

Digital Hegemonies: The Localness of Search Engine Results

Andrea Ballatore ,* Mark Graham ,[†] and Shilad Sen [‡]

^{*}Department of Geography, Birbeck, University of London

[†]Oxford Internet Institute, University of Oxford

[‡]Department of Mathematics, Statistics, and Computer Science, Macalester College

Every day, billions of Internet users rely on search engines to find information about places to make decisions about tourism, shopping, and countless other economic activities. In an opaque process, search engines assemble digital content produced in a variety of locations around the world and make it available to large cohorts of consumers. Although these representations of place are increasingly important and consequential, little is known about their characteristics and possible biases. Analyzing a corpus of Google search results generated for 188 capital cities, this article investigates the geographic dimension of search results, focusing on searches such as “Lagos” and “Rome” on different localized versions of the engine. This study answers these questions: To what degree is this city-related information locally produced and diverse? Which countries are producing their own representations and which are represented by others? Through a new indicator of localness of search results, we identify the factors that contribute to shape this uneven digital geography, combining several development indicators. The development of the publishing industry and scientific production appears as a fairly strong predictor of localness of results. This empirical knowledge will support efforts to curb the digital divide, promoting a more inclusive, democratic information society. *Key Words:* digital place, Google, Internet geography, localness, search engines.

每天有亿万的用户，依赖搜索引擎来搜索信息，以决定旅游、购物与其他无数经济活动的地点。搜索引擎透过晦涩的过程，凑组在世界各地生产的数码内容，并提供给大量的消费者。尽管这些地方的再现日益重要且为必然的结果，但它们的特征和潜在的偏见却鲜为人知。本文分析一百八十八座首都城市的谷歌搜索结果，探讨搜索结果的地理面向，并聚焦不同的在地化版本的搜索引擎中诸如“拉各斯”和“罗马”的搜索。本研究回答下列问题：此般城市的相关信息，由在地生产并有所不同的程度为何？哪个国家生产自身的再现，哪些又是由他人进行呈现？透过崭新地方性指标的研究结果，我们指认导致形塑此一不均数码地理的因素，并结合若干发展指标。出版业与科学生产的发展，似乎是指认在地性的研究结果的强大指标。该经验知识将支援遏止数码差距的努力，提倡更具包容性且民主的信息社会。关键词：数码地方，谷歌，互联网地理，在地性，搜索引擎。

Miles de millones de usuarios de Internet dependen cada día de los motores de búsqueda para obtener información sobre lugares con la cual tomar decisiones relacionadas con turismo, compras y otras incontables actividades económicas. En un proceso opaco, los motores de búsqueda ensamblan el contenido digital producido en una variedad de localizaciones alrededor del mundo y lo ponen a disposición de grandes cohortes de consumidores. Si bien estas representaciones de lugar son cada vez más importantes y consecuentes, poco se sabe de sus características y posibles sesgos. Mediante el análisis de un *corpus* de resultados de búsquedas en Google generado para 188 ciudades capitales, este artículo investiga la dimensión geográfica de los resultados de búsqueda, enfocándose en cierto tipo de búsquedas como “Lagos” y “Roma” en versiones del motor de búsqueda localizadas de manera diferente. El estudio responde a estas preguntas: ¿En qué grado es producida la información relacionada con esta ciudad de modo local y diverso? ¿Qué países están produciendo sus propias representaciones y cuáles son representados por otros? A través de un nuevo indicador del carácter local de los resultados de la búsqueda, identificamos los factores que contribuyen a configurar esta geografía digital desigual, combinando varios indicadores del desarrollo. El desarrollo de la industria editorial y de la producción científica surge como un fuerte predictor del carácter local de los resultados. Este conocimiento empírico apoyará los esfuerzos para frenar la divisoria digital,

© 2017 Andrea Ballatore, Mark Graham, and Shilad Sen.

This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives License (<http://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited, and is not altered, transformed, or built upon in any way.

promoviendo una sociedad de la información más incluyente y democrática. *Palabras clave:* lugar digital, Google, geografía del Internet, el carácter de lo local, motores de búsqueda.

Until the lion learns how to write, every story will glorify the hunter.

—Ewe (Ghanaian proverb)

Our spatial interactions and experiences with the world are increasingly digital. The cities and towns that we live in were once constructed from bricks, concrete, glass, and steel. Now, though, the places we inhabit also consist of digital augmentations (Graham 2013): digital code and digital content like data from Wikipedia, photographs from Flickr, restaurant reviews from Yelp, and algorithms controlled by Google that make information visible or invisible. The digital augmentations of places, in other words, are beginning to matter as much as their material counterparts, and much research has shown how they can have significant impacts on how we economically, socially, and politically interact with our environments (see Kitchin and Dodge 2011; Graham, De Sabbata, and Zook 2015). Digital code and content thus do not just reflect the world but also produce it.

This does not mean that informational augmentations of places represent anything profoundly new. Indeed, the ability to represent place has long been a domain of conflict. This geographic information (whether digital or predigital) has always been produced under conditions of power (Crampton 2008) and is inherently both a product and a producer of power relations (Harley 1989; Pickles 2004). These conditions thus tend to reinforce and legitimate the dominant and powerful (Gramsci 1971). Histories and narratives about the Global South, for instance, could thus be written by the colonizer rather than the colonized (Said 1978).

Because the Web has long been envisaged as a participatory tool, and because its usership now approaches 3.5 billion people (almost half of the world's population), many have hoped that the construction of information geographies could become more open, participatory, and democratic. Harvard Law Professor Lawrence Lessig (2003), speaking at the World Summit on the Information Society, noted, "For the first time in a millennium, we have a technology to equalize the opportunity that people have to access and participate in the construction of knowledge and culture, regardless of their geographic placing." For Lessig, and other commentators such as Benkler (2007), Bruns (2008),

Jenkins (2006), Tapscott and Williams (2006), and Shirky (2011), as places become more digital, they could also become more participatory.

Unfortunately, many of those hopes have not been realized. As the state increasingly rolls back from the cartographic project, spatial information is increasingly controlled by for-profit companies that have entirely different motives than their public predecessors (Leszczynski 2012). The field of critical geographic information systems (GIS), for instance, has long concerned itself with the fact that digital representations of places are rarely equitable or genuinely participatory (see Pickles 1995). Elwood (2006) noted that little has changed for those at the bottom of the digital divide, and financial and skill barriers continue to influence who gets a say in digital representations of the world and who does not (see also Craig and Elwood 1998). More recently, work has shown how user-generated digital content such as Wikipedia not just largely represents the Global North but is also overwhelmingly produced by users in the Global North (Graham, Hale, and Stephens 2011; Graham et al. 2014; Graham, De Sabbata, and Zook 2015; Graham, Straumann, and Hogan 2016).

The key organizing mechanism behind our contemporary digital architectures is not user-generated content, however, but rather the search engine (Hillis, Petit, and Jarrett 2013). Most people use a digital search as their gateway not just to a disembodied Web but also to their lived everyday geographies: using it to learn about a destination, shop, navigate, and perform countless other activities. This is because information search usually takes the form of a one-to-many relationship (Graham 2010). A user has a requirement for information, and the available set of published information is much larger than the user has utility for. Search engines thus do not just direct a user to relevant information but filter that information to distinguish between more and less useful content. Although search engines only index a fraction of the dynamic and growing corpus of Web content (National Aeronautics and Space Administration 2015), they provide a central access point to it.

As geography becomes ever more digital, search engines thus increasingly mediate not just information but also spatial knowledges and experiences. This

article therefore seeks to better understand the geography of information in search engines. Specifically, it focuses on Google—the world’s most powerful, most dominant information mediator—to examine the locality of content about places around the world, asking whether Google directs users to locally produced information or nonlocally produced information. It does this by focusing on search results generated in the 188 countries where Google is currently available, when searching for capital cities. In doing so, it brings novel empirical data about one of the most practiced ways in which we access digital geographic information to bear onto older questions of power and geography.¹

Search Engines, Mediation, and Power

Although they portray themselves as neutral aggregators of information, search engines have created an informational infrastructure with precise characteristics, logics, and biases (Graham, Schroeder, and Taylor 2014; Ballatore 2015). Scanning, interpreting, and organizing large volumes of online information to be served to users, global mediators such as Google and Bing play a crucial role in determining which Web sites, news, blogs, videos, and photographs become visible to whom. Although it is impossible to estimate how much of existing digital content is actually captured by search engines, they provide a highly visible access point to the “surface Web” crawled by their bots. This mediation between content producers and consumers remains opaque, unfolding behind closed doors with far-reaching consequences. On this point, Grimmelmann (2010) stated that “search engines are the new mass media . . . capable of shaping public discourse itself” (436).

Although a variety of Web search services and technologies exist, the search market is controlled by a handful of large actors. A comScore (2015) report shows that majority of desktop searches in the United States in 2015 were executed on Google (64 percent), Bing (20 percent), and Yahoo! (13 percent), capturing a staggering 97 percent of the market. The dominant position of Google is even more pronounced in the United Kingdom, where it attracts 88 percent of searches, and in other European countries, where it commands similar market shares (theEword 2015). Other search engines firmly dominate large markets, notably Baidu in China, Yandex in Russia, and Naver in South Korea, but they are confined to their home countries. By contrast, Google still has the lion’s share of the global search market, attracting a large majority

of searches in most countries and collecting 54 percent of the global search advertising market in 188 countries. This monopoly has not gone unnoticed and has attracted the attention of European antitrust agencies and a variety of critics, worried about such a concentration of information, power, and capital (e.g., Vaidhyanathan 2011; Grimmelmann 2013).

Social scientists and humanists have analyzed search engines from cultural, cognitive, and political viewpoints, pointing out how these tools exert a powerful influence on society (Wouters and Gerbec 2003; Spink and Zimmer 2008; Halavais 2010; Mager 2012; Brosard and Scheufele 2013; Hillis, Petit, and Jarrett 2013; Graham et al. 2014; König and Rasch 2014; Ballatore et al. 2016). The importance of these new mediators is observable in the industry of search engine optimization (SEO), which reflects how these tools are now at the core of the media landscape, representing a considerable portion of global advertising markets. The main activity of SEO consists of shaping and adapting Web content to make it more visible on specific engines, reverse-engineering their algorithms. In this sense, content producers are influenced by the mediation of search engines as much as consumers, and biases in results have real effects (Vaughan and Thelwall, 2004).

Political analyses of search engines focus on the forms of power they exert. Notably, Epstein and Robertson (2015) identified what they termed a *search engine manipulation effect* (SEME) as the influence that biased search results can have on political choices. In different contexts, engines might help groups spread counternarratives and fringe ideologies (Ballatore 2015) or, by contrast, further entrench dominant positions (Introna and Nissenbaum 2000).

The personalization of results is another researched aspect of search engines. Based on sophisticated and rich personal profiles, Google Search produces different content for different users, increasing and decreasing the visibility of links to deliver more relevant results. This process, according to Pariser (2011), might result in a “filter bubble,” in which users are systematically exposed only to content that matches their political and cultural inclinations: In 2011, different users searching for “Tahrir Square” were shown either news reports about the revolution that started there or Web sites of travel agencies that did not engage with the political context. The Google Search personalization was strongly criticized for its lack of transparency, and the company recently introduced an option to disable it. Although precise quantification is difficult, a study suggests that on average about 12 percent of

results differ because of personalization, and this applies only to users logged into their Google account (Hannak et al. 2013).

Research on search engine effects has raised valid concerns, but most studies have overlooked the spatiality of the information retrieval process. In this sense, digital representations of physical places play an increasingly important societal and cultural role (Ballatore 2014; Graham et al. 2014). No prior research has addressed the representation of places in search engine results, analyzing where the Web content assembled by the algorithms is generated from.

The Geography of Google Search Results

Google systematically collects massive amounts of data worldwide and serves it to billions of users in 188 countries. The characteristics and effects of Google's presences and absences can be investigated from approaches developed in the subfield of Internet geography by, for example, studying the spatial layout of its large network infrastructure and the patterns of usage of its services across the globe (Dodge and Zook 2009). This study focuses, in particular, on Google Search, the company's search engine and most used service, analyzing the geography of its content.

To use an example, when searching for information about the city of Ankara on Google by typing the text string "Ankara" into a search box, a user obtains a set of links that includes Wikipedia articles, government Web sites, tourist guides, and news stories (Figure 1). This Google page represents a highly visible entry point to obtain information about the city's geography, economy, politics, history, and culture. Despite recent

efforts by Google (2014) in promoting transparency, it is extremely hard to know with any precision how and why these Web resources are selected over competing ones, as hundreds of signals about the relevance of resources are combined into the ranking presented to users. Indeed, it is worth mentioning that Google Search actually uses the geography of Web content in calculating its relevance. This is visible in the option in Google's interface that allows the user to filter results by country when using a localized version of the product. For example, when searching for "Ankara" on the UK version (google.co.uk), the system allows users to choose between results from "any country" and results exclusively generated from the United Kingdom.

Default options, however, exert a powerful influence when users face complex choices, and most users do not alter the default settings of software tools, including search engines (Nielsen 2005). This friction against changing preset options is referred to as the *default effect* by psychologists, and it has been identified in a wide variety of contexts (Dinner et al. 2011). Default search results are therefore what the vast majority of users searching for "Ankara" will see. For this reason, it is important to pay particular attention to the default results returned by Google Search, even though users have a small degree of control over them. The effect of personalization should thus not deter research into Google results, because only about 12 percent of links differ from the nonpersonalized, default results (Hannak et al. 2013).

The central question of this study is this: Where is the Web content that is returned by Google Search produced? In our case of Ankara, some URLs point to local content produced and hosted in Turkey, whereas

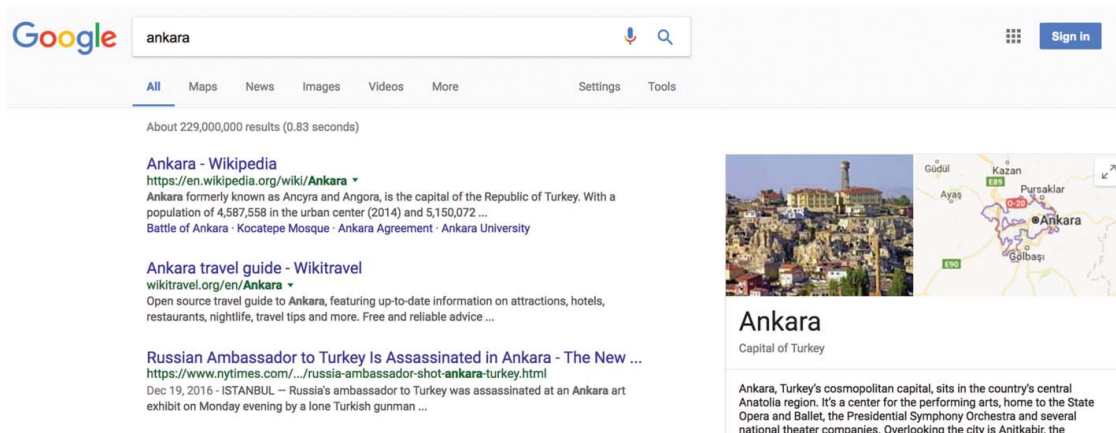


Figure 1. An example of Google search results for "Ankara" on 10 February 2017. (Color figure available online.)

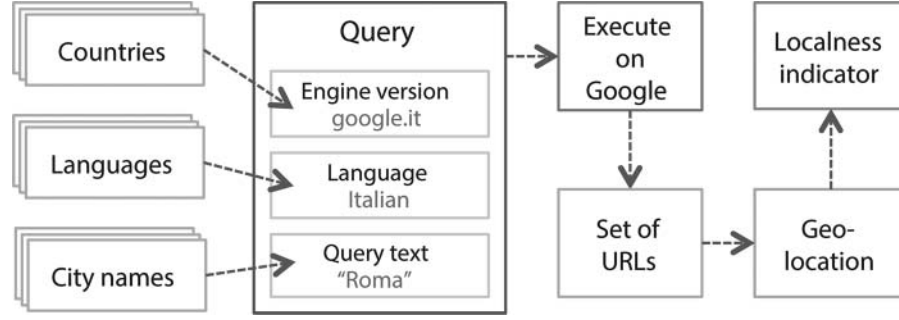


Figure 2. Overview of the study of the localness of Google search results.

others refer to content from the United States and other countries. In other words, the search results have a certain degree of localness that can be quantified and can reveal crucial facets of Google’s information geographies. For this empirical investigation, we consider the representation of the world’s capital cities. The general workflow of the study is outlined in Figure 2. Next we illustrate the study’s methodology, data collection, and analysis. The research data sets are available online.²

Method

To investigate the localness of Google search results in a systematic way, we collected search results for capital cities, adopting the methodology outlined by Balatore (2015). This approach consists of extracting search results at different times and at different geolocations, reducing the effects of personalization and spatiotemporal biases in the data. The repetition and randomization of the queries produce results that are more stable and reproducible than individual observations at a specific location and time. To closely simulate a user’s experience on Google, we extracted results from Google’s HTML search results pages (i.e., the page that a user sees), approaching the typical

usage of the service, rather than from the Google Custom Search API, which is known for returning different results (Google 2017).

Google Search accepts many parameters that allow a user to increase the relevance of results. This study focuses on the effects of three parameters related to the geography of search: the country-specific localization setting for the search engine (e.g., *google.com* or *google.it*), the desired language of results (e.g., English or Italian), and the capital city query text (“Rome” or “Roma”). Other parameters were left to their default options. For each country, we generated a set of searches for the capital of the country. We also generated one U.S.-centric Google query for the U.S. version of Google (*google.com*) with the capital’s English name, abbreviated henceforth as U.S. Google.

Second, we generated queries for the capital city in each language officially supported by Google in each country, which we refer to as local Google queries. For example, because Arabic and French are official languages in Morocco, we included two local Google queries with results in Arabic and in French, respectively. The local Google data only include languages that are supported in the target country and not any other languages. Table 1 shows a sample of these

Table 1. Sample of three Google queries and results

| Country | Engine version | Language | Query text | Local Google | Top three URLs |
|---------|----------------|----------|------------|--------------|--|
| Egypt | google.com | en | “Cairo” | False | 1. en.wikipedia.org/wiki/Cairo 2. www.lonelyplanet.com/egypt/cairo 3. en.egypt.travel/city/index/cairo |
| Uganda | google.co.ug | en | “Kampala” | True | 1. en.wikipedia.org/wiki/Kampala 2. www.lonelyplanet.com/uganda/kampala 3. www.kcca.go.ug |
| Vietnam | google.com.vn | fr | “Hanoi” | True | 1. fr.wikipedia.org/wiki/Hanoi 2. tripadvisor.fr/Tourism-g293924... 3. routard.com/guide_voyage_lieu/... |

Google queries, with one query on the U.S. Google and two on local versions of Google. One query (“Washington, DC,” in English) is present in both the U.S. and local data sets.

The set of search queries used in this study is shaped by variations in different countries’ access and representations within Google and the Web more broadly. Whereas some countries have unlimited access to the search engine (e.g., the United States), others face total censorship if proxies are not used (e.g., Iran). Google provides localized Web sites for 188 countries (e.g., Google.it in Italy or Google.co.ke in Kenya), supporting 103 languages, all of which are included in the study. The list of countries and languages was extracted directly from the Google Web site on February 2015, and the capital city names in different languages were obtained from the multilingual gazetteer GeoNames. At the time of data collection, Google was not available in China, Iran, North Korea, and Cuba and therefore these countries were excluded from the study.

Data Collection and Validation

At the core, Google Search is an online service that, given a set of input parameters, returns search engine result pages (SERP). The parameters are passed to the service through a uniform resource locator (URL), visible in the user’s browser (an example of this would be <https://www.google.com/search?q=sri+lanka>). To collect results systematically, it is therefore essential to generate relevant and stable URLs. *Relevant* refers to results that are about the intended topic (e.g., about the French capital Paris and not about the American celebrity Paris Hilton), and *stable* refers to the degree to which the results are the same regardless of the search location. To identify such URLs, we ran a series of trials in which we manually executed a sample of queries from different geolocations (United Kingdom, United States, France, and Italy) and compared the results. The URL that showed the highest stability and relevance has four URL query parameters: “q,” which specifies the textual query (e.g., $q = \textit{Ankara}$); “hl,” which specifies the query’s results language ($hl = \textit{en}$ for English); “gl,” which specifies the provenance of results ($gl = \textit{us}$ for U.S. results); and “oe,” which specifies the text encoding ($oe = \textit{utf8}$).³ The text encoding was set to *utf8* for all queries to make the results consistent and easily machine-readable, and the other parameters were changed for each query.

The URL query parameter *gl*, which emphasizes results from a particular country (e.g., $gl = \textit{us}$ for U.S.

results), plays an important role in this study. Without this parameter, Google prioritizes results from the current geolocation of the user. For example, searching for English results about “Rome” from a machine located in France returns mostly results from France, regardless of the other parameters. Setting the *gl* parameter to “US” enables the observation of the typical results that a North American user sees when searching for “Rome” in the United States, accessing more stable, representative results as opposed to transient, personalized results (Ballatore 2015). The same results could be generated from machines physically located in the target countries, obtaining extremely similar results. To obtain a more spatially diffuse sample, the URLs were accessed through the Tor network, obtaining a different IP address for each URL.

As a result of this process, 188 URLs were generated for U.S. Google (one for each country) and 357 for local Google for each country in the languages supported by Google (1.9 languages per country on average), for a total of 545 URLs. Because the first page of results attracts more than 91 percent of clicks (Chitika 2013), for each query we collected the results on the first page, typically varying between eight and twelve URLs. To reduce the temporal bias of the results, each query was executed four times over three months, obtaining four separate snapshots of the 545 queries.⁴ In total, 33,736 result URLs were collected. Over time, as observed in Ballatore (2015), the composition of the results varies up to 20 percent, particularly with respect to news stories. By aggregating the snapshots over time, the more stable results are reinforced, because they tend to occur in all four of the snapshots, whereas more transient results occur in fewer snapshots.

After the collection phase, the resulting data set was analyzed for quality control. To reduce noise in the data, the results of the 545 queries were inspected to ensure that they captured the object of the study; that is, the typical results of Google queries in different countries. This inspection showed that, out of 545 queries, twenty-four had invalid results, mainly because of errors in the multilingual data from GeoNames.⁵ As a result of this validation, these twenty-four cases (4.4 percent) were removed from the data set, resulting in 32,327 search results. The validated data set can therefore be considered reliable for this study and is summarized in Table 2. The linguistic composition of the searches on local versions of Google broadly reflects the size of linguistic communities online, with the notable exception of China, which was not included in the study: English (27.4 percent),

Table 2. Overview of Google capitals data set

| Data set characteristic | Value |
|--|---|
| Countries officially supported by Google | 188 |
| Languages included | 103 |
| Dates when URLs were captured | 24 March 2015, 15 April 2015, 22 April 2015, 9 May 2015 |
| Google queries | 545 URLs collected four times |
| Query parameters | Engine version, language of results, query text |
| Results collected per query | All results on the first page |
| Collected URLs | 32,327 (11,091 U.S.; 21,236 local) |
| Languages for local Google | English (27.4%), French (9.8%), Spanish (6.5%), Arabic (5.1%), Russian (3.5%), Portuguese (2.3%), German (2.2%), others (44.4%) |

French (9.8 percent), Spanish (6.5 percent), Arabic (5.1 percent), Russian (3.5 percent), Portuguese (2.3 percent), and German (2.2 percent); all other languages combined amount to 44.4 percent.

The data set also provides insight into the Web sites that dominate the representation of places in Google. Table 3 shows the top fifteen Web sites that obtain the highest visibility in both U.S. Google and the local versions of Google, showing a combination of crowd-sourced reference Web sites (Wikipedia,

Table 3. Fifteen most visible domains in the search results for U.S. Google (11,091 URLs) and local Google (21,236 URLs)

| U.S. Google top domains | % | Local Google top domains | % |
|-------------------------|-------|--------------------------|-------|
| wikipedia.org | 10.11 | wikipedia.org | 10.17 |
| wikitravel.org | 5.18 | lonelyplanet.com | 2.56 |
| tripadvisor.com | 5.10 | facebook.com | 2.11 |
| lonelyplanet.com | 4.65 | usembassy.gov | 1.51 |
| facebook.com | 2.25 | youtube.com | 1.38 |
| youtube.com | 1.98 | wikitravel.org | 0.76 |
| timeanddate.com | 1.20 | localtimes.info | 0.71 |
| nationsonline.org | 0.82 | tripadvisor.com | 0.64 |
| google.com | 0.79 | accuweather.com | 0.58 |
| britannica.com | 0.73 | booking.com | 0.44 |
| wikivoyage.org | 0.67 | timeanddate.com | 0.29 |
| theguardian.com | 0.53 | diplo.de | 0.28 |
| booking.com | 0.51 | hilton.com | 0.23 |
| usembassy.gov | 0.41 | gismeteo.ru | 0.21 |
| nytimes.com | 0.41 | 24timezones.com | 0.20 |
| Other domains | 64.7 | Other domains | 77.7 |

Wikitravel, Wikivoyage), social media platforms (Facebook), newspapers (*New York Times* and *The Guardian*), and travel agencies (Booking.com). Wikipedia has the strongest presence, occupying about 10 percent of the URLs. The tourist Web sites Wikitravel and TripAdvisor rank second and third highest in the United States but only sixth and eighth in the rest of the world. Overall, the top fifteen Web sites provide between 35 percent (U.S. Google) and 22 percent (local Google) of the content, with a tail of less prominent Web sites. Figure 3 summarizes the global distribution of URLs at the country level, highlighting the strong influence of the United States and Western Europe as Web content producers.

Geolocation of Web Pages

Our analysis requires us to accurately identify the geographic origin of each Web page returned as a Google Search result. We operationalize this by using Sen et al.’s (2015) notion of the *geoprovenance* of a URL, or the country primarily responsible for publishing the information on a particular Web page. Although defining geoprovenance of some pages can be difficult or ambiguous, Sen et al. found that human coders agreed 93 percent of the time on the geoprovenance of a Web page, indicating that the specific definition used in this article provides a reliable measure of the geographic origin of information.

To scale the identification of URL geoprovenance to all 32,327 URLs in our data set, we used Sen et al.’s (2015) geoprovenance inference algorithm, which accurately predicts the country that published a specific URL, adopting its reference implementation.⁶ In summary, the algorithm relies on five signals predictive of the geoprovenance of a particular URL:

1. The administrative contact’s mailing address from a whois⁷ database search for the URL.
2. A search for any known organizational headquarters locations associated with the domain using the Wikidata database (e.g., the IBM Web site is linked to the company’s headquarters in Armonk, New York; see <https://www.wikidata.org>).
3. Identification of suffixes associated with specific countries in a URL’s top-level domain (e.g., “.uk” for “bbc.co.uk”).
4. Geolocation of the country physically housing the server hosting facility for a particular URL, obtained by geocoding a URL’s IP address.

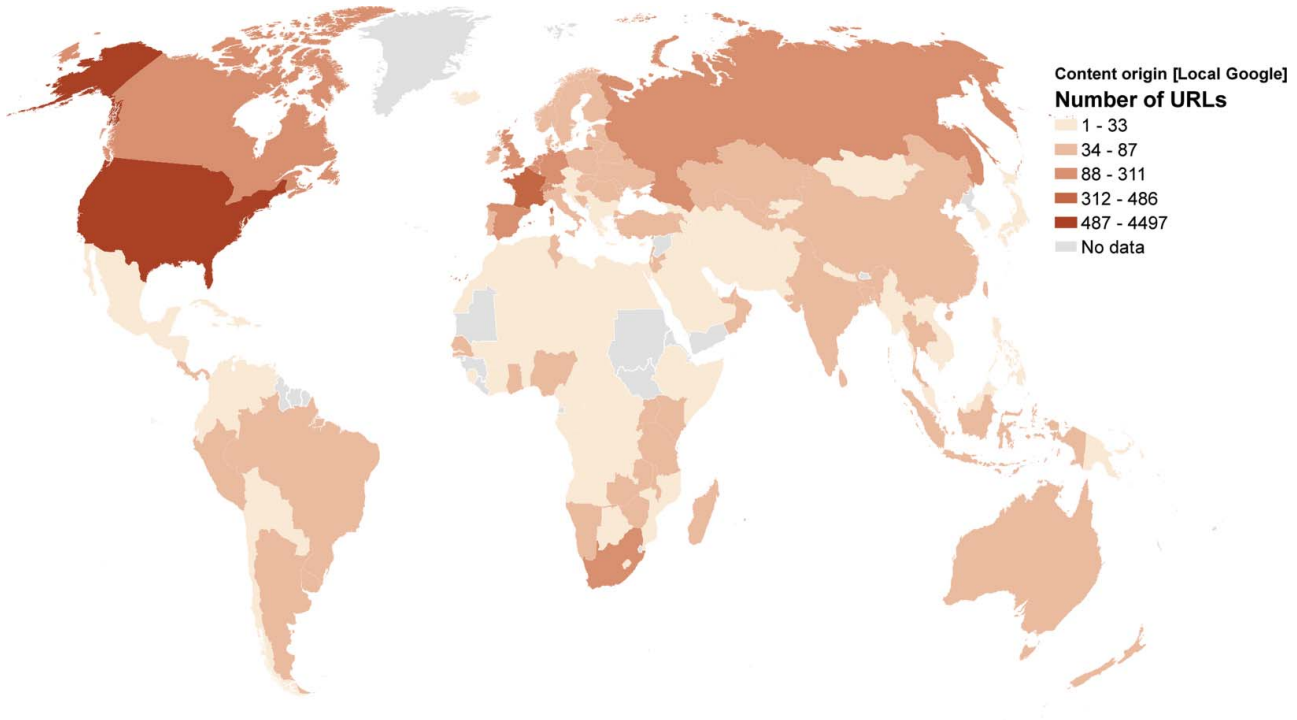


Figure 3. Number of URLs generated from each country, grouped by natural breaks. The top group contains only the United States. (Color figure available online.)

5. The language of the content of a specific URL is used to identify candidate countries that have proficiency in the language.

When combined using a machine learning algorithm, these five different signals achieve 91 percent accuracy, approaching human levels of agreement. In addition to predicting the country that published a URL, the geoprovenance inference algorithm assigns a *confidence indicator* to each URL prediction ranging from 0.0 (no confidence) to 1.0 (full confidence). The confidence indicator is calibrated against Sen et al.'s human-created ground truth data set to estimate the probability a predicted publisher country is correct. Thus, we use this confidence indicator to validate the quality of the inferred publisher countries in our data set.

Across our data set, the average confidence for a set of search results ranges from 0.42 to 0.95 with a median of 0.88, indicating high confidence in publisher country predictions for the vast majority of cases, with few outliers with low values, such as Nigeria and Mali. The map in Figure 4 shows the confidence in the classification spatially, as the mean probability per country of a correct classification. We consider these values appropriate for this study, with a probability of $p > 0.8$ for most major countries.

Localness of Google Search Results

This analysis quantifies the localness of search results for each country. It considers results for both the U.S. version of Google and all local versions. To study this geography of content, we define a localness indicator L as the ratio between local results and the total number of results. Hence, L ranges from 0 (all search results are nonlocal) to 1 (all search results are local). More formally, we define localness L of a country c and URLs U as the ratio of URLs originated from country c U_c and total URLs:

$$L(c, U) = \frac{|U_c|}{|U|}.$$

L is a simple ratio that assumes equal weight of the first page results. This simplified assumption makes L easy to calculate and apply across data sets. Although more complex, weighted indicators could indeed be closer to the actual prominence of links, but they would also be less interpretable. Because our analysis seeks to understand variation at the country level, the unit of analysis is a single country, and we do not place more importance on a country with a large population than one with a small population. We did take several pre-processing steps to make the data set less biased. Although we weight countries equally, there are a




Figure 4. Confidence in localness L , as the mean probability of correct classification of content origin. (Color figure available online.)

large number of very small countries, such as micronations in the Pacific, with fewer than 1 million inhabitants. To prevent these countries from influencing our analysis (micronations tend to have a great deal of incomplete data), we excluded them, leaving 144 countries and ninety-nine languages in the analysis.

We also took several steps to clean the search results data. An inspection of the geolocations of the URLs revealed that all URLs of images were pointing to Google cache services, making them unreliable. For this reason, all URLs to images were removed from the computation of the localness indicator (12,056). Similar caution is needed when considering the effect of Wikipedia URLs in the results. Wikipedia is a large international crowd-sourcing project with a complex content geography that cannot be reduced to its main host country, the United States (Sen et al. 2015). Because one Wikipedia page occurs in every single result set, though, its effect on the localness indicator can be safely ignored. Hence, Wikipedia URLs were not removed from the data set.

We computed the indicator L for the 144 countries. Because each case includes a Google query at four different times, the localness for a country is expressed as *mean* L of the four temporal snapshots and the corresponding standard deviation (SD). As shown in Figure 5, the value of L varies widely across the data set. Figure 5 compares the distribution of L in U.S. Google




Figure 5. Distribution localness indicator L in U.S. Google ($N = 144$) and local Google ($N = 297$).

Table 4. Descriptive statistics of the data set

| Data set | N | Median L | Skewness L | Kurtosis L | Median PL | Outliers |
|--------------|-----|------------|--------------|--------------|-----------|--|
| U.S. Google | 144 | 0.22 | 1.70 | 7.15 | 0.90 | United States, United Kingdom, Australia, Canada |
| Local Google | 297 | 0.37 | 0.29 | 1.91 | 0.85 | None |

(median = 0.22) with local Google results (median = 0.37). The gap between the distributions reflects the expected fact that results for U.S. Google in English are less local than the localized versions of Google in different languages. The distribution of local Google is particularly wide, ranging from 0 to 1 without large gaps, indicating that cases exist for every level of localness. As shown in Table 4, U.S. Google data are less local, more skewed, and have four outliers, corresponding with high-income, English-speaking countries (United States, Canada, United Kingdom, and Australia). Localized versions of Google are more local, less skewed, and without noticeable outliers. Table 5 shows the countries grouped in five categories of localness, ranging from very low to very high, showing how cases are spread uniformly across the spectrum.

The global variation in localness L can be observed through a regional lens. Figure 6 shows the distribution of L for local Google grouped by the seven World Bank regions. To make the countries comparable, countries with multiple languages are aggregated into one point, considering the average L across languages.

Whereas North America (median $L = 0.9$) and South Asia (median $L = 0.36$) are tightly clustered, all of the other regions show considerable variation, having both countries with low localness ($L < 0.3$) and countries with high localness ($L > 0.6$). Europe and Central Asia have the second highest localness after North America (median $L = 0.68$), followed by East Asia and Pacific and Latin America and the Caribbean (median $L \sim 0.45$). Substantially lower L values are observable for sub-Saharan Africa (median $L \sim 0.27$) and the Middle East and North Africa (median $L = 0.24$). Figures 7 and 8 show the same distributions spatially, at the country level. These maps are complemented by Figure 4, which represents the uncertainty in the classification, ranging from 0 (no certainty in the countries of origin of the URLs) to 1 (total certainty), based on the probability of a correct classification.

To corroborate this indicator, we also measured the diversity of results in terms of countries of origin. Whereas some countries receive results from a few dominant countries, others receive results from a broader range of sources. To quantify diversity at the country

Table 5. Overview of local Google localness, aggregated by country ($N = 144$)

| Localness | L range | N | Countries |
|-----------|------------|-----|---|
| Very low | (0, 0.2) | 27 | Albania, Angola, Bahrain, Benin, Botswana, Burundi, Cambodia, Cameroon, Central African Republic, Chad, Congo, Democratic Republic of the Congo, Ethiopia, Haiti, Kuwait, Lesotho, Libya, Mozambique, Palestine, Panama, Papua New Guinea, Puerto Rico, Somalia, Timor-Leste, Trinidad and Tobago, Turkmenistan, Vietnam |
| Low | (0.2, 0.4) | 37 | Afghanistan, Algeria, Azerbaijan, Côte d'Ivoire, Cyprus, Dominican Republic, Egypt, Gabon, Georgia, Ghana, India, Indonesia, Iraq, Jamaica, Kyrgyzstan, Laos, Lebanon, Malawi, Malaysia, Mali, Morocco, Namibia, Nepal, Nicaragua, Niger, Nigeria, Philippines, Qatar, Romania, Rwanda, Saudi Arabia, Sierra Leone, Sri Lanka, The Gambia, Togo, Zambia, Zimbabwe |
| Medium | (0.4, 0.6) | 32 | Armenia, Bangladesh, Belarus, Bolivia, Burkina Faso, Costa Rica, Cuba, Denmark, Estonia, Greece, Honduras, Hong Kong, Ireland, Israel, Jordan, Kazakhstan, Kenya, Latvia, Mauritius, Moldova, Myanmar/Burma, Oman, Pakistan, Peru, Senegal, Singapore, Tanzania, Thailand, Uganda, United Arab Emirates, Uzbekistan, Venezuela |
| High | (0.6, 0.8) | 30 | Argentina, Australia, Belgium, Bosnia and Herzegovina, Brazil, Bulgaria, Chile, Colombia, Croatia, Ecuador, El Salvador, France, Guatemala, Hungary, Japan, Macedonia, Madagascar, Mexico, Mongolia, Netherlands, Paraguay, Poland, Portugal, Serbia, South Africa, Spain, Tajikistan, Tunisia, Ukraine, United Kingdom |
| Very high | (0.8, 1.0) | 18 | Austria, Canada, Czech Republic, Finland, Germany, Italy, Lithuania, New Zealand, Norway, Russia, Slovakia, Slovenia, South Korea, Sweden, Switzerland, Turkey, United States, Uruguay |




Figure 6. Distribution localness indicator L for local Google in 144 countries, grouped by World Bank regions, with global median ($N = 144$).

level, we draw an ecological analogy comparing species diversity in an ecosystem with the diversity of countries of origin in the URLs. Hence, we compute a widely used indicator of diversity based on Shannon entropy (Levine and HilleRisLambers 2009) on the URLs, obtaining an index ranging between 0 (low diversity) and 2 (high diversity). Figure 9 shows the variation of this entropy-based diversity globally. It is possible to notice that, although African and South and Southeast Asian countries tend to obtain low levels of local content, their results also tend to be produced from a wider range of places than those of countries in the Global North.

We also observed the country that generates the highest number of URLs in a result set for a country. The map in Figure 10 illustrates the ability of Google to capture local content in North America, most of South America, Europe, and Central Asia. By contrast, most of Africa, the Middle East, and Southeast Asia are dominated by URLs from the United States

and France. Given the dominant position of the United States in this arena, Figure 11 represents the proportion of URLs from the United States in local versions of Google, showing the country's global but deeply uneven reach in this geography of content.

Explaining Localness at the Country Level

The localness of Google Search results varies substantially around the world. Although this variability is largely expected, we seek to understand the factors that influence it. In this section, we build an explanatory model of localness at the country level that includes a variety of socioeconomic indicators. Henceforth, the mean localness L is the dependent variable, and all of the other variables are considered to be explanatory. For explanatory variables, we consider a wide range of variables related to the robustness of digital infrastructure

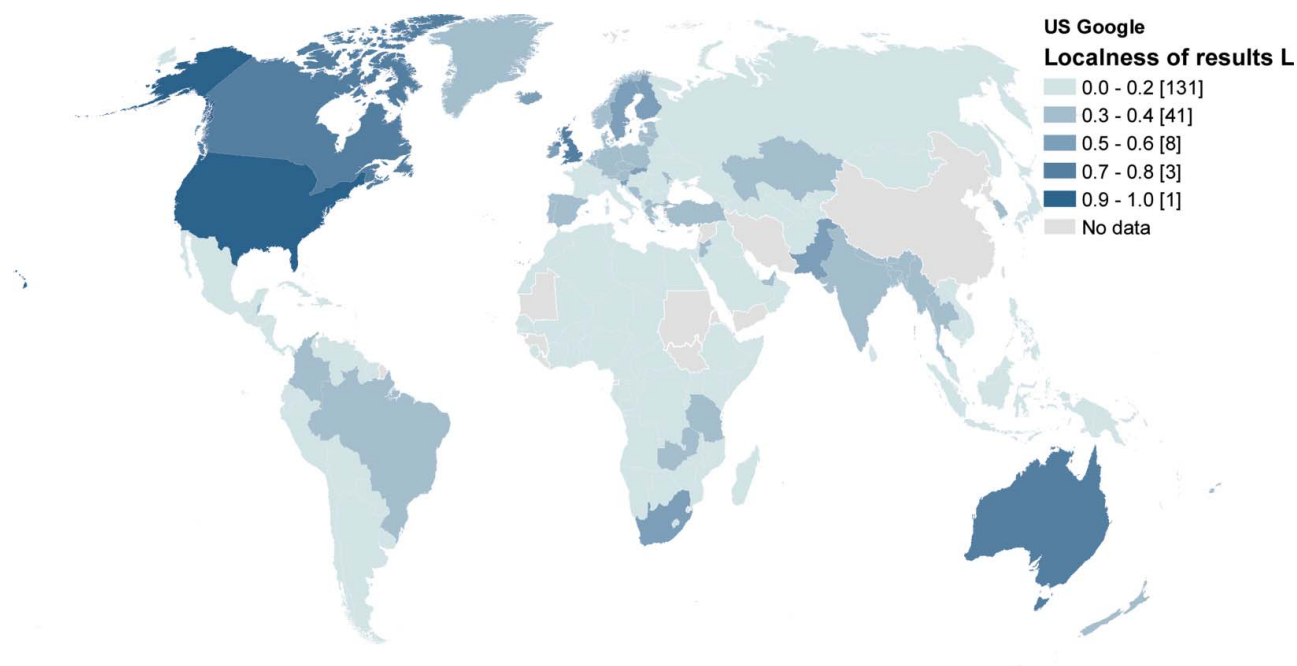


Figure 7. Localness indicator per country for the U.S. version of Google and queries in English. (Color figure available online.)

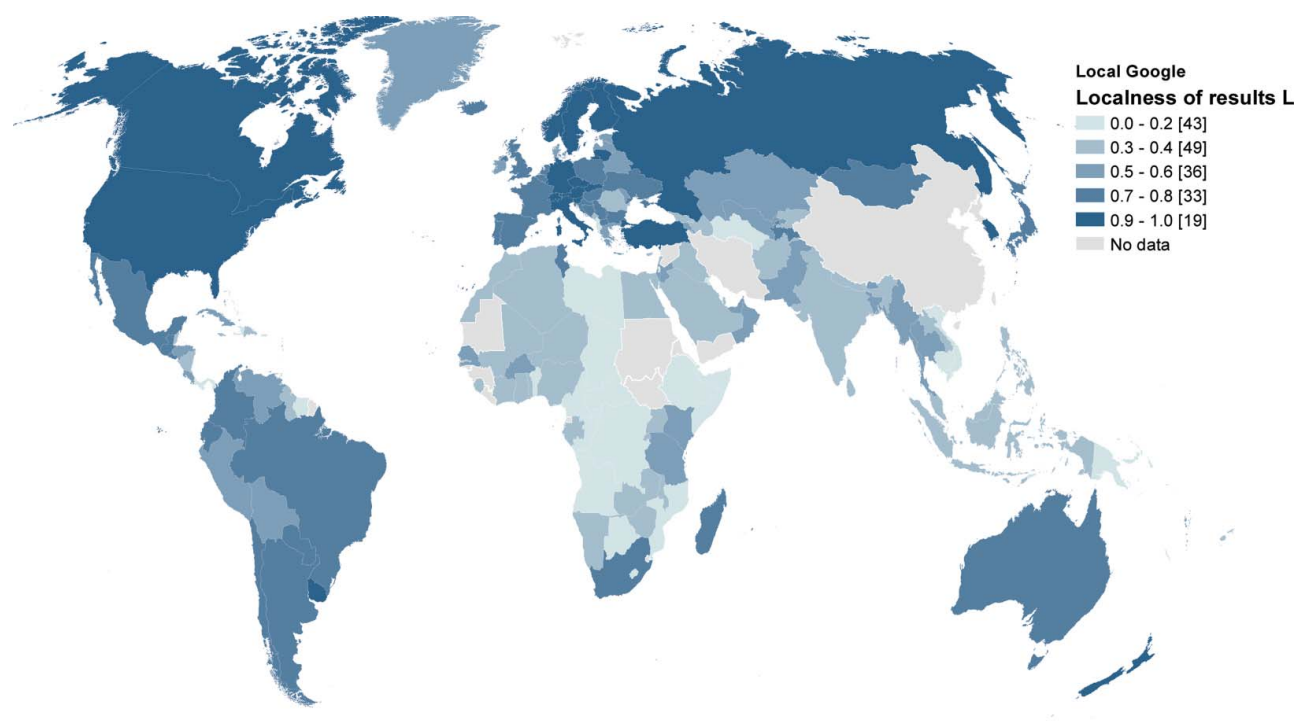


Figure 8. Localness indicator per country for the local versions of Google and queries in local languages. (Color figure available online.)

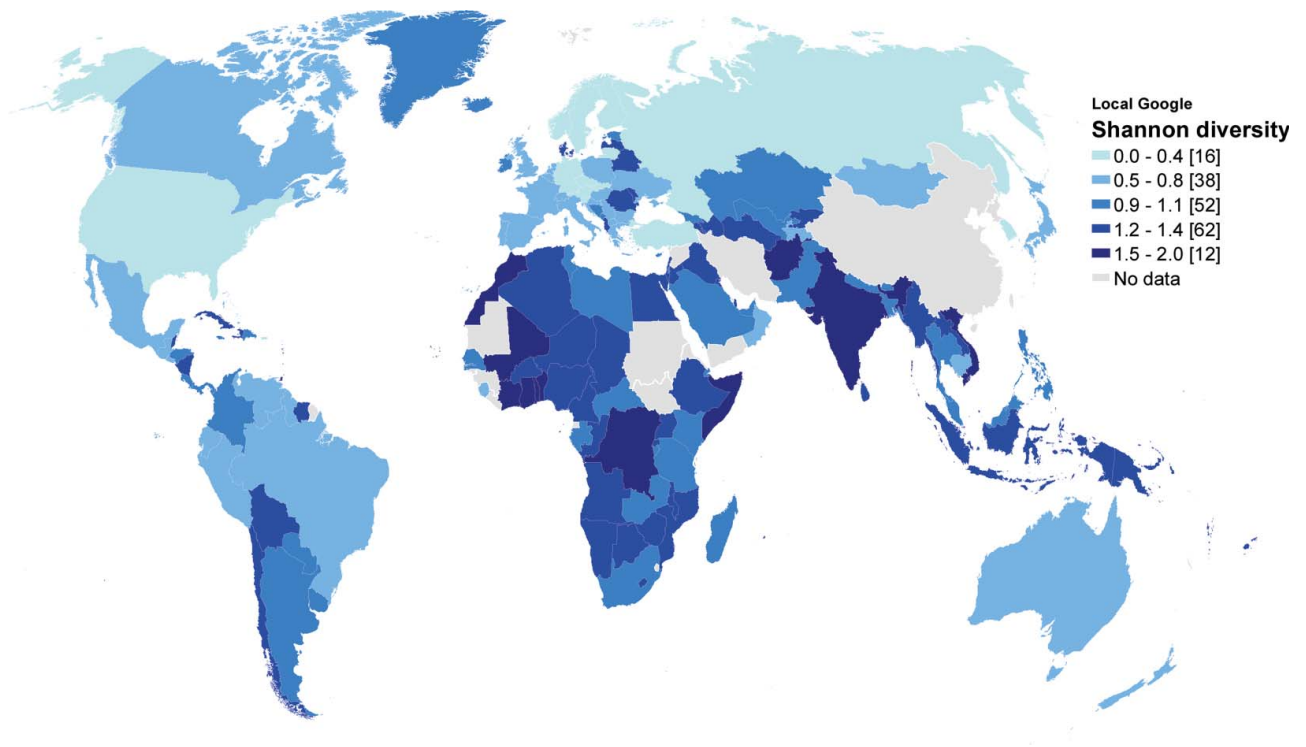


Figure 9. Diversity of results for local versions of Google. Low values indicate results from fewer countries, whereas high values indicate results from many countries. (Color figure available online.)

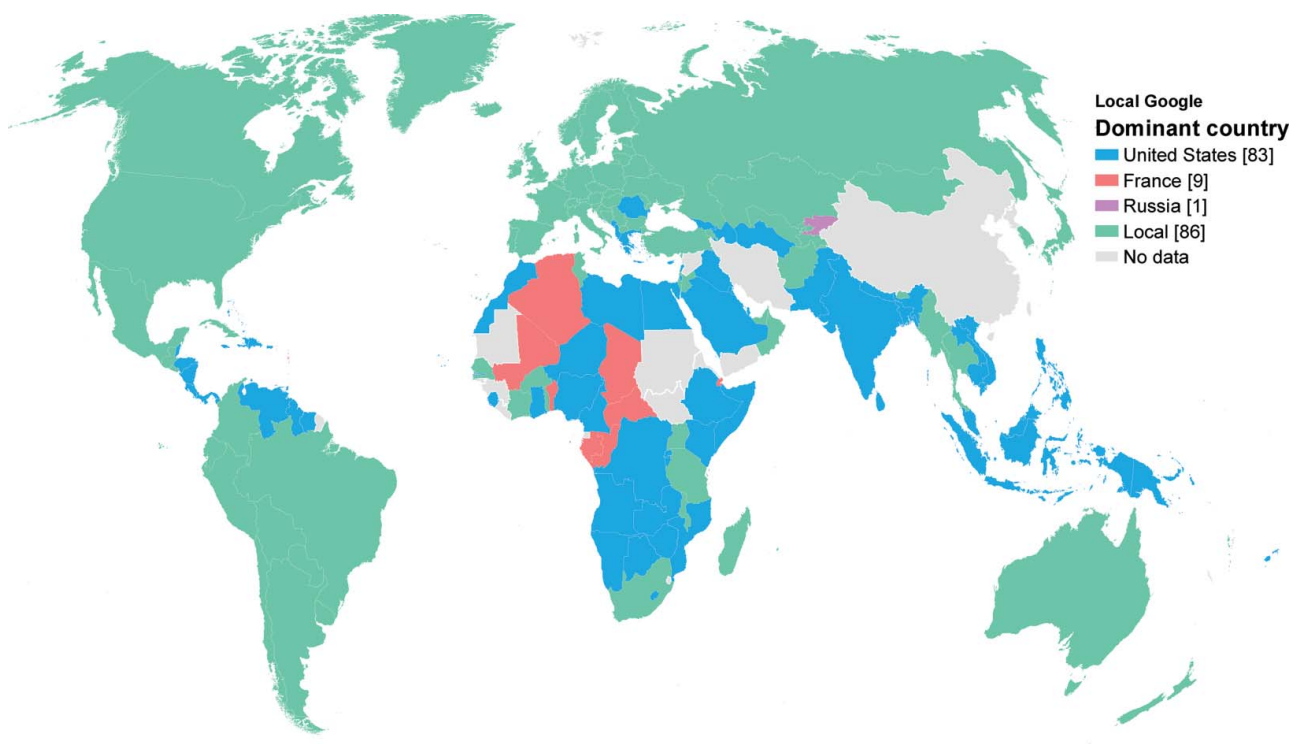


Figure 10. Countries that dominate results in local versions of Google. (Color figure available online.)




Figure 11. Proportion of URLs from the United States in local versions of Google. (Color figure available online.)

(e.g., population with Internet access), education levels, and other socioeconomic indicators published by the World Bank (see <http://data.worldbank.org>) that might relate to the ability for individuals within a certain country to produce searchable content. We used the World Bank data sets from 2011, the most recent complete data set. In addition to socioeconomic indicators, we include indicators related to scholarly publication from the Spanish research group SciMago (see <http://www.scimagojr.com>) that past research has shown to be strongly correlated with the geographic provenance of information on the Web (e.g., Sen et al. 2015). The bibliometric data collected by SciMago draws on more than 21,000 journals in the Scopus database (see <https://www.scopus.com>) and captures the impact of scientific publications around the world.

Table 6 shows a complete list of variables in our analysis and the correlation coefficients between localness L and the explanatory variables, using both Pearson's r and Spearman's ρ . The SciMago indicators on the publishing industry show substantially higher correlations with the localness indicator than the World Bank indicators. In particular, the h -index of a country, calculated as the maximum h such that h articles within the country have been cited at least h times, exhibits strong correlations (>0.55 for both coefficients). Although there are differences between U.S. and local

Google Pearson's r , the differences appear greatly when using Spearman's ρ . This suggests that there are non-linear relationships in the U.S. data that are mostly linear in the local Google data. This might reflect the large inequalities that result in long-tailed or skewed distributions across a variety of socioeconomic indicators (Ostry, Berg, and Tsangarides 2014). For example, the top ten countries by gross domestic product (GDP) account for 66 percent of the world's GDP, and statistical analyses that do not normalize GDP data might be overly influenced by countries with the largest GDPs (e.g., the United States and China).

Localness Regressions

Next, we develop separate regression models that explain localness of results from the U.S. and local Google data sets. Because many columns have missing variables, we filter out the thirty-five observations missing ten or more explanatory variables, leaving 171 complete observations for U.S. Google and 340 for local Google. We use random forests to impute values for any remaining missing values (Stekhoven and Bühlmann 2012), with the imputed values achieving a relatively low normalized root mean squared error (NRMSE) of 0.18 (NRMSE values range from 0.0 to 1.0). Because the dependent variable L represents a proportion, and not an absolute value, we use

Table 6. Bivariate correlations between localness L with U.S. Google ($N = 171$) and local Google versions ($N = 340$)

| Explanatory variables | | Pearson's r with localness L | | Spearman's rho with localness L | |
|-----------------------|-------------------|----------------------------------|----------------|-----------------------------------|----------------|
| Data source | Country variable | U.S. Google | Local Google | U.S. Google | Local Google |
| WBD 2011 | GDP | 0.30*** | 0.30*** | 0.46** | 0.44*** |
| WBD 2011 | GDPPC | 0.42*** | 0.42*** | 0.40*** | 0.40*** |
| WBD 2011 | Internet users | 0.54*** | 0.53*** | 0.49*** | 0.49*** |
| WBD 2011 | Internet servers | 0.42*** | 0.39*** | 0.43*** | 0.45*** |
| WBD 2011 | Population | 0.11 | -0.01 | 0.19* | 0.15 |
| WBD 2011 | Tourism revenue | 0.41*** | 0.37*** | 0.49*** | 0.49*** |
| WBD 2011 | Tourism visitors | 0.27*** | 0.36*** | 0.41*** | 0.43*** |
| SciMago | Documents | 0.42*** | 0.32*** | 0.54*** | 0.55*** |
| SciMago | Citable documents | 0.41*** | 0.32*** | 0.54*** | 0.54*** |
| SciMago | Citations | 0.42*** | 0.31*** | 0.53*** | 0.55*** |
| SciMago | Self-citations | 0.38*** | 0.20** | 0.54*** | 0.56*** |
| SciMago | h -index | 0.56*** | 0.56*** | 0.53*** | 0.56*** |

Note: Values are Pearson correlation coefficients. The highest coefficients are shown in bold. WBD = World Bank data (2011); GDP = gross domestic product; GDPPC = gross domestic product per capita.

* $p < 0.05$.

** $p < 0.01$.

*** $p < 0.001$.

a general linear model with a logistic link function (Long 1997).

To begin, we study to what extent U.S. Google localness can be explained by the World Bank socioeconomic indicators and SciMago publishing metrics. A forward variable selection process that starts with no explanatory variables and iteratively adds the “best” unused variable identifies a ten-variable model that explains 58 percent of the overall deviation in L . This implies that most of the variation in search localness does in fact follow topologies that reflect measurable characteristics. Although the full ten-variable model explains a substantial amount of the variation in L , the collinearities highlighted in Table 5 make the model difficult to interpret. To provide a more interpretable model M1, we identified four variables using forward selection that explain 53 percent of the total variance: h -index, region (a categorical variable with seven World Bank regions), the percentage of Internet users, and a Boolean factor indicating whether or not English is considered a “local” language in the country by Google:

$$\text{M1 : } L \sim h\text{-index} + \text{region} \\ + \text{Internet users} + \text{English is a local language.}$$

The terms in this model can be interpreted as follows. The h -index term indicates a country's level of activity and impact in scholarly publishing. Internet users reflects the level of digital infrastructure in a country

that supports both the creation and consumption of Web resources. For example, Italy and Japan exhibit similar h -indexes, neither has English as a local language, but Japan has more Internet users (79 percent vs. 55 percent) and also higher U.S. localness (16 percent vs. 12 percent). The coefficients of this model, along with the results of an analysis of variance (ANOVA) on the regression results, including degrees of freedom and deviance, are shown in Tables 7 and 8. The residuals of the regression are shown spatially in Figure 12.

Unsurprisingly, countries whose official languages include English exhibit higher U.S. localness. For example, both Norway and Ireland are in the same region (Europe), exhibit high levels of Internet users (93 percent and 75 percent, respectively), and have relatively high h -indexes (439 and 364). Despite Norway's modest advantage in both Internet users and h -index, Ireland's use of English as an official language correlates with its much higher levels of localness (49 percent vs. 29 percent). Finally, as shown in Figure 6, region-level localness exhibits large variations, reflecting many possible region-level linguistic, economic, and other cultural patterns.

We repeated the same procedure for local query results. A seven-variable model explains 53 percent of the deviation, whereas a four-variable model explains just under half (49 percent). Forward selection identifies the same four variables as most explanatory (see Tables 9 and 10). Figure 13 shows the residuals spatially.

Table 7. General linear model on U.S. Google with model M1

| Data set: U.S. Google ($N = 171$) | Estimate | SE | t value | Pr(> t) |
|-------------------------------------|-----------|----------|---------|-----------|
| (Intercept) | 3.11E-01 | 1.16E-01 | 2.675 | 0.0082** |
| h -index | 1.79E-04 | 6.31E-05 | 2.846 | 0.005** |
| Region: North America | 0 | 0 | 0 | N/A |
| Region: East Asia & Pacific | -2.73E-01 | 1.14E-01 | -2.39 | 0.018* |
| Region: Europe & Central Asia | -2.35E-01 | 1.14E-01 | -2.07 | 0.040* |
| Region: Latin America & Caribbean | -3.56E-01 | 1.15E-01 | -3.10 | 0.002** |
| Region: Middle East & North Africa | -3.27E-01 | 1.17E-01 | -2.79 | 0.006** |
| Region: South Asia | -1.56E-01 | 1.20E-01 | -1.30 | 0.195 |
| Region: Sub-Saharan Africa | -2.86E-01 | 1.16E-01 | -2.46 | 0.015** |
| Internet users | 1.76E-03 | 4.97E-04 | 3.53 | 0.0005*** |
| English spoken locally | 9.24E-02 | 3.13E-02 | 2.96 | 0.004** |

* $p < 0.05$.
 ** $p < 0.01$.
 *** $p < 0.001$.

Table 8. Analysis of variance test on U.S. Google ($N = 171$) in Table 6

| Data set: U.S. Google | df | Deviance | Residual df | Residual deviance |
|--------------------------|----|----------|----------------|----------------------|
| Null | | | 170 | 3.89 |
| h -index | 1 | 1.214 | 169 | 2.68 |
| Region | 6 | 0.541 | 163 | 2.14 |
| Internet users | 1 | 0.202 | 162 | 1.94 |
| English spoken locally | 1 | 0.100 | 161 | 1.84 |

Discussion

The general similarities between the U.S. and local Google data sets might reflect the variety in “local” languages supported by Google for each country and Google’s frequent support for English for a country even though it is not an official language. For example, for Switzerland, Google supports five languages: German, English, French, Italian, and Romansh (and all five appear in our Google local data set). English, however, is not one of the official languages of Switzerland.

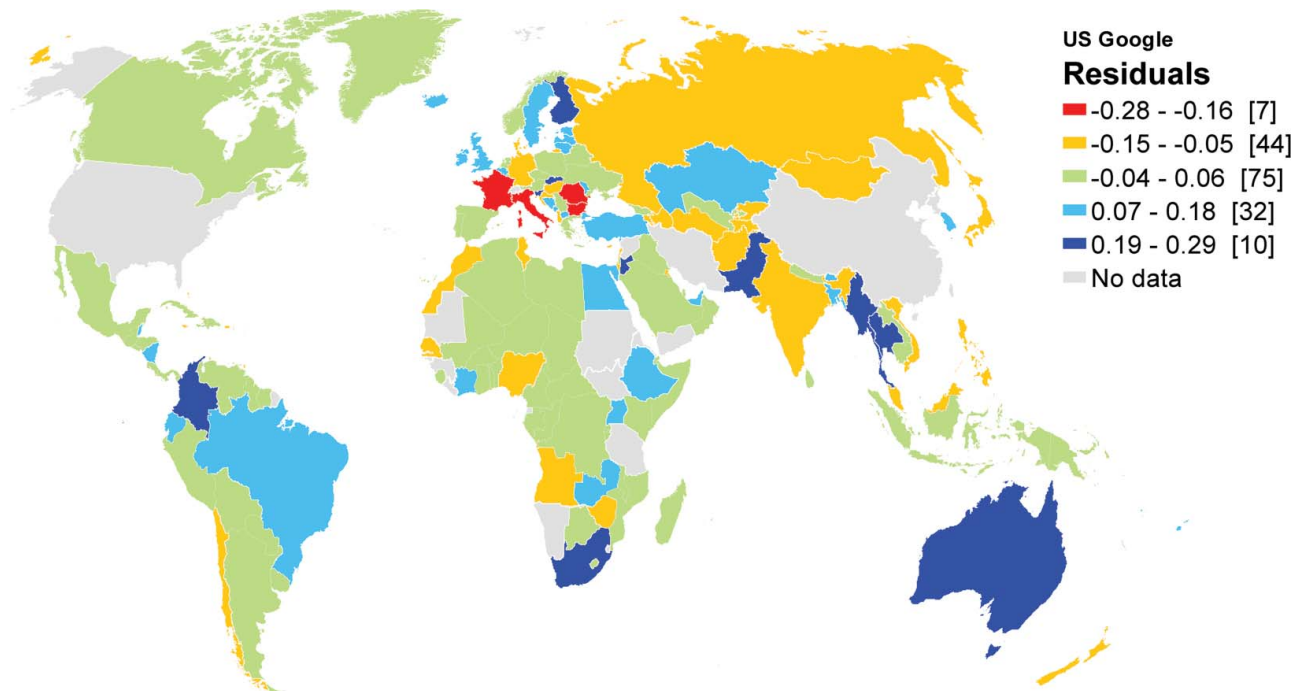


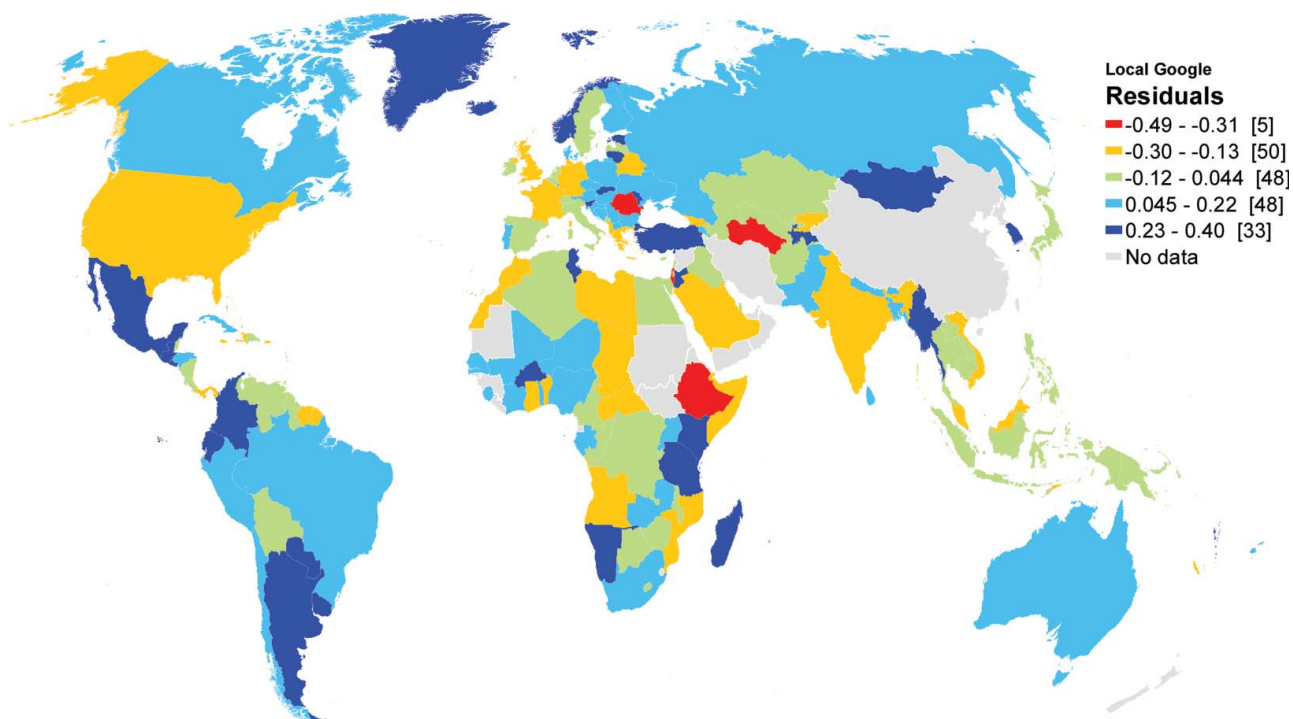
Figure 12. Residuals of regression model for U.S. Google (model M1). (Color figure available online.)

Table 9. General linear model on local Google

| Data set: Local Google ($N = 340$) | Estimate | SE | t value | $\Pr(> t)$ |
|--------------------------------------|-----------|----------|-----------|-------------|
| (Intercept) | 1.12E-01 | 1.36E-01 | 0.82 | 0.41 |
| h -index | 4.72E-04 | 8.83E-05 | 5.35 | 1.65E-07*** |
| Region: North America | 0 | 0 | 0 | N/A |
| Region: East Asia & Pacific | 1.48E-02 | 1.34E-01 | 0.11 | 0.91 |
| Region: Europe & Central Asia | 1.65E-01 | 1.31E-01 | 1.26 | 0.21 |
| Region: Latin America & Caribbean | 6.19E-01 | 1.36E-01 | 0.45 | 0.65 |
| Region: Middle East & North Africa | -9.66E-03 | 1.36E-01 | -0.07 | 0.94 |
| Region: South Asia | 6.60E-02 | 1.39E-01 | 0.47 | 0.63 |
| Region: Sub-Saharan Africa | 6.00E-02 | 1.34E-01 | 0.45 | 0.66 |
| Internet users | 1.90E-03 | 6.50E-04 | 2.92 | 0.004** |
| Language is local | 1.44E-01 | 2.37E-02 | 6.08 | 3.29E-09*** |

* $p < 0.05$.** $p < 0.01$.*** $p < 0.001$.**Table 10.** Analysis of variance test on local Google ($N = 340$) in Table 7

| Data set: Local Google | df | Deviance | Residual df | Residual deviance |
|------------------------|------|----------|---------------|-------------------|
| Null | | | 339 | 24.34 |
| h -index | 1 | 7.93 | 338 | 17.41 |
| Language is local | 1 | 2.33 | 337 | 15.08 |
| Region | 6 | 1.78 | 331 | 13.30 |
| Internet users | 1 | 0.33 | 330 | 12.97 |

**Figure 13.** Residuals of regression model for local Google (model M1). (Color figure available online.)

Similarly, seventy-two records appear in the local Google data set for countries where Google supports English but it is not an official language of the country. Thus, the overlap between the two data sets might partly explain the similarity in the regression models presented earlier.

Although our analysis selected the same four explanatory variables as most important for the local and U.S. analyses, a more careful analysis finds that the models differ significantly. To determine this, we compared two nested models on all 511 records in the combined U.S. and local data set. The *reduced model* included all ten variables. The *full model* included the ten variables, along with interactions between those variables and a Boolean variable indicating whether the record was either a U.S. search result or local search result. An ANOVA comparing the models found that the full model explained significantly more variance (68 percent as opposed to 54 percent), and the models differed significantly ($p < 0.0001$). This indicates that the role of the explanatory variables does differ between the U.S. and local data sets.

A closer inspection of the nested model found two key differences between local and U.S. interactions. First, in the full model, the L values for North America were significantly higher relative to other regions in the U.S. Google results compared to the local Google results. This reflects the United States' close cultural and geographic relationship to Canada (the only North American country in the U.S. Google results). Because the local results were viewed through a less U.S.-centric lens, regions' coefficients were similar across regions. Second, the h -index exhibited a coefficient twice as large for the local search results. We verified this finding in a minimal model with an interaction between the h -index and the Boolean U.S. versus local factor. This result indicates that although a country's scholarly publishing network is a critical feature for predicting L , the importance of a country's scholarly network is dampened when viewed through a U.S. search lens.

To summarize, we find that geographic region, strength of scholarly publishing, Internet infrastructure, and language barriers all play important roles in shaping the geography of Google Search results. These four variables explain roughly half of the variation in L , leaving much room to evaluate the role of other country-specific attributes. We note that these explanatory variables do exhibit significant collinearities, making it difficult to neatly unpack the roles of individual factors. We also find

significant differences between the explanatory models for the U.S. and local results; countries not in North America exhibit much higher L values in local results, and a country's scholarly network is more predictive of L in local results.

Conclusion

This investigation of the geography of Google Search results shows that wealthy and well-connected countries tend to have much more locally produced content that is visible about them than poor and poorly connected countries. Even cities located in countries with huge populations such as Lagos show a tendency toward having relatively little local content about them in Google Search results. This means that a user in the United States or Germany searching for cities is far more likely to be given access to locally produced content than a Tanzanian or Cambodian.

In our empirical study, the results of only eight countries in Africa (and four low-income countries: Tajikistan, Madagascar, Burkina Faso, and Tanzania) have a majority of content that is locally produced. This gives rise to a form of digital hegemony, whereby producers in a few countries get to define what is read by others. The United States, in particular, is a dominant content-producing force, even when excluding Wikipedia, which is a highly visible U.S.-based but globally assembled resource (Figures 10 and 11). In the results for sixty-one countries, the United States supplies over half of the first page content on Google. This means that not only are U.S. Internet users surrounded by an extremely locally produced Internet, but U.S.-produced content is highly visible in much of the rest of the world. This does not necessarily mean that the United States is an informational hegemon everywhere in the world, however. France has a somewhat smaller sphere of influence, mainly limited to countries in Africa, whereas Russia produces a visible effect only on results about Kyrgyzstan (see Figure 10).

It is important to note that, because our data set focused only on capital cities, caution should be taken when extending the results to higher spatial granularities. Our localness indicator does not take into account the actual interaction of users with search results and the variety of devices and media across which individuals currently access search engines. Despite the precautions that we took to access representative samples of search results, some noise is still present and some results might show high volatility.

Much more empirical work is needed to study finer patterns within countries and to build more accurate models to investigate the consumption of geographic information on search engines in different geographic locales.

More broadly, the point remains that most countries in the Global South continue to be defined by a diverse range of sources originating from a diverse range of places. The issue here is not that Internet users are exposed to a diverse range of sources from a diverse range of places—indeed, as Pariser (2011) noted, there are significant concerns for people and media ecosystems that lack access to such diversity. The issue is rather that that diversity itself has a particular bias and those sources tend to be almost entirely from the Global North; very few of the sources come from anywhere in the Global South. For instance, although the search results for Google’s Ghanaian page for its capital “Accra” include pages from six countries, five of them are firmly located in the Global North.⁸ When looking at countries in the Global North, the results for Denmark’s capital are similarly diverse, with five out of six source countries also being located in the Global North.⁹ By contrast, a country like the United States suffers from the inverse problem: having almost no exposure to geographic representations made by nonlocals.

The key question, then, is why. What explains this informational hegemony, or the dominance of the Global North in producing digital representations about not just themselves but also about much of the Global South? Interestingly, our explanatory models indicate that network connectivity and economic development in a country are not enough to make content about that place more local in Google Search results. The presence of a strong publishing industry, using SciMago publication data as proxy, is the strongest predictor of the production of visible online content. The importance of the *h*-index in the model also shows that the impact of scientific publications is a better predictor of localness than the mere number of publications. Thus, we suggest that socioeconomic systems that produce high-quality research also tend to produce highly visible online content. There are no countries in the Global South that score well on such metrics, and there are consequently no countries in the Global South that play a major role in constructing contemporary Internet geographies.

Having taken a first step in this direction, more quantitative and qualitative research is needed to better understand why exactly scientific knowledge

production explains so much of the variance in Google’s local digital representations. More relational variables and different spatial granularities will have to be considered. Until then, though, we hope that the finding that wealth and network connectivity alone are not sufficient factors is worth demonstrating, especially for Internet activists who hope to bring about more genuinely participatory and representative digital environments. This point increasingly matters because places are ever more defined by their digital presences, and the ways in which places are represented digitally increasingly shape how people understand and reproduce those very places (Graham, De Sabbata, and Zook 2015). Google plays an enormous role in constructing these digital representations of places. Because of their dominant role in mediating a majority of the world’s Internet use and the fact that few people ever explore beyond a first page of search results, they essentially determine which digital augmentations of place are made visible or invisible, with tangible effects in the physical world.

This article demonstrated that Google Search results are actively reproducing new forms of informational hegemony around the globe. A few countries in the Global North play an inordinately large role in defining the digital augmentations of the Global South. Google’s methods for ranking and representing are notoriously opaque (Vaidhyanathan 2011; Graham et al. 2014), but we do know that two key factors come into play. First, much of the reason of the lack of local voice in the Global South is likely simply because the production of Internet content happens at a much lower rate compared to that in the Global North (Graham, De Sabbata, and Zook 2015). Second, since the company’s creation in 1998, Google’s algorithms have tended to favor highly central Web content: Pages linked to by a lot of other pages are prioritized, and those largely ignored are demoted in the rankings. This creates a worrying situation whereby it becomes difficult for those on the information peripheries to break out of their digital marginality.

Building on earlier research looking at the geographies of information, this article has analyzed not just where digital content comes from but how it is ranked in the world’s most powerful digital mediator. Much more will need to be done to understand not just the ways in which people are afforded voice about their own communities and countries but also the myriad factors that serve to amplify or constrain it. Until then, we hope that other research can use this article as a beginning to ask not just why some parts of the

world are denied locally produced representations but how we might bring about more representative and participatory digital augmentations of place.

Acknowledgments


The idea for this article emerged at the Specialist Meeting on Spatial Search at the University of California, Santa Barbara. The authors thank the Center for Spatial Studies for organizing and funding the event.


Funding

We received funding from the European Research Council under the European Union's Seventh Framework Programme for Research and Technological Development (FP/2007–2013) and ERC Grant Agreement No. 335716. This work is also funded by National Science Foundation Grant 1527173 and a grant for computational resources from Amazon.com, Inc.

ORCID

Andrea Ballatore  <http://orcid.org/0000-0003-3477-7654>

Mark Graham  <http://orcid.org/0000-0001-8370-9848>

Shilad Sen  <http://orcid.org/0000-0001-7419-1864>

Notes

1. The research data sets are available at <https://github.com/andrea-ballatore/SearchGeography>.
2. The research data sets are available at <https://github.com/andrea-ballatore/SearchGeography>.
3. For an example of the U.S. Google query for Zimbabwe, see <https://www.google.com/search?q=Harare&hl=en&gl=us&oe=utf8>.
4. The data collection was executed on 24 March 2015, 15 April 2015, 22 April 2015, and 9 May 2015.
5. We considered results to be invalid when the query was formulated inconsistently with the experimental design; for example, if the query language was Italian but the string was "Rome" instead of "Roma."
6. The relevant files are available at <https://github.com/shilad/geo-provenance/tree/master/py>.
7. Whois is a protocol that allows access to a store of the addresses of people and firms that register every domain name at <https://whois.icann.org>.
8. The geographic composition of the 183 URLs returned for the capital of Ghana (Accra) is as follows: United

States: 67, Ghana: 43, The Netherlands: 29, Switzerland: 20, United Kingdom: 20, Sweden: 4.

9. This is the geographic composition of the seventy-seven URLs returned for the capital of Denmark (Copenhagen): Denmark: 32, Faroe Islands: 22, United States: 15, Namibia: 4, Poland: 3, Norway: 1.

References

- Ballatore, A. 2014. The myth of the Digital Earth between fragmentation and wholeness. *Wi: Journal of Mobile Media* 8 (2). <http://wi.mobilities.ca/myth-of-the-digital-earth> (last accessed 8 April 2017).
- . 2015. Google chemtrails: A methodology to analyze topic representation in search engine results. *First Monday* 20 (7): 1–20.
- Ballatore, A., W. Kuhn, M. Hegarty, and E. Parsons. 2016. Special issue introduction: Spatial approaches to information search. *Spatial Cognition & Computation* 16 (4): 245–54.
- Benkler, Y. 2007. *The wealth of networks: How social production transforms markets and freedom*. New Haven, CT: Yale University Press.
- Brossard, D., and D. A. Scheufele. 2013. Science, new media, and the public. *Science* 339 (6115): 40–41.
- Bruns, A. 2008. *Blogs, Wikipedia, Second Life, and beyond: From production to produsage*. New York: Peter Lang.
- Chitika. 2013. The value of Google result positioning. <https://chitika.com/google-positioning-value> (last accessed 23 January 2017).
- comScore. 2015. comScore releases January 2015 U.S. desktop search engine rankings. <http://www.comscore.com/Insights/Market-Rankings/comScore-Releases-January-2015-US-Desktop-Search-Engine-Rankings> (last accessed 23 January 2017).
- Craig, W. J., and S. A. Elwood. 1998. How and why community groups use maps and geographic information. *Cartography and Geographic Information Systems* 25 (2): 95–104.
- Crampton, J. 2008. Will peasants map? Hyperlinks, map mashups and the future of information. In *The hyperlinked society: Questioning connections in a digital age*, ed. J. Turow and L. Tsui, 206–26. Ann Arbor: University of Michigan Press.
- Dinner, I., E. J. Johnson, D. G. Goldstein, and K. Liu. 2011. Partitioning default effects: Why people choose not to choose. *Journal of Experimental Psychology: Applied* 17 (4): 432.
- Dodge, M., and M. Zook. 2009. Internet-based measurement. In *The international encyclopedia of human geography*, ed. R. Kitchin and N. Thrift, 569–79. New York: Elsevier.
- Elwood, S. 2006. Critical issues in participatory GIS: Deconstructions, reconstructions, and new research directions. *Transactions in GIS* 10 (5): 693–708.
- Epstein, R., and R. E. Robertson. 2015. The search engine manipulation effect (SEME) and its possible impact on the outcomes of elections. *Proceedings of the National Academy of Sciences* 112 (33): E4512–21.
- Google. 2014. Google transparency report. <https://www.google.com/transparencyreport> (last accessed 23 January 2017).

- . 2017. Custom search vs. Google.com. <https://support.google.com/customsearch/answer/70392?hl=en> (last accessed 23 January 2017).
- Graham, M. 2010. Neogeography and the palimpsests of place. *Tijdschrift voor Economische en Sociale Geografie* 101 (4): 422–36.
- . 2013. The virtual dimension. In *Global city challenges: Debating a concept, improving the practice*, ed. M. Acuto and W. Steele, 117–39. London: Palgrave.
- Graham, M., S. De Sabbata, and M. A. Zook. 2015. Towards a study of information geographies: (Im) mutable augmentations and a mapping of the geographies of information. *Geo: Geography and Environment* 2 (1): 88–105.
- Graham, M., S. A. Hale, and M. Stephens. 2011. *Geographies of the world's knowledge*. London: Convoco!
- Graham, M., B. Hogan, R. K. Straumann, and A. Medhat. 2014. Uneven geographies of user-generated information: Patterns of increasing informational poverty. *Annals of the Association of American Geographers* 104 (4): 746–64.
- Graham, M., R. Schroeder, and G. Taylor. 2014. Re:Search. *New Media & Society* 16 (2): 187–94.
- Graham, M., R. Straumann, and B. Hogan. 2016. Digital divisions of labor and informational magnetism: Mapping participation in Wikipedia. *Annals of the American Association of Geographers* 105 (6): 1158–78.
- Gramsci, A. 1971. *Prison notebooks*. New York: International.
- Grimmelmann, B. J. 2010. Some skepticism about search neutrality. In *The next digital decade: Essays on the future of the Internet*, ed. B. Szoka and A. Marcus, 435–59. Washington, DC: TechFreedom.
- . 2013. What to do about Google? *Communications of the ACM* 56 (9): 28–30.
- Halavais, A. 2010. *Search engine society*. Cambridge, UK: Polity.
- Hannak, A., P. Sapiezynski, A. Molavi Kakhki, B. Krishnamurthy, D. Lazer, A. Mislove, and C. Wilson. 2013. Measuring personalization of web search. In *Proceedings of the 22nd International Conference on World Wide Web*, 527–37. Rio de Janeiro, Brazil.
- Harley, J. B. 1989. Deconstructing the map. *Cartographica: The International Journal for Geographic Information and Geovisualization* 26 (2): 1–20.
- Hillis, K., M. Petit, and K. Jarrett. 2013. *Google and the culture of search*. London and New York: Routledge.
- Introna, L. D., and H. Nissenbaum. 2000. Shaping the Web: Why the politics of search engines matters. *The Information Society* 16 (3): 169–85.
- Jenkins, H. 2006. *Convergence culture: Where old and new media collide*. New York: New York University Press.
- Kitchin, R., and M. Dodge. 2011. *Code/space: Software and everyday life*. Cambridge, MA: MIT Press.
- König, R., and M. Rasch, eds. 2014. *Society of the query reader: Reflections on Web search*. Amsterdam: Institute for Network Cultures.
- Leszczynski, A. 2012. Situating the geoweb in political economy. *Progress in Human Geography* 36 (1): 72–89.
- Lessig, L. 2003. An information society: Free or feudal. *Lecture given at the World Summit on the Information Society*, Geneva, Switzerland. <http://www.itu.int/wsis/docs/pc2/visionaries/lessig.pdf> (last accessed 14 April 2016).
- Levine, J. M., and J. HilleRisLambers. 2009. The importance of niches for the maintenance of species diversity. *Nature* 461 (7261): 254–57.
- Long, J. S. 1997. *Regression models for categorical and limited dependent variables*. Thousand Oaks, CA: Sage.
- Mager, A. 2012. Algorithmic ideology: How capitalist society shapes search engines. *Information, Communication & Society* 15 (5): 769–87.
- National Aeronautics and Space Administration. 2015. “Deep Web search” may help scientists. <https://www.nasa.gov/jpl/deep-web-search-may-help-scientists> (last accessed 23 January 2017).
- Nielsen, J. 2005. The power of defaults. <http://www.nngroup.com/articles/the-power-of-defaults> (last accessed 8 April 2017).
- Ostry, M. J. D., M. A. Berg, and M. C. G. Tsangarides. 2014. *Redistribution, inequality, and growth*. Washington, DC: International Monetary Fund.
- Pariser, E. 2011. *The filter bubble: What the Internet is hiding from you*. New York: Penguin.
- Pickles, J., ed. 1995. *Ground truth: The social implications of geographic information systems*. New York: Guilford.
- . 2004. *A history of spaces: Cartographic reason, mapping, and the geo-coded world*. Hove, UK: Psychology Press.
- Said, E. 1978. *Orientalism*. New York: Pantheon.
- Sen, S., H. Ford, D. Musicant, M. Graham, O. Keyes, and B. Hecht. 2015. Barriers to the localness of volunteered geographic information. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, 197–206. New York: ACM.
- Shirky, C. 2011. *Cognitive surplus: Creativity and generativity in a connected age*. London: Penguin.
- Spink, A., and M. Zimmer. (Eds.). 2008. *Web search: Multidisciplinary perspectives* (Vol. 14). Berlin: Springer.
- Stekhoven, D. J., and P. Bühlmann. 2012. MissForest—Non-parametric missing value imputation for mixed-type data. *Bioinformatics* 28 (1): 112–18.
- Tapscott, D., and A. D. Williams. 2006. *Wikinomics: How mass collaboration changes everything*. New York: Penguin.
- theEword. 2015. Search engine market. http://theeword.co.uk/info/search_engine_market (last accessed 23 January 2017).
- Vaidhyanathan, S. 2011. *The Googlization of everything (and why we should worry)*. Berkeley: University of California Press.
- Vaughan, L., and M. Thelwall. 2004. Search engine coverage bias: Evidence and possible causes. *Information Processing & Management* 40 (4): 693–707.
- Wouters, P., and D. Gerbec. 2003. Interactive Internet? Studying mediated interaction with publicly available search engines. *Journal of Computer-Mediated Communication* 8. doi:10.1111/j.1083-6101.2003.tb00221.x

ANDREA BALLATORE is a Lecturer in GIS and Big Data Analytics in the Department of Geography, Birkbeck, University of London, London WC1E 7HX, UK. His research interests include volunteered geographic information,

Internet geography, geographic information retrieval, geo-semantics, and media studies.

MARK GRAHAM is the Professor of Internet Geography at the Oxford Internet Institute, University of Oxford, Oxford OX1 3JS, UK. He is also a Faculty Fellow at the Alan Turing Institute, a Visiting Fellow at the London School of Economics and Political Science, an Associate in the University of Oxford's School of Geography and the Environment, and a Senior Research Fellow at Green

Templeton College. His current research focuses on digital labor, digital geographies, and inequalities in our digital world.

SHILAD SEN is Associate Professor of Computer Science at Macalester College, St. Paul, MN 55015–1801. He is also a Data Science Research Fellow at the Target Corporation. His research describes online dynamics, particularly those related to inequality, and also develops new systems that mine patterns in online behavior.