GI_Forum 2022, Issue 1
Page: 91 - 97
Short Paper
Corresponding Author:
arobinson@psu.edu
DOI: 10.1553/giscience2022_01_s91

Visualizing Viral Cartography with MapReverse

Anthony C. Robinson[1] and Xi Zhu[1]
[1]The Pennsylvania State University, USA

Abstract

Social media provide a fast-moving ecosystem for the rapid dissemination of digital maps. Some maps become viral, reaching massive audiences across diverse networks. A new challenge for cartographers is to characterize the design and dissemination of viral maps in social media. Here, we present a prototype visualization tool called MapReverse which allows users to explore collections of map images that are similar to a given viral map. Similar maps are identified using a combination of reverse image search and machine learning image analysis. In some cases, we can identify original sources for maps; more commonly we can identify many modifications of the original viral maps. This system provides a glimpse into the evolution of cartographic design as found in viral social media.

Keywords:

virality, cartography, design, image analysis, disinformation

1 Introduction

Maps are powerful abstractions of our world, and people are often inclined to trust the information they see on a map. It is possible that public trust in maps is amplified by the fact that so many people now have access to reliable and effective personal navigation apps and devices. We hypothesize that many map users may not distinguish between maps used for navigation and those that are used to tell stories about thematic data (e.g. election results or climate change). Users may place a great deal of trust in thematic maps, unaware that they embed a broader range of uncertainties related to their provenance and design than a map built for navigation. As a result, thematic maps can be a particularly useful vehicle for those seeking to muddy the waters when it comes to political, social, economic and environmental issues, especially since manipulating an image to make subtle but crucial changes to its content (e.g. changing elements of a map legend) is now trivial for non-experts to achieve. Furthermore, disseminating maps has never been easier, and we have observed quite a few viral maps that did not originate from trustworthy sources but nonetheless reached very large audiences via social media (Robinson, 2019; Griffin, 2020; Shannon & Walker, 2020).

In our present work, we have sought to capture, characterize and visualize examples of manipulated maps that appear in social media. Using examples of viral maps collected from Twitter, we apply multiple commercial machine learning platforms for image analysis to search for similar images, to identify websites on which those maps appear, and to reveal which terms and topics co-occur in those maps. We have developed a prototype visualization tool called MapReverse (http://mapreverse.org/) that supports visual exploration of viral maps and their variations. In several instances, we have found that viral maps are in fact manipulated images based on earlier work. In other examples, we can see that a viral map may be original work, but that it can then become the canvas for derivations downstream.

In both scenarios, maps can be wielded as tools of disinformation to destabilize social conversations, and they could potentially be used to disrupt crisis situations. For example, we observed election-related maps that had been manipulated to intervene in the U.S. political discourse (McNutt et al., 2021). We have also found examples of viral maps that relate to crises (Cha et al., 2020). It would be possible, for example, to create and widely disseminate maps that provide misleading information about the spread of a wildfire, the best evacuation routes, or the spread of an infectious disease. Our work seeks to develop methods that will support the rapid capture of maps in social media, analyse their contents, and visualize the results so that patterns of disinformation become clearer and the provenance of a given map may be more readily understood by map readers.

In the following sections, we review recent related work and describe the development and functionality of the MapReverse prototype. We then highlight an example of what discoveries can be made using MapReverse. Finally, we conclude with ideas for future research challenges related to geovisual analytics of viral cartography.

2 Related Work

The persuasive power of maps is not new, but today it is easier than ever to make and share maps with broad audiences in a short period of time (Muehlenhaus, 2014). The rapid dissemination and diversification of cartographic production are generating new scenarios in which thematic maps are quickly crafted and contested (or not), impacting global audiences. Viral content in social media may now feature maps as the centre for discussion, and there are instances in which one viral map leads to the emergence of many new and modified maps in its wake (Robinson, 2019).

Advances in machine learning designed to analyse the attributes of image collections have begun to be applied in cartographic research. Li and Xiao (2019) compared the application of several machine learning methods to identify potential maps from image collections and to automate the identification of the regions they depict as well as their projections. Others have used machine learning image analysis techniques to extract embedded information from large map archives of scanned historic paper maps (Uhl et al., 2018). Tavakkol et al. (2019) have taken a combined approach, using crowdsourcing along with machine learning to vectorize, geo-rectify and describe scans of historical maps.

This research also builds on recent efforts to characterize issues surrounding trust and disinformation in visualization and cartographic design. The power of graphics to shape public opinion and discourse is well known, but the mechanics of this power are not well understood. MapReverse offers a partial response to calls in information visualization and cartography to interrogate the intersections of visualization design, trust and disinformation (Griffin, 2020; McNutt et al., 2021). Recent work in this area by Ledermann (2018) provides one example of how the digital provenance of a web map can be evaluated in the context of understanding cartographic veracity.

3 MapReverse: Visual Analytics for Viral Cartography

3.1 Reverse Image SearchFigure 1: Eye-Tracking Metrics for Task 5 (Adding a new POI)

Our work combines search results from the TinEye reverse image search engine. TinEye can ingest an input image and compare it to its collection of over 50 billion images. In addition, TinEye allows the retrieval of a match score for similar images (albeit one that is calculated by a proprietary algorithm that we cannot review), enabling the user to identify and compare potential antecedent or descendent maps of a particular viral map.

3.2 Image Analysis

We augment results from TinEye by taking its top matches and analysing those images with Google Cloud Vision (GCV). GCV is a service which analyses the structure and contents of a given input image. It can highlight elements of the visual construction of an image (e.g. graphics, colours, text) as well as identify relevant web entities (e.g. the topics discussed on web pages that embed that image). Additionally, GCV provides details on the domains in which an image is found.

3.3 MapReverse Prototype

To visualize the results of the image analysis carried out using TinEye and GCV, we developed the MapReverse prototype (http://MapReverse.org/). MapReverse uses a combination of React.js and D3.js to support interactive cross-filtering and highlighting across multiple coordinated views of viral map data. The interface allows users to view a matrix showing the results of a reverse image search for a given viral map, ordered by similarity scores from highest to lowest (Figure 1-A). MapReverse provides controls to filter those results by similarity scores and/or the timeframe in which those images were first indexed (Figure 1-B). Image details, including basic metadata, are provided along with the image at its original size, as shown in Figure 1-C. Finally, it is possible to view and filter by web entities and domains (Figure 1-D).

Figure 1: MapReverse supports the coordinated visualization of multiple views and the analysis of viral maps and related maps found via reverse image search and annotated via machine learning.

4 Case Study Example

In the Summer of 2019, Hurricane Dorian threatened the east coast of the United States. As the storm intensified and its forecasted track was updated, the U.S. National Hurricane Center (NHC) issued a series of maps showing the so-called 'cone of uncertainty' associated with Dorian's movement (Ruginski et al., 2016). As the hurricane continued its movement and intensification, these maps were iteratively updated to show changes to areas under threat. It is common for the representation of the cone of uncertainty to change many times as a storm develops.

One map showing Hurricane Dorian's progress became a centre of controversy following statements by the president of the United States referring to Alabama as a state under potential threat from Dorian. This reflected information from an older, no longer accurate forecasted cone of uncertainty. The weather service office in Alabama then issued a statement assuring residents that they should not be impacted by Dorian. The president insisted he was correct, refusing to acknowledge that the forecast had changed over time, and setting off a controversy that intensified over several days. At one point, the president displayed an obviously doctored print of an older cone of uncertainty map that had been crudely extended to include part of Alabama (Watson, 2019). The episode became known as *Sharpiegate*, named after the marker pen that had clearly been used by the president to alter the map.

Here, we show how MapReverse can be used to examine aspects of this story. Although the manipulation of this map was done on a physical copy of a digital map, as an example it helps reveal the potential strengths as well as weaknesses associated with the application of reverse image search and machine learning image analysis on social media maps. We began by using the original un-modified NHC map of the forecasted track for Hurricane Dorian; we then retrieved similar images using that seed (Figure 1-A). The MapReverse image matrix is ordered by similarity score from top to bottom, and many of the most algorithmically similar results are simply different crops of the original image, sometimes using different resolutions. Next in order, we see a large number of NHC maps from other storms, which is not surprising given the intentional visual similarity associated with NHC warning map design (Figure 2-A).

Figure 2: An example showing cross-filtered results in MapReverse to help identify temporally and semantically relevant map images stemming from the 2019 *Sharpiegate* controversy.

Filtering MapReverse by image similarity score helps narrow the pool further, but we do not see examples here of the sharpie-modified map (Figure 2-B). If we switch and use time as a filtering facet, we can select September 2019 and then click the entity marker titled 'Hurricane Dorian-Alabama Controversy' to highlight temporally and semantically relevant image results (Figure 2-C). The sharpie-modified map does eventually appear, but it is ranked very low in terms of image similarity score, probably because the modified map was done using a physical copy which was then photographed by the media, making it a poorer match in an algorithmic sense to the original NHC digital map. This reveals an important caveat associated with at least some of the available commercial image analysis tools: these tools may be easily fooled by simply taking a photograph of a digital image before re-sharing it in order to avoid potential filters which might alert to a modification.

5 Research Challenges

Developing a geovisual analytics platform for evaluating the sources and modifications associated with viral maps raises several important future research questions. For example, we do not yet know a great deal regarding the dissemination networks that drive the spread of viral maps. We hope to address this in future iterations of the MapReverse system. It also remains difficult to identify the original sources for certain images, and in some cases this is exacerbated by the fact that web images are often hosted by third parties (e.g. not directly via the website in which a map is viewed). A great deal of web content is ephemeral, with links dying as time progresses. Web archives must in the future be integrated with systems developed to analyse the origins and trajectories of social media maps.

Additionally, we note that it is a major challenge to algorithmically identify novel maps in social media sources. The simple act of photographing and then re-posting a map with a nearly identical design may cause its measured similarity to decline, as shown in our *Sharpiegate* case study. Adversarial examples are commonly utilized in the development of machine learning methods (Moosavi-Dezfooli et al., 2016), and we need new adversarial examples that reflect cartographic contexts.

Finally, we do not know how users interpret information from maps they find in social media. Maps clearly elicit a level of trust in many readers, but we do not know whether this trust is changing; nor can we describe in detail what 'trust' in mapping means. User evaluations are needed to begin tackling this major area of concern. Previous work by Hoyer (2016) provides a potential set of geomedia literacy goals that could be used to frame such research.

References

Cha, M., Gao, W., & Li, C.-T. (2020). Detecting Fake News in Social Media: An Asia-Pacific Perspective. *Communications of the ACM*, *63*(4), 68–71. https://doi.org/10.1145/3378422

Griffin, A. L. (2020). Trustworthy maps. *Journal of Spatial Information Science*, *20*, 5–19. https://doi.org/10.5311/JOSIS.2020.20.654

Hoyer, T. (2016). On the question of how Web 2.0 features support critical map reading. *GI_Forum 2016, Volume 4,* 295–301. https://doi.org/10.1553/giscience2016_01_s295

Ledermann, F. (2018). Analysing Digital Maps Online: A Reverse Engineering Approach. *GI_Forum 2018*, *Volume 6*(2), 119–130. http://dx.doi.org/10.1553/giscience2018_02_s119

Li, J. & Xiao, N. (2019). *Using Machine Learning Methods to Identify and Classify the Regions and Projections of Online Maps. GeoComputation 2019*. https://doi.org/10.17608/k6.auckland.9848360.v2

McNutt, A. M., Huang, L., & Koenig, K. (2021). Visualization for Villainy. *ArXiv:2109.06007 [Cs]*. http://arxiv.org/abs/2109.06007

Moosavi-Dezfooli, S.-M., Fawzi, A., & Frossard, P. (2016). *DeepFool: A Simple and Accurate Method to Fool Deep Neural Networks*. 2574–2582. https://doi.org/10.1109/CVPR.2016.282

Muehlenhaus, I. (2014). Going Viral: The Look of Online Persuasive Maps. *Cartographica: The International Journal for Geographic Information and Geovisualization*, *49*(1), 18–34. https://doi.org/10.3138/carto.49.1.1830

Robinson, A. C. (2019). Elements of viral cartography. *Cartography and Geographic Information Science*, *46*(4), 293–310. https://doi.org/10.1080/15230406.2018.1484304

Ruginski, I. T., Boone, A. P., Padilla, L. M., Liu, L., Heydari, N., Kramer, H. S., Hegarty, M., Thompson, W. B., House, D. H., & Creem-Regehr, S. H. (2016). Non-expert interpretations of hurricane forecast uncertainty visualizations. *Spatial Cognition & Computation*, *16*(2), 154–172. https://doi.org/10.1080/13875868.2015.1137577

Shannon, J., & Walker, K. E. (2020). Ventures into Viral Cartography: Waffle House, Educational Attainment, and the Social Life of Maps. *The Professional Geographer*, *72*(1), 66–77. https://doi.org/10.1080/00330124.2019.1653774

Tavakkol, S., Chiang, Y.-Y., Waters, T., Han, F., Prasad, K., & Kiveris, R. (2019). Kartta Labs: Unrendering Historical Maps. *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on AI for Geographic Knowledge Discovery*, 48–51. https://doi.org/10.1145/3356471.3365236

Uhl, J. H., Leyk, S., Chiang, Y.-Y., Duan, W., & Knoblock, C. A. (2018). Map Archive Mining: Visual-Analytical Approaches to Explore Large Historical Map Collections. *ISPRS International Journal of Geo-Information*, *7*(4). https://doi.org/10.3390/ijgi7040148

Watson, J. (2019). *US President Donald Trump and Acting US Secretary of Homeland Security Kevin McAleenan update the media on Hurricane Dorian preparedness as they show a map, from August 29, 2019, in the Oval Office at the White House in Washington, DC, September 4, 2019.* [Photography]. AFP/Getty Images. https://www.gettyimages.in/detail/news-photo/president-donald-trump-and-acting-us-secretary-of-homeland-news-photo/1165958521