when asking deeper questions on influence: which media producers are doing most of the reporting on a topic, as well as the amount of referencing, citing, or sharing of the topic’s media on social platforms by the general population. In addition to direct access to datasets via the API, the web tools provide links to CSV datasets underlying each visualization.

Source Manager (<https://sources.mediacloud.org/>) provides researchers with information about individual news sources, and collections of news sources, within our system. Collection patterns, ingest rates, and other background information about the data contained are exposed to researchers. Critically, Source Manager tracks collection quality, so a researcher knows if data are incomplete for any range within the time period she is studying.

## Related Work

There are three groups of platforms or tools that are similar in scope to Media Cloud. The first are news databases. Of these, the most comparable are LexisNexis, Google News, and mediastack. The second are event databases, of which the most comparable are GDELT and Event Registry. The third are crawler tools, of which the most comparable is IssueCrawler. We compare Media Cloud to each group in turn.

LexisNexis is the most prominent commercial news database, providing archives of thousands of media sources which has made it the go-to tool for social science media research (Deacon 2007). It differs from Media Cloud

in three important ways. The first is cost and accessibility. Media Cloud is free and open source, whereas LexisNexis is both expensive and proprietary. While some universities have subscriptions to LexisNexis, this may not include bulk API access without which a user must laboriously go through the web interface, selecting up to ten stories per page and downloading stories in batches of up to 100.

A second difference is coverage. LexisNexis is exhaustive for many sources, since it often has direct relationships with publishers. Particularly noteworthy is that it is a record of print archives in addition to online content, and it allows access to the full text content (images, graphics, audio, and video are not included) of media sources that may be behind paywalls or otherwise not on the open web. But it focuses almost exclusively on formal sources (professional media outlets); apart from a few blogs, informal sources are largely absent. Sites that do not publish in a news format, such as the World Health Organization and Wikipedia, are not covered at all. The third key difference is in story discovery. LexisNexis has no ability to go beyond the sources it covers, whereas Media Cloud uses spidering to iteratively follow hyperlinks, including to non-news sites such as Wikipedia.

The lack of public documentation of LexisNexis also makes it difficult to freely compare to other tools. It claims coverage of 80 thousand sources, in 75 languages covering over 100 countries,<sup>10</sup> but further details as well as documentation of its API is not publicly available. Even if formal comparison is difficult without access, we emphasize the difference in use cases: for an authoritative search of formal news sources, LexisNexis is superior. Media Cloud is instead built for studying the media ecosystem on the open web, including hyperlink structure and non-formal sources.

A more recent competitor is Google News (now just “News API”). A 2008 study (Weaver and Bimber 2008) compared Google News and LexisNexis, finding inter-database agreement ranging from 29% to 83%, mostly due to LexisNexis excluding wire services. Today, News API boasts of access to 50 thousand sources.<sup>11</sup> It offers a free version, although with the significant restriction of only being able to search for content in the past three months. The commercial options are better but still limited: coverage there is in the past three years. News API focuses on providing streaming live, top headlines: for this it is superior to Media Cloud, although it is not an appropriate tool for research that seeks to consider older sources or collect more sources via snowball sampling.

Lastly among news databases, mediastack (<https://mediastack.com/>) is a commercial product from the company apilayer, although it does allow for 500 free requests per month. Commercial options start from \$25/month for 10 thousand calls. mediastack covers “general news, business news, entertainment news, health news, science news, sports news and technology news” with both a live stream (updated every minute, and for which “each of the sources

<sup>8</sup><https://github.com/mediacloud/api-tutorial-notebooks>

<sup>9</sup>For screenshots, see online appendix in <https://arxiv.org/abs/2104.03702>.

<sup>10</sup><https://internationalsales.lexisnexis.com/nexis-daas/>; archived snapshot at <https://perma.cc/XC5C-3XLS>

<sup>11</sup><https://newsapi.org/>; archived snapshot at <https://perma.cc/P63U-D8BJ>

used is monitored closely and around the clock for technical or content anomalies”) and a historical news data REST API, although without any description of how long this historical coverage extends or how it is collected (e.g. as an archive of its live stream, or something more). It does not have documentation of its collection process nor any academic papers we could find describing or making use of it, and only says that it covers over 7.5 thousand sources over more than 50 countries and in 13 languages. The countries are publicly listed (<https://mediastack.com/sources>), although the sources are not, which makes it difficult to compare to other news databases.

Next, there are event databases (Wang et al. 2016). First, GDELT, the Global Database of Events, Language and Tone (Leetaru and Schrodt 2013), also makes use of news stories on the open web (unlike other event databases, such as ICEWS), and is free and run from within academia. GDELT stores full-text data in some cases, but mainly focuses on collecting metadata. However, it is incompletely documented, and has received considerable criticism (Ward et al. 2013; Weller and McCubbins 2014; Kwak and An 2014a,b, 2016; Wang et al. 2016) around what its collection contains, with critics citing problems with duplication, lack of coverage of non-Western news sources, and dictionary-based pattern matching that makes it unable to make use of words not in its dictionary. It also does not include any spidering functionality for following hyperlinks beyond the sources it tracks.

Event Registry (<https://eventregistry.org/>) was originally an academic project (Leban et al. 2014; Rupnik et al. 2017) at the Institut Jožef Stefan (IJS) in Slovenia. However, in 2017 it was spun off into a commercial entity (Leban 2020), with the original tool no longer available. Recent academic papers published with Event Registry (e.g., Chołoniowski et al. 2019; Sitar, Mladenić, and Erjavec 2020) have been from IJS authors. It may have been an option for academic research until as late as 2019 (Fronzetti Colladon and Grippa 2020), but at the time of writing the site did not list any academic or free access options.

The paper initially describing the Event Registry (Leban et al. 2014) focused on *events* as the end-goal, with news stories treated as a measurement instrument rather than the construct of interest. The tool described there sought to identify underlying events from a collection of articles using a clustering approach, and then extracted event metadata (location, date, parties involved) using custom named entity recognition and disambiguation. The paper describes relying on the IJS NewsFeed (Trampuš and Novak 2012), also from the Institut Jožef Stefan. This seems to still be active (<http://newsfeed.ijs.si/>), although the page notes the stream is only internally accessible due to copyright issues. It also appears open source, although not updated since 2012.<sup>12</sup>

In a 2016 comparison of GDELT and Event Registry (Kwak and An 2016), a rate limit on Event Registry made a direct comparison impossible. Kwak and An did not attempt an article-by-article comparison, but found an overlap of about 13 thousand websites between GDELT’s 63 thou-

sand and Event Registry’s 23 thousand, leading to a conclusion that the two offered different views of the world, without any clear way to tell how they might differ.

The commercial tool currently described on the Event Registry website seems far more extensive than what was described in the 2014 paper, and potentially than what was tested against GDELT in 2016. Instead of relying on the IJS NewsFeed, Event Registry appears to have a custom collection (although perhaps built on the IJS NewsFeed) of “over 150,000 diverse sources” in over 40 languages that it has in fact spun out into a separate news API (<https://newsapi.ai/>). The site says that the tool allows clients to “access news and blog content” with search options by “keyword, tags, categories, sentiment, location and more” alongside the ability to “group content into events” that was the focus of the original paper. Articles are returned in groups of 100, with options to include the full body of text alongside metadata fields.<sup>13</sup> Beyond this, there is no information about what the tool is like or how much it costs; the website asks potential clients to get in contact for both demos and pricing options.

Now that Event Registry is now a commercial tool that lacks public documentation or free access, direct comparison with Media Cloud is difficult. However, that means that for academic and non-commercial usage where access, transparency, and documentation are important, Media Cloud has a clear advantage.

Lastly, there are web crawling tools. IssueCrawler (Rogers 2006, 2010, 2019) is possibly the earliest (Rogers and Marres 2000) and longest-running academic project to provide web crawling functionality for social science researchers. It has a strong theoretical motivation and backing based in science and technology studies (Rogers and Marres 2000; Marres and Rogers 2008; Rogers 2010, 2012), has been used by researchers external to IssueCrawler for quite some time for studying conversations around civil society (McNally 2005; Siapera 2007; Chowdhury 2013; Schmid-Petri 2017), and shares functionality of the spidering component of Media Cloud. However, on the technical side, its spidering is far more limited: IssueCrawler does not perform automatic deduplication, does not have its own internal seed set (previous sources of seed sets are now defunct), does not restrict stories by keyword, and requires repeated iterations of crawls to be done manually (Rogers 2018). It is free, although its source code does not appear to be available or open-source. It may be preferable for those who are seeking a minimal tool to do only spidering and want to do their own deduplication, filtering, and iteration in order to have fine-grained control, or who wish to use some of IssueCrawler’s other link analysis tools.

Commercial tools such as Brandwatch (into which the academic-friendly tool Crimson Hexagon recently merged; Crimson Hexagon previously sold full Twitter data access to academics at a discounted rate), Signal (which accesses social media data, and also gives subscribers access to LexisNexis), or Parse.ly offer some of the same functionality as Media Cloud. These commercial tools are focused on

<sup>12</sup><https://github.com/JozefStefanInstitute/newsfeed>

<sup>13</sup><https://eventregistry.org/documentation?tab=searchArticles;> archived snapshot at <https://perma.cc/TWN7-D2MR>



providing ready-made tools for analysis to non-academic clients, and lack documentation or access that would allow comparison with open-source and/or non-commercial tools. Since 2017, Event Registry also falls into this category.

We also note that there exist one-off datasets, including those published in ICWSM (Horne, Khedr, and Adali 2018). These obviously differ from Media Cloud and other tools in that they are static.

## About the Data

### Data Platform

Media Cloud is built as an open platform for generating research datasets. The underlying Media Cloud database includes over 1.7 billion stories belonging to 1.1 million distinct media sources, which the platform makes available through the tools and query language described above. The growth in the number of stories over time is shown in figure (2). While the entirety of the underlying data is merely a collection of media sources and topics in which the Media Cloud has been interested over the past decade, and so is not a useful representative sample of a general population of interest, the API and web tools described above enable researchers to create their own datasets based on their own interests. In this section, we describe two sample datasets generated using the platform’s publicly available tools.

### Geographic Media Collection

Media Cloud exposes a breadth of geographic coverage that is unavailable in other tools, with extensive coverage outside of the U.S. (figure 3). The geographic media collections are a curated lists of the media sources published in a particular country. These were originally populated from a data ingest of the ABYZ News Links<sup>14</sup> archive in 2018. Subsequently, our team has checked each collection for validity, and consulted with domain experts in many countries to improve coverage. The list of media sources from some countries have been less well vetted, but they are still present. A static snapshot of these geographic collections is available

<sup>14</sup><http://www.abyznewslinks.com/>




Figure 2: The total number of stories collected by Media Cloud per 7-day period (Sundays to Saturdays) since collection started in 2010, as well as the number of stories collected in English, Spanish, and French.

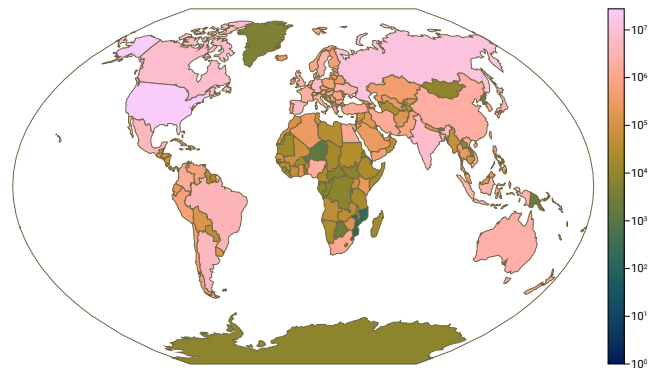


Figure 3: The number of stories for 2020 published from each country in Media Cloud, in log scale. The most stories ( $10^{7.44}$ ) are from the U.S.

at the Harvard Dataverse at <https://doi.org/10.7910/DVN/4F7TZW>, which includes over 1.5 thousand geographic collections and 12.7 thousand associated sources. This dataset can be used to drive both research and other media analysis systems.

New sources are added to the system through a variety of mechanisms. Because the media landscape is constantly changing, we continuously collaborate with local partners to validate and curate geographic collections. In addition, while Topic-specific data ingestion is underway for a research project, any new media source domains encountered will be added to the system as placeholders. Finally, system users outside of the core team can suggest media sources and collections to be added to the regular data ingest pipeline. These are vetted and approved by Media Cloud staff, based on overall system resources and validation of proposed sites as online news sources that write primarily about current events. Beyond collating by country of publication, we also specifically curate thematic collections related to research areas of interest: misinformation, reproductive health rights, partisan slant, etc.

### Voter Fraud Topic Data

A key feature of Media Cloud Topic Mapper is its ability to generate large datasets of open web stories relevant to a particular topic. For this paper, we include a sample dataset of stories, media, and links covering the vote-by mail-controversy in the U.S. during the 2020 presidential election. This dataset, available at the Harvard Dataverse at <https://doi.org/10.7910/DVN/D1XLWY>, was the basis for a paper that found that the primary drivers of voter fraud misinformation were President Trump and political and media elites rather than social media (Benkler et al. 2020).

The dataset was generated using the Topic Mapper with the Boolean query (vote\* or voti\* or ballot) and (mail or absent\*) and (fraud or rigg\* or harvest\*); the date range February 1, 2020 to November 23, 2020; and media sources from the following collections: United States - State & Local, US Center Right 2019, US Left 2019, US Center 2019, US Right 2019, U.S. Top Sources 2018, US Center Left 2019, U.S. Top Digital Native Sources 2018,

and U.S. Top Newspapers 2018. In addition, the dataset was seeded with URLs parsed from a 1.1 million tweets discovered by executing the same query over the same dates on Brandwatch. The resulting dataset includes 120,210 stories from 6,992 media sources, 82,636 story-to-story hyperlinks, and 33,029 media-to-media hyperlinks aggregated from the story links.

### Generating Datasets

Options for generating custom datasets are built into Media Cloud at all levels, including:

- Explorer: The Explorer tool allows the user to download all metadata for all stories matching any query executed on the system.
- Topic Mapper: The Topic Mapper tool allows the user to download all metadata for all stories and media discovered through the topic query and spider process.
- API: The open API allows programmatic download of metadata about all stories, media, feeds, topics, and tags in the system, including the ability to download word counts for any query, down to individual stories.
- visualizations: almost every one of the hundreds of visualizations made available by the web tools provides a CSV download of the data underlying the visualization.

### Adherence to FAIR Principles

The Media Cloud system adheres to the FAIR principles. As noted previously, we have uploaded our sample datasets to the Harvard Dataverse. In this paper, we describe the rich metadata and unique identifiers, while the upload snapshots serve as concrete examples. Datasets are accessible via the web using our published APIs and by downloading standard CSV files via the web apps. Easy access to story metadata is provided on our story details page.<sup>15</sup> In addition to the standard methods of consumption (e.g. API, CSV files), the data are made further interoperable via API clients available in Python<sup>16</sup> and R<sup>17</sup>. To be *reusable* by other researchers, our datasets and metadata must be clearly described, which we have done in this paper. A key tenet of *reuse* is provenance, and in the case of stories this is represented by the included URLs.

### Key Areas of Use

Media Cloud fosters a community of researchers and technologists that engage with us via academic conferences, a dedicated Slack channel, a community email list, and directly via support emails. This community continues to grow year after year with over 4 thousand active (e.g., performed at least one action that year) users in 2020 (figure 4). Some of these users, excluding internal Media Cloud users, performed hundreds of thousands of actions (figure 5).

As an academic engine for datasets used to support published research, Media Cloud has been growing in mentions since 2009. Figure (6) shows this pattern over time, with 45

<sup>15</sup><https://explorer.mediacloud.org/#/general/story/details>

<sup>16</sup><https://github.com/mediacloud/api-client>

<sup>17</sup><https://github.com/jandix/mediacloudr>




Figure 4: Yearly count of active (performing at least one action) unique users, excluding internal Media Cloud users.




Figure 5: An empirical complementary cumulative distribution function (CCDF) for the count of actions by unique users, again excluding internal Media Cloud users, per year.




Figure 6: Count of mentions of mediacloud.org in scholarly articles published between 2009 to 2020, via a January 2021 search in Google Scholar. Articles by the Media Cloud team are shown separately in light gray.

papers mentioning Media Cloud in 2020 (only one of which was a publication from the Media Cloud team).

Media Cloud supports researchers beyond the immediate Media Cloud team. As seen in figure (6), mentions of Media Cloud are on the rise, especially amongst external researchers. Recent papers mentioning Media Cloud include one describing it as a solution to archiving content from hyper-partisan sources (Tucker et al. 2018), a recent study of COVID-19 news and its association with sedentary activity that uses Media Cloud (Huckins et al. 2020), and a work suggesting that Media Cloud could be used to track the flow of ideas as applied to macroeconomics (Brazeal 2011).

Projects that use Media Cloud come from a diverse set of scholars, activists, journalists, and educators. Some areas of research Media Cloud datasets have already supported include:



**Mapping national political discourse.** Media Cloud supports deep exploration of national media stratification, yielding significant findings of difference between countries. Researchers have documented asymmetric polarization in the U.S. media landscape (Benkler, Faris, and Roberts 2018), while in France researchers have found a split between institutionalists and the “anti-elite” (Zuckerman et al. 2019b), but no evidence for a left/right split in influential news media.

**Online mis/dis-information.** Media Cloud supports tracking narratives emerging from unreliable sources, surfacing coordinated attempts to mainstream untrue news stories, and investigating how media are manipulated to seed fear, uncertainty and doubt by those in political power (Benkler et al. 2020). In the sub-domain of public health misinformation, our dataset has supported research into Ebola-related misinformation online (Roberts et al. 2017), the spread of conspiracy theories related to water fluoridation in the United States (Helmi, Spinella, and Seymour 2018), investigations into false claims about risk from casual exposure to fentanyl (Beletsky et al. 2020), and mis/disinformation about climate activist Greta Thunberg (Dave et al. 2020).

**Media self-analysis.** Media Cloud has been used as a tool for self-reflection within the journalism industry. Journalism schools have used the dataset to model the agenda of election coverage in the U.S. over the course of primaries for the U.S. presidential election (Bajak, Kennedy, and Wihbey 2019), and journalists have similarly used Media Cloud to track and compare their own coverage of controversial stories (Rakich and Mehta 2020). Our research team compared how different global outlets covered the 2019 Christchurch shooting against industry recommendations for covering terrorist attacks (Baumgartner et al. 2019). A variety of teachers are using Media Cloud in higher-ed classroom settings to introduce approaches to studying media attention online using our highly accessible web tools.

**Human rights and social justice media impact.** Many human rights foundations specifically fund non-profits to create impact on the media narratives around an issue they want to address. Media Cloud has helped these types of foundations and advocacy organizations track and measure media mentions (Media Ecosystems Analysis Group 2020). Our own team has worked with Define American to analyze the language of news stories about immigration to America (Ndulue et al. 2019), and analyzed media stories about police killings of unarmed people of color in the U.S., finding significant changes in the framing over time based on an analysis of news from before and after the killing of Michael Brown in 2014 (Zuckerman et al. 2019a).

**Media-based art and advocacy.** Artists have used the Media Cloud corpus to create interactive pieces showing news reporting on global conflicts (Bhargava 2003). Other

researchers (D’Ignazio et al. 2020) are supporting advocacy groups across the Americas in their efforts to track gender-related killings of women by creating automated data pipelines built on Media Cloud to surface potentially related news stories.

## Discussion and Conclusion

This paper introduces the Media Cloud platform, documenting how it is accessed, how it is produced, and how it is being used to generate open, accessible research datasets. Media Cloud provides global coverage, easy access, multiple levels of entry, and dataset creation for researchers studying attention, representation, influence and language in online news. Here, we discuss some logistical challenges, potential theoretical limitations, and future directions.

First, content collection and curation on the open web grow more difficult (Bossetta and Segesten 2015). As mentioned previously, RSS has long been a central element of the regular data ingest into Media Cloud, but it is less frequently supported. In response, we have taken efforts to increase the robustness of our content discovery via spidering and are also experimenting with ingesting sitemaps. In addition, tracking and maintaining the geographic media collections requires significant staff time and local domain expertise.

Another barrier is the growth of news content surfacing on social media platforms. Studying the media ecosystem requires understanding how news content moves across these platforms alongside the social media users who are potential news consumers (van Dijck 2013). We are now beginning to directly index some platforms, including YouTube via their API, and incorporating independent research archives such as the Pushshift Reddit Dataset (Baumgartner et al. 2020). This allows us to connect traditional stories with social sharing metrics of influence and also surface discrepancies between mainstream media and social sharing practices.

We also wrestle with questions of whether the online media we index can be considered a reasonable proxy for the stories and topics a media source covers more broadly. Are the stories on the New York Times’ website a full and accurate representation of the content that appears within the paper version of the newspaper? While print alone is becoming less important in overall news consumption (The New York Times, for example, reported having 831 thousand print subscribers as compared to 4.7 million digital subscribers in 2020; Lee 2020), this challenge is still acute for television and radio news, which continue to be important sources of information for large parts of the population (Barthel et al. 2020). For these, a given source’s online content may not be a good proxy (e.g., how much talk radio is transcribed and available online?). We believe Media Cloud provides an accurate picture of the online presence of these sources, but acknowledge the lack of coverage of offline media as a potential gap in our methodology and it is something we seek to address.

A key issue with the meaningfulness of any collection is a *denominator problem*: in order to know what percentage of online media is covering or discussing a topic one way or another, we need to know the total universe of online media.

Practically, such exhaustive sampling is impossible both because some content is not openly available and because Media Cloud's ingest and spidering can ultimately only collect a tiny sample of the open web. Thus, instead of attempting to solve the denominator problem, Media Cloud focuses on *prominence* using a relational approach, similar to IssueCrawler's claim that "hyperlinking and discursive maps provide a semblance of given socio-epistemic networks on the web" (Rogers and Marres 2000; see also Rogers 2010). Specifically, we assert that visibility via snowball link sampling, and then some form of centrality or threshold (e.g., above some minimum in-link degree, or in some upper percentile of in-link degree) within the resulting network, is a proxy for a story being prominent in the attention of media consumers. But the validity of such uses of the link economy to measure prominence relies on several key assumptions:

- Stories hyperlink to other stories on which they rely or which they reference;
- Stories are influenced by the stories to which they link;
- The influence of stories summarizes the media landscape;
- The media landscape determines the bounds of what people can consume; and
- The most prominent media items (the "center" of the media landscape) are the most influential on news consumers.

We have found the first assumption violated in many countries, notably in specific work on Saudi Arabian, Egyptian, Brazilian, Indian, and Colombian media, where web pages rarely link to external web pages. We have nonetheless had success analyzing influence in these non-hyperlinking countries by using alternative measures of influence such as social media metrics and third party panel-based metrics.

For the last assumption, there is always the possibility of stories with a large influence on the public but not linked to by other media sources. This points to a larger, deeper issue: Media Cloud is only capable of measuring production, and not consumption. Our ultimate interest in media, especially as forming an essential pillar of democratic society, is: *what media are having what influence, on whom?* But this is a question of consumption, not production. The connection between them is a long-standing concern in media research (e.g., Protesse et al. 1992), and divergences threaten the validity of any attempt to use media ecosystems to study civil society; e.g., Bruns (2007) pointed out for IssueCrawler that links (production) do not say anything about the traffic that may go over those links (consumption). Even studying tweets is evidence only of what some Twitter users produce, providing only indirect evidence of influence on those users, who themselves are a tiny and biased sample of citizenry.

It ultimately falls to users of Media Cloud to think through if and how its collection of online media provides a sufficiently valid way to measure a construct of interest, just as we are constantly doing internally. Our hope is that by being transparent about our methodology, it is easier for other researchers to reason through these issues as well.

The technology maintenance itself also poses challenges for the project. Our corpus of stories continues to grow and stress the underlying storage architectures. Keeping a collec-

tion of almost 2 billion news stories highly available for instant search results, while evolving the research it supports, requires significant investments in computational hardware and personnel. Then, although it is a positive thing that Media Cloud is seeing increased usage by researchers of many backgrounds, this adds additional stress. Lastly, our project has been generously supported by foundations to research specific topics, but we have not raised long-term funding to maintain the platform as an open resource for the broader research community, even though the core team believes strongly in that goal.

Despite these challenges, the growing user-base and mentions in scholarly literature demonstrate Media Cloud's widespread usefulness, and attest to its increasing visibility and recognition. We hope that this description will support Media Cloud's continued growth, empower more researchers to study the online media ecosystem, and inspire new and creative uses of Media Cloud.

## Acknowledgments

We thank Becky Bell and Jason Baumgartner. We also thank the collection curators.

## References

- Bajak, A.; Kennedy, D.; and Wihbey, J. 2019. How news media are setting the 2020 election agenda: Chasing daily controversies, often burying policy. *Storybench* URL <https://www.storybench.org/how-news-media-are-setting-the-2020-election-agenda-chasing-daily-controversies-often-burying-policy/>.
- Barthel, M.; Mitchell, A.; Asare-Marfo, D.; Kennedy, C.; and Worden, K. 2020. Measuring news consumption in a digital era. Report, Pew Research Center. URL <https://www.journalism.org/2020/12/08/measuring-news-consumption-in-a-digital-era/>.
- Baumgartner, J.; Bermejo, F.; Ndulue, E.; Zuckerman, E.; and Donovan, J. 2019. What we learned from analyzing thousands of stories on the Christchurch shooting. *Columbia Journalism Review* URL <https://www.cjr.org/analysis/christchurch-shooting-media-coverage.php>.
- Baumgartner, J.; Zannettou, S.; Keegan, B.; Squire, M.; and Blackburn, J. 2020. The Pushshift Reddit dataset. In *Proceedings of the Fourteenth International AAAI Conference on Web and Social Media*, ICWSM-2020, 830–839. URL <https://ojs.aaai.org/index.php/ICWSM/article/view/7347>.
- Beletsky, L.; Seymour, S.; Kang, S.; Siegel, Z.; Sinha, M. S.; Marino, R.; Dave, A.; and Freifeld, C. 2020. Fentanyl panic goes viral: The spread of misinformation about overdose risk from casual contact with fentanyl in mainstream and social media. *International Journal of Drug Policy* 86: 102951. doi:10.1016/j.drugpo.2020.102951.
- Benkler, Y. 2006. *The wealth of networks: How social production transforms markets and freedom*. Yale University Press.
- Benkler, Y.; Faris, R.; and Roberts, H. 2018. *Network propaganda: Manipulation, disinformation, and radicalization in American politics*. Oxford University Press. doi:10.1093/oso/9780190923624.001.0001.
- Benkler, Y.; Roberts, H.; Faris, R.; Solow-Niederman, A.; and Etling, B. 2013. Social mobilization and the networked public sphere: Mapping the SOPA–PIPA Debate. Research Publication No. 2013-16, Berkman Center for Internet & Society. doi:10.2139/ssrn.2295953.

- Benkler, Y.; Tilton, C.; Etling, B.; Roberts, H.; Clark, J.; Faris, R.; Kaiser, J.; and Schmitt, C. 2020. Mail-in voter fraud: Anatomy of a disinformation campaign. Research Publication No. 2020-6, Berkman Klein Center for Internet & Society. URL <https://ssrn.com/abstract=3703701>.
- Bhargava, R. 2003. Drums of war. <https://home.connectionlab.org/project/drums-of-war/>. Accessed: 2021-04-15.
- Bhargava, R.; Dave, A.; and Papakyriakopoulos, O. 2020. Investigating attention and influence online with Media Cloud. Tutorial presented at the Fourteenth International AAAI Conference on Web and Social Media. URL [https://youtube.com/playlist?list=PL48YHRNVv8PS5JltJQu5hgJmH98nDxP\\_U](https://youtube.com/playlist?list=PL48YHRNVv8PS5JltJQu5hgJmH98nDxP_U).
- Bossetta, M.; and Segesten, A. D. 2015. Tracing Eurosceptic party networks via hyperlink network analysis and #FAIL!ng: Can web crawlers keep up with web design? In *#FAIL! Workshop at the Websci'15 Conference*. URL <https://failworkshops.wordpress.com/tracing-eurosceptic-party-networks-via-hyperlink-network-analysis-and-failng/>.
- Brazeal, G. 2011. How much does a belief cost: Revisiting the marketplace of ideas. *Southern California Interdisciplinary Law Journal* 21(1): 1–46. URL <https://gould.usc.edu/why/students/orgs/ilj/assets/docs/21-1%20Brazeal.pdf>.
- Bruns, A. 2007. Methodologies for mapping the political blogosphere: An exploration using the IssueCrawler research tool. *First Monday* 12(5). doi:10.5210/fm.v12i5.1834.
- Choloniewski, J.; Sienkiewicz, J.; Leban, G.; and Hołyst, J. A. 2019. Modeling of temporal fluctuation scaling in online news network with independent cascade model. *Physica A: Statistical Mechanics and its Applications* 523: 129–144. doi:10.1016/j.physa.2019.02.035.
- Chowdhury, M. G. 2013. Understanding British Muslim environmentalism using qualitative research techniques and Issue Crawler mapping software. *Interdisciplinary Journal of Research on Religion* 9(6). URL <http://www.religjournal.com/pdf/ijr09006.pdf>.
- Dave, A.; Ndulue, E. B.; Schwartz-Henderson, L.; and Weiner, E. 2020. Targeting Greta Thunberg: A case study in online mis/disinformation. Policy Paper, German Marshall Fund. URL <https://www.gmfus.org/publications/targeting-greta-thunberg-case-study-online-misdisinformation>.
- Deacon, D. 2007. Yesterday's papers and today's technology: Digital newspaper archives and 'push button' content analysis. *European Journal of Communication* 22(1): 5–25. doi:10.1177/0267323107073743.
- D'Ignazio, C.; Bhargava, R.; Zuckerman, E.; and Beck, L. 2014. CLIFF-CLAVIN: Determining geographic focus for news articles. In *NewsKDD: Data Science for News Publishing, at KDD 2014*. URL [http://agava.ijs.si/~blazf/NewsKDD2014/submissions/newskdd2014\\_submission\\_10.pdf](http://agava.ijs.si/~blazf/NewsKDD2014/submissions/newskdd2014_submission_10.pdf).
- D'Ignazio, C.; Fumega, S.; Suresh, H.; Cruxên, I.; So, W.; Martínez, M. D. L. A.; and García-Montes, M. 2020. Femicide & machine learning: Detecting gender-based violence to strengthen civil sector activism. In *4th Workshop on Mechanism Design for Social Good, MD4SG '20*.
- Dorfman, L. 2003. Studying the news on public health: How content analysis supports media advocacy. *American Journal of Health Behavior* 27(Suppl 3): S217–226. doi:10.5993/ajhb.27.1.s3.5.
- Finkel, J. R.; Grenager, T.; and Manning, C. 2005. Incorporating non-local information into information extraction systems by Gibbs sampling. In *Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics, ACL '05*, 363–370. doi:10.3115/1219840.1219885.
- Fishman, M. 1978. Crime waves as ideology. *Social Problems* 25(5): 531–543. doi:10.2307/800102.
- Fronzetti Colladon, A.; and Grippa, F. 2020. Brand intelligence analytics. In *Proceedings of the 9th International COINs Conference, COINs 2019*, 125–141. doi:10.1007/978-3-030-48993-9\_10.
- Gans, H. J. 2004/1979. *Deciding what's news: A study of CBS Evening News, NBC Nightly News, Newsweek, and Time*. Northwestern University Press.
- Gonçalves, A. A. 2014. *Conflicting frames: The dispute over the meaning of rolezinhos in Brazilian media*. Master's thesis, Massachusetts Institute of Technology. URL <http://hdl.handle.net/1721.1/92660>.
- Graeff, E.; Stempeck, M.; and Zuckerman, E. 2014. The battle for 'Trayvon Martin': Mapping a media controversy online and offline. *First Monday* 19(2). doi:10.5210/fm.v19i2.4947.
- Helmi, M.; Spinella, M. K.; and Seymour, B. 2018. Community water fluoridation online: An analysis of the digital media ecosystem. *Journal of Public Health Dentistry* 78(4): 296–305. doi:10.1111/jphd.12268.
- Horne, B.; Khedr, S.; and Adali, S. 2018. Sampling the news producers: A large news and feature data set for the study of the complex media landscape. In *Proceedings of the Twelfth International AAAI Conference on Web and Social Media, ICWSM-2018*, 518–527. URL <https://www.aaai.org/ocs/index.php/ICWSM/ICWSM18/paper/view/17796>.
- Huckins, J. F.; daSilva, A. W.; Wang, W.; Hedlund, E.; Rogers, C.; Nepal, S. K.; Wu, J.; Obuchi, M.; Murphy, E. I.; Meyer, M. L.; Wagner, D. D.; Holtzheimer, P. E.; and Campbell, A. T. 2020. Mental health and behavior of college students during the early phases of the COVID-19 pandemic: Longitudinal smartphone and ecological momentary assessment study. *Journal of Medical Internet Research* 22(6): e20185. doi:10.2196/20185.
- Kayser, J. 1953. *One week's news: Comparative study of 17 major dailies for a seven-day period*. UNESCO.
- Kwak, H.; and An, J. 2014a. A first look at global news coverage of disasters by using the GDELT dataset. In *Proceedings of the 6th International Conference on Social Informatics, SocInfo 2014*, 300–308. doi:10.1007/978-3-319-13734-6\_22.
- Kwak, H.; and An, J. 2014b. Understanding news geography and major determinants of global news coverage of disasters. arXiv:1410.3710. URL <https://arxiv.org/abs/1410.3710>.
- Kwak, H.; and An, J. 2016. Two tales of the world: Comparison of widely used world news datasets GDELT and EventRegistry. In *Proceedings of the Tenth International AAAI Conference on Web and Social Media, ICWSM-2016*, 619–622. URL <https://www.aaai.org/ocs/index.php/ICWSM/ICWSM16/paper/view/13014>.
- Leban, G. 2020. Event Registry. Seminar: AI for Industry and Society, Ljubljana 2020. URL [http://videolectures.net/AIindustrySeminar2019\\_leban\\_event\\_registry/](http://videolectures.net/AIindustrySeminar2019_leban_event_registry/).
- Leban, G.; Fortuna, B.; Brank, J.; and Grobelnik, M. 2014. Event Registry: Learning about world events from news. In *Proceedings of the 23rd International Conference on World Wide Web, WWW '14 Companion*, 107–110. doi:10.1145/2567948.2577024.
- Lee, E. 2020. New York Times hits 7 million subscribers as digital revenue rises. *The New York Times* URL <https://nyti.ms/364KCKJ>.
- Leetaru, K.; and Schrodt, P. A. 2013. GDELT: Global Data on Events, Location and Tone, 1979–2012. URL <http://data.gdelproject.org/documentation/ISA.2013.GDELT.pdf>.

- Marres, N.; and Rogers, R. 2008. Subsuming the ground: how local realities of the Fergana Valley, the Narmada Dams and the BTC pipeline are put to use on the Web. *Economy and Society* 37(2): 251–281. doi:10.1080/03085140801933314.
- McNally, R. 2005. Sociomics! Using the IssueCrawler to map, monitor and engage with the global proteomics research network. *PROTEOMICS* 5(12): 3010–3016. doi:10.1002/pmic.200500163.
- Media Ecosystems Analysis Group. 2020. Women in the news in the age of coronavirus. <https://perma.cc/5CEK-UHSZ>. Accessed: 2021-01-16.
- Ndulue, E. B.; Bermejo, F.; Ramos, K.; Lowe, S. E.; Hoffman, N.; and Zuckerman, E. 2019. The language of immigration reporting: Normalizing vs. watchdogging in a nativist age. Report, Define American. URL <https://www.defineamerican.com/journalismreport>.
- Protest, D. L.; Cook, F. L.; Doppelt, J. C.; Ettema, J. S.; Gordon, M. T.; Leff, D. R.; and Miller, P. 1992. *The journalism of outrage: Investigative reporting and agenda building in America*. Guilford Press.
- Rakich, N.; and Mehta, D. 2020. How the media has — and hasn't — covered Tara Reade's allegation. <https://fivethirtyeight.com/features/how-the-media-has-and-hasnt-covered-tara-reades-allegation-against-joe-biden/>. Accessed: 2021-04-15.
- Roberts, H.; Seymour, B.; Fish, S. A., II; Robinson, E.; and Zuckerman, E. 2017. Digital health communication and global public influence: A study of the Ebola epidemic. *Journal of Health Communication* 22(sup1): 51–58. doi:10.1080/10810730.2016.1209598.
- Rogers, R. 2006. Mapping web space with the Issue-crawler. [http://govcom.org/publications/full\\_list/rogers\\_politics\\_webspace.pdf](http://govcom.org/publications/full_list/rogers_politics_webspace.pdf). Accessed: 2021-04-15.
- Rogers, R. 2010. Mapping public web space with the Issuecrawler. In Reber, B.; and Brossaud, C., eds., *Digital cognitive technologies: Epistemology and the knowledge economy*, 89–99. ISTE Ltd. doi:10.1002/9781118599761.ch6.
- Rogers, R. 2012. Mapping and the politics of web space. *Theory, Culture & Society* 29(4–5): 193–219. doi:10.1177/0263276412450926.
- Rogers, R. 2018. Issuecrawling : Building lists of URLs and mapping website networks. In Lury, C.; Fensham, R.; Heller-Nicholas, A.; Lammes, S.; Last, A.; Michael, M.; and Uprichard, E., eds., *Routledge Handbook of Interdisciplinary Research Methods*, 169–175. Routledge. doi:10.4324/9781315714523-26.
- Rogers, R. 2019. *Doing digital methods*. SAGE.
- Rogers, R.; and Marres, N. 2000. Landscaping climate change: A mapping technique for understanding science and technology debates on the World Wide Web. *Public Understanding of Science* 9(2): 141–163. doi:10.1088/0963-6625/9/2/304.
- Rubinovitz, Y. 2017. *News Matter: Embedding human intuition in machine intelligence through interactive data visualizations*. Master's thesis, Massachusetts Institute of Technology. URL <http://hdl.handle.net/1721.1/112544>.
- Rupnik, J.; Muhič, A.; Leban, G.; Fortuna, B.; and Grobelnik, M. 2017. News across languages: Cross-lingual document similarity and event tracking. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence, IJCAI-17*, 5050–5054. doi:10.24963/ijcai.2017/720.
- Schmid-Petri, H. 2017. Politicization of science: How climate change skeptics use experts and scientific evidence in their on-line communication. *Climatic Change* 145(3–4): 523–537. doi:10.1007/s10584-017-2112-z.
- Siapera, E. 2007. Multiculturalism, progressive politics and British Islam online. *International Journal of Media & Cultural Politics* 2(3): 331–346. doi:10.1386/macp.2.3.331\_1.
- Sittar, A.; Mladenčić, D.; and Erjavec, T. 2020. A dataset for information spreading over the news. In *Proceedings of the 2020 Slovenian KDD Conference on Data Mining and Data Warehouses*, SiKDD 2020. URL <https://ailab.ijs.si/dunja/SiKDD2020/Papers/01%20-%20A-Dataset-for-Information-Spreading-over-the-News.pdf>.
- Trampuš, M.; and Novak, B. 2012. Internals of an aggregated web news feed. In *Proceedings of the 2012 Slovenian KDD Conference on Data Mining and Data Warehouses*, SiKDD 2012. URL [http://ailab.ijs.si/dunja/SiKDD2012/Papers/Trampus\\_Newsfeed.pdf](http://ailab.ijs.si/dunja/SiKDD2012/Papers/Trampus_Newsfeed.pdf).
- Tucker, J. A.; Guess, A.; Barberá, P.; Vaccari, C.; Siegel, A.; Sanovich, S.; Stukal, D.; and Nyhan, B. 2018. Social media, political polarization, and political disinformation: A review of the scientific literature. Report, Hewlett Foundation. URL <https://hewlett.org/library/social-media-political-polarization-political-disinformation-review-scientific-literature/>.
- van Dijck, J. 2013. *The culture of connectivity: A critical history of social media*. Oxford University Press.
- Wang, W.; Kennedy, R.; Lazer, D.; and Ramakrishnan, N. 2016. Growing pains for global monitoring of societal events. *Science* 353(6307): 1502–1503. doi:10.1126/science.aaf6758.
- Ward, M. D.; Beger, A.; Cutler, J.; Dickenson, M.; Dorff, C.; and Radford, B. J. 2013. Comparing GDELT and ICEWS event data. URL <https://www.researchgate.net/publication/303211430>.
- Weaver, D. A.; and Bimber, B. 2008. Finding news stories: A comparison of searches using LexisNexis and Google News. *Journalism & Mass Communication Quarterly* 85(3): 515–530. doi:10.1177/107769900808500303.
- Weller, N.; and McCubbins, K. 2014. Raining on the parade: some cautions regarding the Global Database of Events, Language and Tone dataset. *Political Violence @ a Glance* (2014, February 20). URL <https://politicalviolenceataglance.org/2014/02/20/raining-on-the-parade-some-cautions-regarding-the-global-database-of-events-language-and-tone-dataset/>.
- Zuckerman, E. 2003. Global attention profiles – a working paper: First steps towards a quantitative approach to the study of media attention. Research Publication No. 2003-06, Berkman Center for Internet & Society. doi:10.2139/ssrn.487943.
- Zuckerman, E. 2013. *Rewire: Digital cosmopolitans in the age of connection*. WW Norton & Co.
- Zuckerman, E.; Matias, J. N.; Bhargava, R.; Bermejo, F.; and Ko, A. 2019a. Whose death matters? A quantitative analysis of media attention to deaths of Black Americans in police confrontations, 2013–2016. *International Journal of Communication* 13: 4751–4777. URL <https://ijoc.org/index.php/ijoc/article/view/8782>.
- Zuckerman, E.; Patino, B.; Bosser, F.; Cardon, D.; François, C.; Frau-Meigs, D.; Fogel, J.-F.; Jay, O.; Morenne, B.; Nimmo, B.; Pariser, E.; Puyfaucher, L.; Reille-Soult, V.; Sauty de Chalon, M.-L.; Shefet, D.; Schmelck, C.; Wardle, C.; Zylberberg, H.; Lenoir, T.; Corbel, A.; Jolys, F.; Van Muylders, J.; Cointet, J.-P.; Plique, G.; and Ooghe-Tabanou, B. 2019b. Media polarization “à la française”? Comparing the French and American ecosystems. Report, Institut Montaigne. URL <https://www.institutmontaigne.org/en/publications/media-polarization-la-francaise>.