MapWatch: Detecting and Monitoring International Border Personalization on Online Maps

Gary Soeller
Northeastern University
soelgary@ccs.neu.edu

Karrie Karahalios
University of Illinois at Urbana-Champaign
kkarahal@illinois.edu

Christian Sandvig
University of Michigan
csandvig@umich.edu

Christo Wilson
Northeastern University
cbw@ccs.neu.edu

ABSTRACT

Maps have long played a crucial role in enabling people to conceptualize and navigate the world around them. However, maps also encode the world-views of their creators. Disputed international borders are one example of this: governments may mandate that cartographers produce maps that conform to their view of a territorial dispute.

Today, online maps maintained by private corporations have become the norm. However, these new maps are still subject to old debates. Companies like Google and Bing resolve these disputes by localizing their maps to meet government requirements and user preferences, i.e., users in different locations are shown maps with different international boundaries. We argue that this non-transparent personalization of maps may exacerbate nationalistic disputes by promoting divergent views of geopolitical realities.

To address this problem, we present MapWatch, our system for detecting and cataloging personalization of international borders in online maps. Our system continuously crawls all map tiles from Google and Bing maps, and leverages crowdworkers to identify border personalization. In this paper, we present the architecture of MapWatch, and analyze the instances of border personalization on Google and Bing, including one border change that MapWatch identified live, as Google was rolling out the update.

1. INTRODUCTION

Maps have long played a crucial role in enabling people to conceptualize and navigate the world around them. In ancient societies, birds-eye projections of the world facilitated the administration of cities and empires, as well as long distance travel, long before the advent of satellite imagery and GPS. To this day, printed (and now digital) maps allow us to conceptualize the geographic, spatial, physical, and political features of the world.

WWW 2016, April 11–15, 2016, Montréal, Québec, Canada.
ACM 978-1-4503-4143-1/16/04.
http://dx.doi.org/10.1145/2872427.2883016 .

However, cartography is not a purely objective endeavor: maps encode the world-views of their creators. For example, Renaissance maps of Europe commissioned by the wealthy elite codified social structures centered around palaces, churches, and feudal divisions [31]. Similarly, maps produced during the World Wars contained intentional inaccuracies as fuel for propaganda [32] and misdirection [30, 31].

One acute problem that modern cartographers grapple with is the portrayal of disputed international borders. Accurately displaying uncertainty is challenging from a technical perspective, but more importantly the choice to display uncertainty *at all* is a value judgement that the countries involved may disagree with. Indeed, some countries legislate how their borders must be rendered, regardless of international consensus or the actual situation on the ground.

Anecdotal evidence shows that modern, online mapping services deal with conflicting territorial claims through *personalization*. Specifically, users may be shown localized maps with different international borders. Users have noticed this phenomenon in the past around the borders of China and India in Google Maps [3, 20, 24, 46, 51]. This personalization occurs automatically, there is no indication on the map that alternate localizations exist, and there is no option to turn the personalization off.

We argue that this non-transparent personalization of maps is problematic in two respects. First, since users are unaware of the personalization, it may exacerbate nationalistic disputes by reinforcing divergent views of geopolitical realities. Second, it is troubling that private corporations have become the primary arbiters of geographic information to the public, yet we have no idea when or why these companies choose to alter or personalize maps.

To address these problems, we present MapWatch, our system for detecting modifications to and personalization of international borders in online maps. Our system continuously crawls *all* map tiles from Google and Bing Maps from the perspective of users in 250 different countries and territories. We observe hundreds of tile updates per week, far too many to manually analyze, so MapWatch leverages crowdworkers to identify border changes in map tiles. We show that crowdworkers achieve ~98% labelling accuracy and only cost \$298/month, which makes the process of labelling thousands of tile pairs tractable.

MapWatch has been in operation since January 2015, and in that time we have detected the seven instances of border personalization shown in Figure 1. This includes two new

Figure 1: Personalized regions detected by MapWatch. Country names denote where users are shown the personalized tiles.

instances that have never been documented before. MapWatch detected one of these new personalized borders in Georgia, as Google Maps altered the border in May 2015, revealing interesting details about how Google stages rollouts of updated tiles. MapWatch is the first tool of its kind that temporally monitors the entire world of multiple online maps. It becomes a tool to view world border conflicts, allowing us, or other researchers, to see such conflicts evolve over time via an archive of cartographic artifacts: a kind of Wayback Machine for the world.

We view MapWatch as a service to the public, and will continue publishing border changes on Google and Bing Maps on our project website. Furthermore, the MapWatch source code is open source.[1] Finally, we view MapWatch as an extensible platform, and plan to continue improving it over time, e.g., by adding support for additional map providers like OpenStreetMaps.

Roadmap. The remainder of this paper is organized as follows: in §2, we present background information about cartography and auditing on the Web, and frame the concerns around the national border personalization of online maps. In §3, we introduce the details of MapWatch, followed by analysis of border personalization on Google and Bing Maps in §4. We conclude by discussing the implications of our findings in §5.

2. BACKGROUND

2.1 Borders and Web Cartographers

National borders and computer software have a complicated history. In one infamous incident, a system dialog in Microsoft Windows 95 asked users to select a time zone by displaying a map. Microsoft received a complaint from the government of India about the depiction of the Kashmir border, and responded by removing all maps from the operating system [38]. Software companies later employed "geopolitical product strategy teams" with expertise in international relations to make corporate policy about sensitive map features. In 2000, a Microsoft news release described the value of these teams as primarily one of error correction and harmonization—for instance, ensuring that all Microsoft products correctly identified Mexico as part of North America and not Central America [38].

[1] Visit `http://personalization.ccs.neu.edu` for code and up-to-date examples of personalized map tiles.

Initially, Web platforms typically offered one version of their data to visitors. In the early 2000s major content platforms on the Web "grounded" their products (in Rogers' term [44]), creating national versions to unlock advertising markets by geography. Governments began to scrutinize these localized offerings for compliance with national laws.

National governments produce and often require cartography that is aligned with their geopolitical policies. For decades after India and Pakistan reached a cease fire over the disputed territory of Kashmir, the cease fire boundary appeared on maps made across the world, but it did not appear in India or Pakistan, where each country's printed maps claimed the entire territory [41]. In many countries, maps that admit the existence of disputes are illegal; by law, Chinese companies must produce maps that show Arunachal Pradesh (*a.k.a.* South Tibet) as part of China [20].

In 2010, a Google statement admitted that the depiction of national borders was an ongoing challenge, and emphasized "accurate" maps, implying a single correct depiction of the world was the goal [21]. In the words of a Google executive, "We work to provide as much discoverable information as possible so that users can make their own judgments" [24]. Indeed, the Google Earth product employs a data model that divides national borders into four types: *internationally recognized* ("legal and final"), *disputed* by one or both adjacent entities (whether "uncontested, dormant, jointly-managed, hostile, bellicose...[or at] war"), *treaty* (established by a particular agreement, but not more widely recognized), and *provisional* (a *de facto* boundary that is widely recognized, but not legal). A dialog box with an impressive 17 lines of text explains the territorial status of Arunachal Pradesh in Google Earth [20]. These nuances do not appear in consumer grade mapping products such as Bing Maps or Google Maps.

Instead, consumer grade Web mapping sites operate localized world maps to comply with the law and territorial sensitivities. In 2010, Google stated that it operates 32 versions of the World Map [24]. Today, it is possible to request over 250 localizations from the Google Maps API, but it is unknown how many of these contain localized borders [22].

Other researchers discovered that language and expressed preference personalization can affect the display of geopolitically problematic places, irrespective of the user's location. [11] notes that MSN Livesearch returned different place markers for "Al Quds", Jerusalem's Arabic name, and the query "Jerusalem". Mapquest returned Palestinian cities only if they were queried phonetically by their Arabic names. Flickr asked users who queried "Jerusalem" whether they meant "Israel" or "Occupied Palestinian Territories," but then provided the same data regardless of their selection.

2.2 Maps, Politics, and Authority

These differences are significant because, in Alfred Korzybski's slogan, "the map is not the territory." Geographers and historians emphasize that a map advocates a way of thinking about space, rather than transmitting the single correct representation. Indeed, mapping technology itself—including the prior technology of the printed map—privileges a particular cognitive perspective [9].

In an influential example, the Siamese historian Winichakul described the introduction of the first maps containing a national boundary to Siam. Before 1851, Siamese maps were hand-drawn and lacked scale, containing only religious or navigational information. Boundary stones marked borders, but these stones sat on the ground and not on paper. Winichakul noted that the printed national boundary arrived in Siam with printing itself [9]. The idea of a fixed, printed national boundary served as an advertisement for the Western colonial idea of a nation, and the shapes drawn were often as aspirational as they were indexical. Printed borders were "an intellectual tool for legitimating territorial conquest" [41]. Web maps are thus our present "intellectual tool," with implications that are still emerging. Map software developers are often unaware of this legacy and of the broader universe of cartographic conventions they are implementing or avoiding [49].

Over twenty years ago, Monmonier warned that computer mapping would perturb "distinctions between mapmakers and map users" because an interactive map is produced by a user's own actions as much as by a distant, authoritative cartographer [42]. Indeed, as many cartographic platforms now include volunteered geographic information, users can suggest new borders, or annotate existing ones. (As of this writing, the disputed Spratly Islands in the South China Sea are illustrated on Google Maps by uploaded, geotagged photographs labeled "belong to Vietnam.")

Prior scholarship has emphasized the potential of dynamic or selective maps to modify our perception of space [27], but work has largely focused on navigation systems [55]. Cultural geographers argued algorithmic map displays tend to produce benefits that are unequally available [23], e.g., a national boundary produced by an automated system "challenges the whole concept of 'national borders' for some, while making such borders more rigid...for others" [59].

So far, provocative investigations have examined case studies of specific border disputes and specific platforms [3, 11]. These studies occurred against the backdrop of increasing popular awareness of international border personalization of online maps [46, 51] and concern about the implications of geographic personalization more generally [60].

2.3 Auditing Algorithms

Personalized maps are an instance of the more general problem of research into "black-boxed" Web systems. Complex algorithmic systems can easily create outcomes that are unforeseeable by their designers [50], and indeed mapping platforms have been plagued by high-profile errors detected by outsiders [20]. More fundamentally, researchers have argued for "auditing algorithms" from the outside in situations where a proprietary, secret process can produce undesirable societal consequences, whether these are intended or not [48]. Although true "reverse engineering" of a black-boxed system may not be possible, it may also be unnecessary, as detecting undesirable behavior may be enough to convince a platform to change it [16]. Auditing in this sense takes its name from the social scientific "audit study" where one feature is manipulated in a field experiment [47, 48].

Numerous recent studies have pursued this approach. In the realm of social networks, [18, 19] examined user perceptions of algorithmically curated content vs. uncurated content in the Facebook News Feed. Inspired by concerns about the "Filter Bubble", [28, 35] examined Google's search personalization algorithm. Three studies have examined personalization on e-commerce sites, and demonstrated that many e-retailers implement price discrimination and price steering [29, 39, 40]. [12] examined Uber's dynamic "surge pricing" algorithm and discovered fairness issues. Even more alarmingly, [17,53] identify racial discrimination on AirBNB and Google Ads. Finally, [25,36,37] investigate online tracking techniques leveraged by advertising networks.

As a small number of commercial mapping platforms increasingly dominate Web cartography, these platforms are ripe for systematic auditing to determine the extent of map personalization and its consequences.

3. METHODOLOGY

To detect border personalization on online maps, we need a system that continuously collects and processes map tiles. This is a daunting task, given the huge volume of tiles, the need to check them regularly for updates, and the labor necessary to identify personalized borders in images.

In this section, we describe MapWatch, our system that meets these goals. First, we briefly sketch the mechanisms used by online mapping services to personalize content. Second, we discuss how MapWatch crawls map tiles from Google and Bing Maps. Third, we describe the *baseline* set of personalized tiles that we manually identified in our crawled data. Finally, we explain how MapWatch detects new personalized and global border updates over time.

3.1 Localization Mechanisms

Under typical circumstances, map providers automatically set the default localization for users by inferring their geolocation. This can be accomplished in principle by geolocating the user's IP address, examining the default language on the user's computer (e.g., using the `Accept-Language` HTTP header), or accessing sensors in the user's device (e.g., GPS). In practice, the exact mechanisms used by Google and Bing Maps to infer users' locations are unknown. Anecdotal evidence suggests that Google Maps does use IP address geolocation: we contacted collaborators who live in China, but subscribe to an ISP in Hong Kong, and they reported seeing the Hong Kong-localization.

Google Maps does not offer users the option to explicitly choose the map localization they wish to view[2]. Google users *may* be able to influence the localization they are shown by choosing a specific language or region in their map settings, or by visiting specific web domains (e.g., `maps.google.es`

[2]Software developers can access any localization programmatically using the Google Maps API, but these tools are not useful to typical web users.

versus `maps.google.com`). Bing Maps does allow users to explicitly set their localization, but this requires changing the default settings.

3.2 Crawling Google and Bing

The first challenge in building MapWatch was developing a method to crawl online maps. Specifically, we must crawl *all* tiles from a given mapping service from the perspective of *every* country around the world. For reasons that will become clear, we focus on the 250 countries and territories that have been assigned ISO 3166-1 alpha-2 country codes (e.g., *de* for Germany).

For Google Maps, we use the Google Maps API [22] to crawl tiles. The Maps API allows clients to fetch arbitrary tiles by specifying a desired latitude, longitude, zoom level, and an ISO 3166 two-letter country code. The country code corresponds to tiles that are shown on a localized version of Google Maps, e.g., specifying *de* retrieves tiles that are shown on `https://www.google.de/maps/`. Using this API, we crawl all tiles from all country/territory perspectives by iterating through spatial coordinates and ISO 3166 country codes. We set the zoom level to 6, which corresponds to a scale of about 87km per centimeter on a computer screen for tiles at the equator [6]. As shown in Figure 2, tiles at this zoom level provide an appropriate level of detail to evaluate country-scale features like borders.

To collect map tiles from Bing Maps, we began by using the Bing Maps API. This API, however, enforced strict rate limits when not using the (very expensive) enterprise version. We then discovered that all Bing Map tiles are directly accessible from Microsoft's Content Delivery Network (CDN) via URLs that include two parameters: a unique tile ID and an ISO 3166 country code. Through the CDN, we manually determined the range of IDs that corresponded to tiles matching the Google Maps zoom level 6. Given the tile IDs and country codes, we crawled all tiles from all possible perspectives on Bing Maps.

We exhaustively crawled Google and Bing Maps and indexed each tile in a database using its location, country code, and MD5 image hash value. We then recrawled each service weekly, comparing the MD5 hash value of each new tile against our historical data to find updated tiles; we stored and indexed only the updated tiles. In total, MapWatch fetches 156,060 tiles from Google Maps during each crawl; this process takes 6 days, due to rate limits imposed by the Maps API. MapWatch fetches 65,280 tiles from Bing Maps during each crawl, which takes 2.5 days. We chose the period of a week for our periodic crawls to maintain consistency across our Google and Bing datasets.

Ethics. Many Web platforms discourage all automated crawling outside the API, and even within the API a platform may prohibit "research" in its online Terms of Service document (usually for competitive reasons). We agree with [57] that non-commercial research for the public good that deals with issues of societal importance must be able to access public Web resources for research purposes as long as automated processes do not produce an unreasonable load. This was our guiding philosophy in this research design.

Alternate Crawling Strategies. The *Disputed Territories* project [3] was the first to attempt to automatically curate personalized map tiles from Google Maps. To avoid crawling all map tiles, the system focuses on areas within the Natural Earth Data disputed territories shapefiles [7].

Figure 2: Screenshot of one page from a MapWatch HIT, displaying the response form and GIF of the map tiles.

We made an explicit choice not to follow the crawling strategy used by *Disputed Territories* for two reasons. First, the Natural Earth Data shapefiles may be incomplete. Second, online map services may personalize or otherwise alter borders for reasons unrelated to disputes. Thus, we choose to crawl all tiles in MapWatch, and focus on scaling our system to handle the increased number of tiles.

3.3 Obtaining a Baseline Set of Tiles

The goal of MapWatch is to temporally monitor online mapping services and detect border changes. As our initial baselines, we used the tiles from our first full crawl in January 2015 for Google and Bing Maps.

To develop the baseline set, we compared the MD5 hashes of all pairs of map tiles from the same coordinates but with different country codes. A discrepancy in MD5 hash values from two tiles at the same coordinates indicates that users in the two different countries see different tiles for the same region. In total, this process uncovered 24 unique region discrepancies on Google Maps and 12 on Bing Maps.

Discrepancies in tile pairs, however, do not always imply altered international borders. For example, in some cases, city names are rendered in a country's dominant language, while the borders remain the same. To filter non-border discrepancies, we manually analyzed the divergent images. Ultimately, we found 15 region discrepancies on Google Maps and nine on Bing Maps resulting from personalized borders. These cases formed our baseline sets; we analyze them in detail in §4. Of the remaining Google tiles, ten contained differences in city or state/province names, and one contained a small difference in a body of water. Similarly, the remaining two tiles from Bing also contained differences in city names and bodies of water.

3.4 Continuous Tile Monitoring

The final step in the development of MapWatch was the detection of border changes over time. For the initial baselines, we were able to manually label the 37 divergent locations. However, between January and September 2015, we discovered 8,283 updated tiles on Google Maps and 612 on Bing Maps. Upon manual examination of a small subset of these updates, we determined they were not due to border changes. Instead, the changes were geographic alterations (e.g., changes to small bodies of water and forests) or slight shifts in the locations of cities. We could not manually label this volume of tiles, and thus explored alternate strategies to make this task more feasible.

First, we attempted to filter out tiles if the number of differing pixels fell below a set threshold. This approach failed as almost all of the tiles had a significant number of altered pixels. *Second*, we used Optical Character Recognition (OCR) to filter out tiles that differed due to textual changes (e.g., the addition of a city name). Unfortunately, OCR performs best on high-contrast images (e.g., black text on a white background); OCR over our map tile corpus resulted in low recall, and was therefore not feasible.

Crowdsourcing. We ultimately included a crowdsourced tile labelling pipeline into MapWatch. Specifically, after each completed crawl, the updated tiles automatically inhabit Human Intelligence Tasks (HITs) and the HITs post to Amazon Mechanical Turk (AMT) [8]. Each HIT consists of ten updated, unlabelled tiles and one *control tile*. Each page of the HIT displays an animated GIF alternating between two tiles and a form asking if there is a border change between the GIF tiles (see Figure 2). The geographically synced tiles in the alternating animated GIF enable easy and timely identification of border differences. Workers select *yes* or *no*, then move to the next page, until all eleven tile pairs have been analyzed.

We paid each worker \$0.85 per HIT. Workers were allowed ten minutes to complete each HIT; most HITs were completed within 2–3 minutes. We did not collect any personal information from workers.

Quality Control. To ensure the integrity of the tile labels collected from AMT, we implemented three standard quality control techniques. *First*, each HIT contains a control tile pair with a border change (manually identified by the authors). We use this control to discard unreliable data (e.g., from workers who answer *no* to all questions, or answer questions randomly). We reject a HIT if the control question is answered incorrectly, and assign the HIT to a new worker. Workers who answer *no* to all questions are blocked from completing further HITs. *Second*, each unlabelled tile is added to three HITs; if at least one worker identifies a border change, we manually analyze the tile. *Third*, we only accept workers with a 97% or higher approval rate to participate in our HITs.

Between January and September of 2015, MapWatch generated 1,638 HITs. Of those, we approved 1,533, yielding an 95.6% success rate. We reposted the rejected HITs; 100% were approved on the second attempt.

Validation. To assess the effectiveness of the crowdsourced workers, we examined the quality of their labels. Of the 4,914 pairs of analyzed tiles, workers identified 127 true positives (i.e., tiles with border updates) and 298 false positives. Although the precision of the workers is only 30%, this represents a 91.4% reduction in our workload, since we manually verify all potential border updates.

To estimate the workers' overall accuracy, we randomly examined 200 tile pairs and compared our ground-truth labels to the workers' labels. We found that 195 of the pairs were true negatives and five were false positives, thus suggesting that the workers' accuracy is ∼98% for our system.

Overall Results. During the nine months of MapWatch's operation, we spent ∼\$298 per month on HITs and ∼1.5 hours per month approving HITs and manually verifying potential border changes labelled by workers. MapWatch successfully detected 62 unique border updates on Google Maps and 27 on Bing Maps. Although most of these updates were minor, one was politically significant; we discuss this change in detail in §4.2.

4. ANALYSIS

In this section, we analyze the cases of border personalization detected by MapWatch, starting with results from our baseline crawl, and finishing with results detected by crowdworkers during the continuous crawls.

4.1 Baseline Border Personalization

Our baseline set is composed of 24 tile pairs containing borders personalized for specific countries. These correspond to six different conflicts. In the rest of this section, we describe the personalization in each location and briefly explain the history of the conflict.

Sino-Indian Border Conflict. The Sino-Indian War in 1962, fought between China and India, resulted from border disputes in the Aksai Chin and Arunachal Pradesh regions. To this day there is no division of territory that appeases all sides. Pakistan also claims territory in the area around the Northwest region of Kashmir. Laws in both India and China mandate that maps display their respective territorial claims [3, 20, 24], and thus both Google and Bing Maps personalize the borders for users in these two countries. Figure 3 displays the respective personalized borders around Arunachal Pradesh on Google and Bing Maps.

Indo-Pakistani Border Conflict. The Indo-Pakistani War in 1947 was the first time India and Pakistan fought over territorial control of the Kashmir region [4]. Since 1947, the two sides have fought several other wars over the same land. Despite diplomatic attempts to resolve the issue, an agreement never materialized as China did not accept the terms after the Sino-Indian War. Figure 4 shows the personalized borders around the Aksai Chin region.

Crimea Annexation. In March 2014, in the aftermath of the Ukranian revolution, Russia secretly invaded and annexed Crimea. At the end of March 2014, the UN passed a general assembly resolution [15] stating the invalidity of the annexation and asking all international organizations and agencies not to recognize Crimea as a part of Russia. However, as of April 2014, the localized Russian Google Maps view includes Crimea within its borders [3,46,51]. As shown in Figure 5, Russians see a solid border indicating that Russia has sovereignty over the peninsula, while Ukrainians see no border. The rest of the world sees a dashed border indicating that control over Crimea is disputed.

Argentine Antarctica. In 1904, Argentina became the first country to set up permanent occupation in Antarctica [1]. The Argentinian government claimed land between the 25 and 74 West meridians and between the 60 and 90 South parallels. The United Kingdom and Chile also lay claim to parts of this territory. Due to the *Antarctic Treaty System* signed by each of these countries, to date, no military conflict exists. The treaty emphasized the preservation of Antarctica for scientific research. Figure 6 shows the bor-

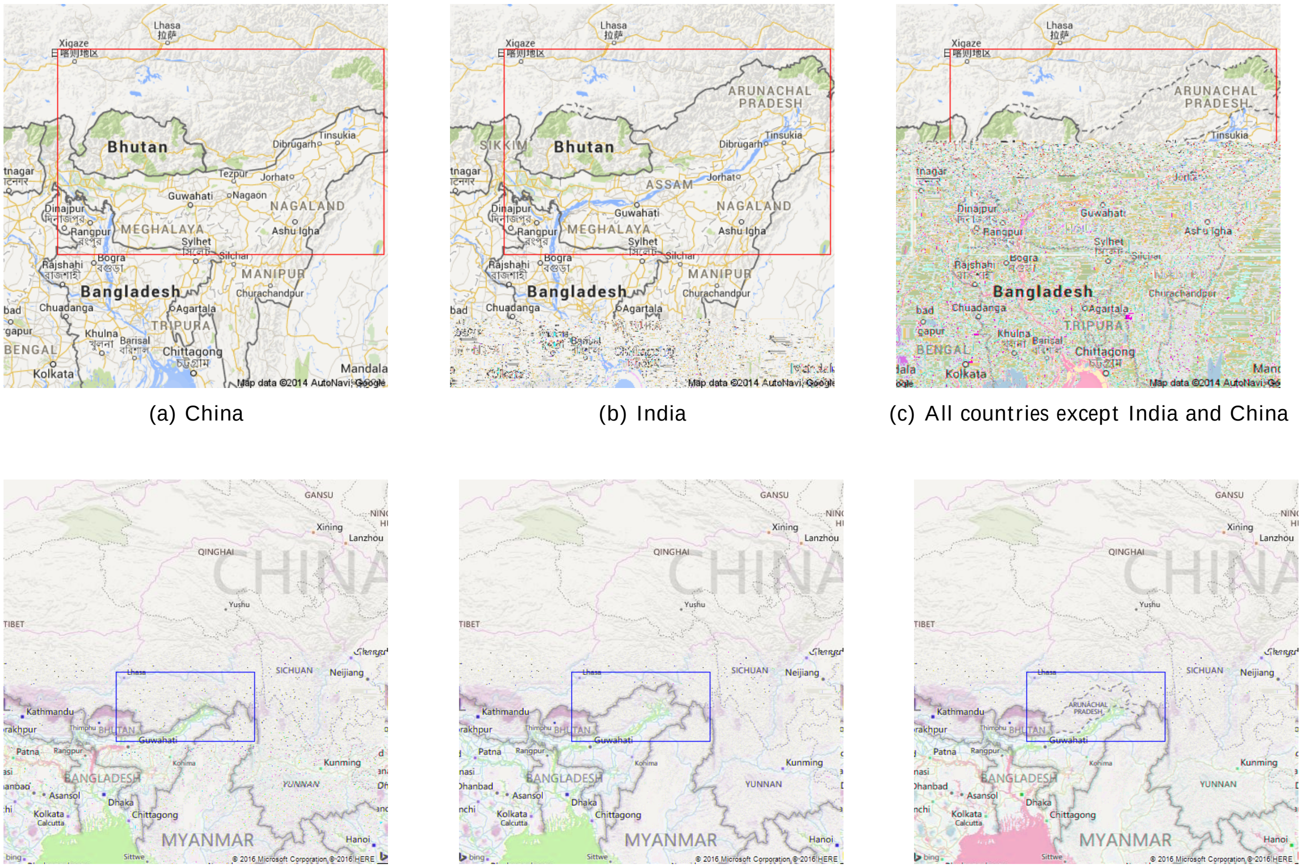

(a) China (b) India (c) All countries except India and China

(d) China (e) India (f) All countries except India and China

Figure 3: Border changes around Arunachal Pradesh between India and China on Google (top) and Bing Maps (bottom). Figures 3a to 3c: Map Data ©2014 AutoNavi Google. Figures 3d to 3f: Microsoft product screen shots reprinted with permission from Microsoft Corporation.

,

der along the longitudinal lines in Antarctica claimed by Argentina. Interestingly, these borders appear only to Argentinian users on Bing Maps; other countries do not see them. Google Maps does not display these borders at all.

South China Sea. In 1940, the Chinese government released a map displaying the original *nine-dashed line*. The *nine-dashed line* marks a border around the edge of the South China Sea labelling this as Chinese territory. This

(a) China

(b) India

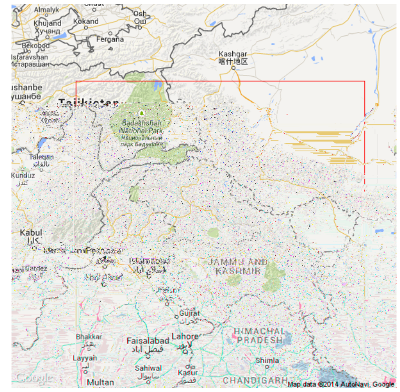

(c) All countries except India and China

(d) China

(e) India

(f) All countries except India and China

Figure 4: Border changes around Jammu, Kashmir, and Himachal Pradesh on Google (top) and Bing Maps (bottom). Figures 4a to 4c: Map Data ©2014 AutoNavi Google. Figures 4d to 4f: Microsoft product screen shots reprinted with permission from Microsoft Corporation.

The Strange Case of Ascension Island. Given the nature of the conflict in Georgia, it is not immediately clear why Google Maps shows users from Ascension Island, an island in the South Atlantic between Brazil and Africa, a personalized border in Georgia. In fact, closer analysis reveals that 62% of all temporal border updates detected by MapWatch on Google Maps first appear within the localization for Ascension Island.

Considering that Ascension Island has <1000 residents (implying that controversial tile updates become visible to a small population), this suggests that Google may be using the tiles in the Ascension Island localization as a testbed for map changes before large-scale roll-outs. In fact, it is unknown whether Google automatically directs any users to the Ascension Island-localized version of maps; the tiles may only be served in practice through the Google Maps API. Additionally, we suspect that Google chose the *ac* country code over similarly obscure country codes like *bv* (uninhabited Bouvet Island) or *cx* (Christmas Island) because it is first alphabetically among all ISO 3166 country codes.

5. DISCUSSION

This research found that online Maps can display very different views of the world to different users. For instance, we expect that during this study, two Internet users in China who were located in the same building could see different depictions of China's border. This would occur if they subscribed to two different ISPs (due to ISP-based geolocation) or if they used different online map platforms (Google Maps vs. Bing). These differences can be large: we found that China's territory, already the third largest land territory of any country in the world, was shown to be about 21% larger by pixel count when it was depicted on Google Maps localized for mainland Chinese consumption. (That is, ISO country code "cn" vs. other locations. Most of this area difference is due to the ten-dash line.)

Border localization is done without the user's knowledge, and the processes that determine which borders are shown remain opaque. If the major mapping platforms all agree on a particular localization, we expect there would be no way for a user of those platforms to easily detect that they are being shown a different version of the world from other people, as the map interfaces we examined do not easily allow localization to be set by the user as an option (see §3.1 for details). In some cases, savvy comparative Web-browsing could bring a difference to light, but in other cases this would also be impossible: for instance there is presently no `maps.google.ac` Web site available to compare to the output of `maps.google.com`.

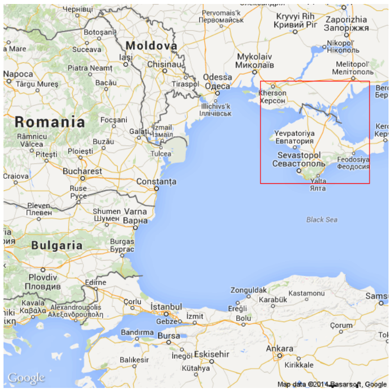

(a) Russia

(b) Ukraine

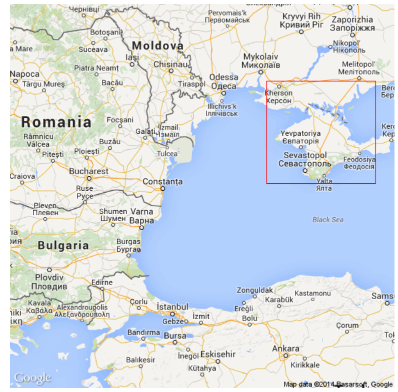

(c) All countries except Russia and Ukraine

Figure 5: Border changes separating the Crimean Peninsula from Ukraine on Google Maps. Map Data ©2014 AutoNavi Google.

Figure 6: Vertical borders in Antarctica along the 25° and 74° West meridians shown only to Argentinian users on Bing Maps. Microsoft product screen shot reprinted with permission from Microsoft Corporation.

(a) China

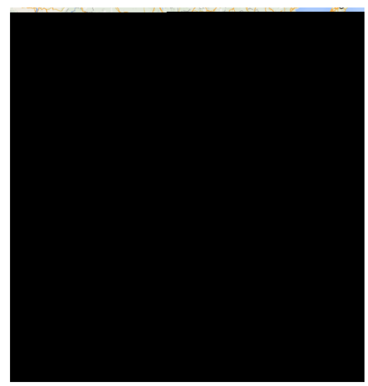

(b) All countries except China

Figure 7: Full view of the personalized maritime border around the South China Sea on Google Maps. Map Data ©2016 Google, ZENRIN

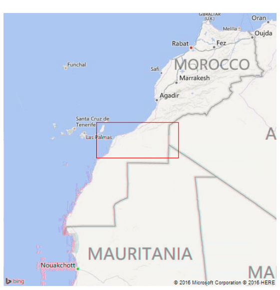

(a) Morocco

(b) All countries except Morocco

Figure 8: View of Morocco and Western Sahara from Bing Maps. Microsoft product screen shots reprinted with permission from Microsoft Corporation.

(a) Ascension Island, Russia, and Ukraine

(b) All countries except Ascension Island, Russia, and Ukraine

Figure 9: An update discovered by crowdworkers, showing a border dispute between Georgia and Russia. Map Data ©2015 Basarsoft, Google

This paper presented MapWatch, a platform that allows these localization differences to be systematically detected. In the remainder of this section we will highlight and discuss MapWatch's intervention into this situation. We will then conclude by discussing the larger implications of personalization suggested by this study.

Dynamic Borders. The temporal monitoring efforts of MapWatch reveal a new orientation toward the technology of the map itself. While maps have always been used as propaganda, never before has it been possible to reconfigure the world so quickly, and to very carefully target particular cartographic messages to specific populations. The Google Maps updates for South Ossetia (2015) and Crimea (2014) indicate that Google Maps is updating national borders 1-2 months after significant world events. In contrast, some print publishers updated the Crimean border 19-21 months after a relevant treaty [45]. While some updates are surely driven by legal compliance, there also appears to be a band of discretion where map providers may pursue their own motives.

Archiving and Preservation. Unlike printed maps which are saved in physical libraries, changes to a map of the world provided as an interactive online service may be ephemeral, and could easily occur without notice. Internet Archive founder Brewster Kahle once said, "The average lifespan of a Web page is 100 days. This is no way to run a culture." In this spirit, a key contribution of MapWatch is archival. Just as researchers have used over-time animations of the Internet Archive to visualize the Web's evolution [44] MapWatch can provide the future with a dynamic history of the Web's cartographic past.

Algorithmic Auditing. Although the expertise required to build MapWatch is not trivial, we found that the ongoing operating cost is quite low. This is important as it serves as a proof-of-concept for the problem of black-box algorithmic auditing (see §2.3). Our ability to realize this project suggests that meaningful auditing may be possible for other important Web-based platforms. This is significant, as some commentators have characterized the personalized Web as a new domain of inscrutable opacity that we are helpless to examine [61].

Borders vs. Other Features. This specific auditing project provides the first holistic picture of how often national borders are personalized on major Web mapping platforms. However, public statements by Google indicated that the company operated at least 32 versions of the world map in 2010. Our project's "baseline check" (§3) was therefore an estimate of what proportion of these differences involve borders (detected by MapWatch) vs. other features. These other features probably include labels (e.g., place names) and non-border map features too small to be seen at the present zoom level (e.g., sensitive government installations). For instance, during manual analysis we noticed that "Arabian Gulf" is displayed for thirteen localizations (*sa*, *eg*, *iq*, *jo*, *kw*, *bh*, *qa*, *om*, *ae*, *lb*, *ly*, *dz*, and *tn*) while the rest of the world sees the label "Persian Gulf."

With continued development, we may be able to expand MapWatch to detect differences in other features beyond borders—such as places removed from maps by government request. Other scholars have highlighted the study of cartographic censorship as a key research need, and warned that online maps may be eroding democratic norms of access to public information via an incremental series of "small encroachments that make little sense" [43] such as blurred buildings and hidden map features.

Compliance vs. Discretion. Five of the seven personalized borders we observed with MapWatch so far concern economically powerful countries like India, China, and Russia. All three have large Internet-connected populations, and are major markets for software and digital services. This raises questions about the leverage these countries have to force map providers to implement personalized borders. For example, Google has physical offices in Russia and India, in addition to significant revenue, that could be threatened if they refused to acquiesce to governmental demands.

That Google Maps implements extensive personalization around China is particularly interesting. Google famously withdrew their services from China in 2010 rather than comply with state-sponsored censorship. Google's services have effectively been banned in China since then [33]. Yet, in the mapping case, Google has clearly produced a map consistent with the wishes of the Chinese government. Whether this is for reasons of legal compliance, a response to a formal request from China, or a desire on Google's part to reach détente with the Chinese government is not known. Microsoft has not personalized the South China Sea on Bing Maps, even though they have extensive business operations in China.

Map providers are sometimes beholden to requests from governments, but even when they are not compelled to edit a border, their exercise of discretion also has important implications. Flickr required users to disambiguate "Israel" from "Occupied Palestinian Territories"—taken to its extreme conclusion, this kind of selection suggests cartography as a preference. It indicates the future possibility of an algorithmic cartography that is designed to show the user the shape of the world that they are most comfortable with. With continued development, future iterations of MapWatch can be designed to detect this future.

Process and Transparency. Our speculation about motives and processes is exacerbated by the lack of transparency surrounding these services. Specifically, we have no idea which countries are requesting changes to maps, or the process by which companies handle these requests. MapWatch improves this situation by revealing map changes shortly after they are implemented, which can help us to deduce the *who* and possibly even the *why* behind them.

Limitations. MapWatch has two noteworthy limitations. First, although MapWatch uses several techniques to minimize labeling errors (see § 3.4), it is possible that the crowdsourced workers may miss some border updates (i.e., false negatives). Fortunately, we observe that map tiles on Google and Bing are updated relatively frequently, which gives the workers multiple opportunities to identify changes. In the future, we may be able to further improve MapWatch's accuracy using more sophisticated computer vision and machine learning techniques (e.g., deep convolutional neural networks).

Second, our existing implementation only crawls tiles at zoom level 6. This level is coarse enough that we may miss personalized or altered borders that are less then 50km in length. Scaling MapWatch to monitor tiles at higher zoom levels would enable us to monitor additional interesting cartographic features, e.g., shifting intra-state borders in Israel [11], and censored military installations [10]. Currently, we have no records of these significant cultural artifacts. Although crawling tiles at higher zoom levels incurs a commensurate increase in labeling costs (since the number of tiles grows as you zoom in), incorporating automated image classifiers into MapWatch may mitigate this issue.

Future Work. We have already noted several areas of improvement for MapWatch, including automated image labeling, increasing the zoom level, adding more map providers (e.g., OpenStreetMaps), and adding support for detecting textual changes (e.g., "Arabian Gulf"). Additionally, the data collected by MapWatch raises interesting user-interface, visualization, and human-computer interaction questions. As noted by historians like Jo Guldi, there are many maps of geographical space, but far fewer of geographical time [26]. Long-term, we would like to expose historical MapWatch data to users, which will necessitate the development of map visualizations that allow the user to zoom backwards through time, as well as highlighting areas with personalized tiles. Beyond the archival components of this work, this data set is ripe for exploring visualization challenges such as the visualization of uncertainty — an unsolved problem in this domain [2, 14, 52], the historical representation of temporal data [56], and the dynamic visualization of invisible traces that exist in everyday experience [58].

Open Source. We make the MapWatch source code available to the research community at

`http://personalization.ccs.neu.edu/`

Additionally, we offer access to MapWatch to academic researchers upon request. Unfortunately, we cannot open source our map tile database due to image copyright issues.

6. CONCLUSION

This study documented a series of specific investigations into map localization and national borders. In the widest framing, these results raise the question of how to conceptualize personalized systems in general. Mapping has typically been a quasi-public endeavor, with governments playing an essential role in both producing maps and controlling them. Whether this role is funding and operating an extensive geographic survey unit, legislating that public map data must be available for free, or launching observation satellites, governments have invested the significant resources required to produce a comprehensive map of anything. More recently, the rise of digital cartography produced an important new global role for private corporations that provided digital map data and map manipulation tools, such as Navteq, TeleAtlas, and ESRI. In this study, we observed that major Internet platforms such as Microsoft, Apple, and Google have become important new intermediaries in the mapping ecology, and they appear to be making geopolitical decisions that would not have been seen as relevant to their businesses just a few years earlier.

MapWatch demonstrated that different people are shown a different world because of localization by Internet platforms. Localization is a form of personalization, and personalization is often portrayed as the task of satisfying user needs and preferences. As we have seen in this study, national borders on maps are an incisive case study of personalization because they clearly demonstrate the problem with this formulation of personalization. On maps, needs and preferences conflict, and the interests of the user, the map provider, the government, and society in general need not coincide. It is not clear that the world is best served by showing map readers the borders that they prefer. At the same time, nationalism produces strong feelings about borders and online map users are likely to hold strong preferences (such as those Google users who have geo-annotated the Spratly Islands).

This situation foregrounds the role of the system designer as online intermediary. It may be that online platforms that provide personalization should have an obligation to reveal uncertainty or conflict in their representations of the world (as Google Earth does, but Google Maps does not). We may need 1) more transparency in how personalization decisions are made (or the fact that they are made at all) and 2) some ability for every user to control personalization, as minimum standards for reconciling the demands on interactive systems from competing interests. To achieve these goals we encourage the investigation of other personalized online platforms using systems similar to MapWatch. The extent and dimensions of personalization online should be aired publicly and discussed with the gravity that an issue like "What is the size of China?" deserves.

Acknowledgements

This research was supported in part by NSF grants CNS-1319019 and CHS-1408345. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF.

7. REFERENCES

[1] Argentina in Antarctica. `http://www.dna.gov.ar/INGLES/DIVULGAC/ARGANT.HTM`.

[2] The New York Times Hurricane Tracker, August 2011. `http://www.nytimes.com/projects/hurricanes/?nl=todaysheadlines&emc=thab1#!/2011/Irene`.

[3] Disputed Territories. Knight-Mozilla-MIT "The Open Internet" Hack Day, June 2014. `http://opennews.kzhu.io/map-disputes/`.

[4] Kashmir profile, November 2014. `http://www.bbc.com/news/world-south-asia-11693674`.

[5] Russia signs 'annexation' deal with South Ossetia, March 2015.

[6] Slippy map tilenames. OpenStreetMap Wiki, December 2015. `http://wiki.openstreetmap.org/wiki/Slippy_map_tilenames#Resolution_and_Scale`.

[7] Admin 0 - Breakaway, Disputed Areas. Natural Earth, 2016. `http://www.naturalearthdata.com/downloads/10m-cultural-vectors/10m-admin-0-breakaway-disputed-areas/`.

[8] Amazon Mechanical Turk. `http://mturk.com/`.

[9] Anderson, B. R. *Imagined Communities: Reflections on the Origins and Spread of Nationalism*. Verso, 1983.

[10] Begley, J. How do you measure a military footprint?, 2013. `http://empire.is/`.

[11] Ben-David, A. La palestine et ses frontières virtuelles 2.0: Du "non-lieu" à l'espace généré par les utilisateurs. *Réseaux 28*, 159 (2010), 151–179.

[12] Chen, L., Mislove, A., and Wilson, C. Peeking beneath the hood of uber. In *Proc. of IMC* (2015).

[13] Chitty, A. Western Sahara - territorial dispute, self-determination and the UN, May 2015. http://www.exploringgeopolitics.org/publication_chitty_alex_western_sahara_territorial_dispute_self_determination_un_polisario_sahrawi_plebiscite_minurso_morocco_rio_de_oro_terrritory_algeria_mauritania/.

[14] Correll, M., and Gleicher, M. Error bars considered harmful: Exploring alternate encodings for mean and error. *Visualization and Computer Graphics, IEEE Transactions on 20*, 12 (2014), 2142–2151.

[15] Territorial integrity of ukraine, 2014. http://www.un.org/en/ga/search/view_doc.asp?symbol=A/RES/68/262.

[16] Diakopoulos, N. Algorithmic accountability. *Digital Journalism 3*, 3 (2015), 398–415. doi:10.1080/21670811.2014.976411.

[17] Edelman, B. G., and Luca, M. Digital Discrimination: The Case of Airbnb.com. Available at SSRN, 2014.

[18] Eslami, M., Aleyasen, A., Karahalios, K., Hamilton, K., and Sandvig, C. Feedvis: A path for exploring news feed curation algorithms. In *Proc. of CSCW* (2015).

[19] Eslami, M., Rickman, A., Vaccaro, K., Aleyasen, A., Vuong, A., Karahalios, K., Hamilton, K., and Sandvig, C. "i always assumed that i wasn't really that close to [her]": Reasoning about invisible algorithms in the news feed. In *Proc. of CHI* (2015).

[20] Geens, S. Google Maps' Arunachal Pradesh Place Names Turn Chinese, Google Admits Error. Ogle Earth, September 2009. http://ogleearth.com/2009/08/google-maps-arunachal-pradesh-place-names-turn-chinese-google-admits-error/.

[21] Google. Improving the quality of borders in Google Earth and Maps. Google Lat Lon: The Google Maps Blog, 2010. http://google-latlong.blogspot.com/2010/07/improving-quality-of-borders-in-google.html.

[22] Google Maps API. https://developers.google.com/maps/.

[23] Graham, S. Software-sorted geographies. *Progress in Human Geography 29*, 5 (2005), 562–580. doi:10.1191/0309132505ph568oa.

[24] Gravois, J. The Agnostic Cartographer. Washington Monthly, July 2010. http://www.washingtonmonthly.com/features/2010/1007.gravois.html.

[25] Guha, S., Cheng, B., and Francis, P. Challenges in measuring online advertising systems. In *Proc. of IMC* (2010).

[26] Guldi, J. *Roads to Power: Britain Invents the Infrastructure State*. Harvard University Press, 2011.

[27] Hamilton, K., Karahalios, K., Sandvig, C., and Langbort, C. The image of the algorithmic city: A research approach. *Interaction Design and Architecture(s) (IxDA)*, 20 (2014).

[28] Hannak, A., Sapiezynski, P., Kakhki, A. M., Krishnamurthy, B., Lazer, D., Mislove, A., and Wilson, C. Measuring personalization of web search. In *Proc. of WWW* (2013).

[29] Hannak, A., Soeller, G., Lazer, D., Mislove, A., and Wilson, C. Measuring price discrimination and steering on e-commerce web sites. In *Proc. of IMC* (2014).

[30] Harley, J. B. Maps, knowledge, and power. *Geographic thought: a praxis perspective* (1988), 129–148.

[31] Harley, J. B. Deconstructing the map. *Cartographica 26*, 2 (1989), 1–20.

[32] Heffernan, M. The politics of the map in the early twentieth century. *Cartography and Geographic Information Science 29*, 3 (2002), 207–226.

[33] Helft, M., and Wines, M. Google Faces Fallout as China Reacts to Site Shift. The New York Times, 2010. http://www.nytimes.com/2010/03/24/technology/24google.html.

[34] King, C. The Five Day War, November 2008. https://www.foreignaffairs.com/articles/russia-fsu/2008-11-01/five-day-war.

[35] Kliman-Silver, C., Hannak, A., Lazer, D., Wilson, C., and Mislove, A. Location, location, location: The impact of geolocation on web search personalization. In *Proc. of IMC* (2015).

[36] Lecuyer, M., Ducoffe, G., Lan, F., Papancea, A., Petsios, T., Spahn, R., Chaintreau, A., and Geambasu, R. Xray: Enhancing the web's transparency with differential correlation. In *Proc. of USENIX Security Symposium* (2014).

[37] Lecuyer, M., Spahn, R., Spiliopolous, Y., Chaintreau, A., Geambasu, R., and Hsu, D. Sunlight: Fine-grained targeting detection at scale with statistical confidence. In *Proc. of CCS* (2015).

[38] Microsoft. Microsoft's Geopolitical Bug Experts. Microsoft News Center, 2010. https://news.microsoft.com/2000/12/06/microsofts-geopolitical-bug-experts/.

[39] Mikians, J., Gyarmati, L., Erramilli, V., and Laoutaris, N. Detecting price and search discrimination on the internet. In *Proc. of HotNets* (2012).

[40] Mikians, J., Gyarmati, L., Erramilli, V., and Laoutaris, N. Crowd-assisted search for price discrimination in e-commerce: First results. In *Proc. of CONEXT* (2013).

[41] Monmonier, M. *How to Lie With Maps*. University of Chicago Press, 1991.

[42] Monmonier, M. *Drawing the Line: Tales of Maps and Cartocontroversy*. H. Holt, 1995.

[43] Monmonier, M. *No Dig, No Fly, No Go: How Maps Restrict and Control*. University of Chicago Press, 2010.

[44] Rogers, R. *Digital Methods*. MIT Press, 2013.

[45] RT. Oxford's new geography textbook names Crimea as part of Russia, October 2015. http://on.rt.com/6tx6.

[46] Ruggeri, A. The Politics of Making Maps. BBC, June 2014. http://www.bbc.com/travel/story/20140602-the-politics-of-making-maps.

[47] Saltman, J. Implementing local housing laws through social action. *Journal of Applied Behavioral Science 11*, 1 (1975), 39–61.

[48] Sandvig, C., Hamilton, K., Karahalios, K., and Langbort, C. Auditing Algorithms: Research Methods for Detecting Discrimination on Internet Platforms. ICA Pre-Conference on Data and Discrimination, 2014.

[49] Schöning, J., Hecht, B., and Kuhn, W. Informing online and mobile map design with the collective wisdom of cartographers. In *Proc. of ACM DIS* (2014).

[50] Seaver, N. Knowing Algorithms. Media in Transition (MiT) 8, 2013. http://static1.squarespace.com/static/55eb004ee4b0518639d59d9b/t/55ece1bfe4b030b2e8302e1e/1441587647177/seaverMiT8.pdf.

[51] Sparkes, M. Revealed: how Google moves international borders. The Telegraph, June 2014. http://www.telegraph.co.uk/technology/google/10922595/Revealed-how-Google-moves-international-borders.html.

[52] Spiegelhalter, D., Pearson, M., and Short, I. Visualizing uncertainty about the future. *Science 333*, 6048 (2011), 1393–1400.

[53] Sweeney, L. Discrimination in online ad delivery. *ACM Queue 11*, 3 (2013).

[54] Tharoor, I. Could This Map Start a War? The Washington Post, 2014. https://www.washingtonpost.com/news/worldviews/wp/2014/06/27/could-this-map-of-china-start-a-war/.

[55] Thielmann, T. "you have reached your destination!" position, positioning and superpositioning of space through car navigation systems. *Social Geography*, 2 (2007), 63–75. doi:10.5194/sg-2-63-2007.

[56] Tufte, E. R., and Graves-Morris, P. *The visual display of quantitative information*, vol. 2. Graphics press Cheshire, CT, 1983. Visualization of Minard's Napoleon's March to Moscow.

[57] Vaccaro, K., Karahalios, K., Sandvig, C., Hamilton, K., and Langbort, C. Agree or cancel? research and terms of service compliance. In *ACM CSCW Ethics Workshop: Ethics for Studying Sociotechnical Systems in a Big Data World* (2015).

[58] Viegas, F., and Wattenberg, M. Windmap, 2012. http://hint.fm/projects/wind/.

[59] Wood, D. M., and Graham, S. Permeable boundaries in the software-sorted society: Surveillance and the differentiation of mobility. In *Mobile Technologies of the City*, M. Sheller and J. Urry, Eds. Routledge, 2006, pp. 177–191.

[60] York, J. C. The myth of a borderless internet. *Atlantic Monthly* (June 2015). http://www.theatlantic.com/technology/archive/2015/06/the-myth-of-a-borderless-internet/394670/.

[61] Ziewitz, M. Governing algorithms: Myth, mess, and methods. *Science, Technology, and Human Values 41*, 1 (2016). doi:10.1177/0162243915608948.