Analyzing the Geography of Systemic Racism in an Introductory Geographic Information Systems Workshop

Aditya Ranganath

Center for Research Data and Digital Scholarship, University of Colorado (Boulder), Boulder, United States

aditya.ranganath@colorado.edu

Analyzing the Geography of Systemic Racism in an Introductory Geographic Information Systems Workshop[1]

Academic geospatial librarians have the potential to stimulate broader critical understanding and reflection about the racial inequities and injustices that remain inscribed in our social institutions. One way they might do so is by teaching introductory GIS workshops that explore these themes. This paper proposes one such workshop, and provides a link to a detailed sample lesson plan that other instructors can use as a template for their own teaching materials. In particular, the proposed workshop uses a publicly available dataset of traffic police stops, which has been collected and organized by the Stanford Open Policing Project, to explore and document geographic patterns in racially biased policing practices.

Keywords: geographic information systems; spatial data science; systemic racism; policing; census; bias

[1] The sample workshop discussed in this paper can be accessed at the following link: https://doi.org/10.25810/x6yz-6g18

Introduction

Academic geospatial librarians play an important role as educators on their campuses. In fulfilling this role, they are often asked to teach introductory workshops on Geographic Information Systems (GIS) and spatial data to members of the campus community. This presents an opportunity for geospatial librarians to incorporate the empirical analysis of systemic racism and discrimination into their geospatial educational offerings, and thereby stimulate broader critical understanding and reflection about the racial inequities and injustices that remain inscribed in our social institutions.

Indeed, there is growing interest, both within academia and in activist communities dedicated to social justice issues, in using data analysis and visualization techniques to document and understand patterns of systemic racism and inequality in the United States, with the goal of eventually empowering citizens (particularly those in the marginalized communities that are most directly harmed by these injustices) to effectively mobilize for meaningful reform.[2] Spatial analysis has an important role to play in such efforts, and geospatial librarians in turn have the potential to play an important role in introducing scholarly and activist communities on campus to the ways in which GIS and spatial visualization can be deployed to further our understanding of systemic racism. In doing so, they would contribute to an important civic mission, and meaningfully advance the commitment of academic libraries to promote the common good.

However, it is challenging to address substantive themes of systemic racism and social (in)justice within the temporal constraints imposed by a typical 60 to 90 minute introductory GIS workshop. Indeed, while members of the geospatial librarian community

[2] The organization Data for Black Lives is one example of an anti-racist community organization that centers the role of data science and visualization in its work. There are often important connections between such organizations and scholarly research communities. For a discussion of several projects (often involving collaborations between academics and community members) that use data science to understand and redress various forms of social inequity, an excellent place to begin is *Data Feminism,* by D'Ignazio and Klein (2020).

have pioneered compelling projects that document the geography of systemic racism in the United States, such as the University of Minnesota's "Mapping Prejudice" project (Ehrman-Solberg et al. 2020), the scale and complexity of these projects may render them unsuitable for the purpose of teaching basic GIS and data analysis skills in an introductory setting.

This raises an important question for the community of geospatial librarians: how might an instructor teach essential skills of geospatial analysis and visualization, while simultaneously encouraging learners to reflect on substantive themes of systemic racism and inequality, in an introductory GIS workshop that engages participants without overwhelming them? Developing workshop curricula with these priorities in mind is a long-term project for librarians and the broader geospatial education community, and one that will require discussion and collaboration over time.

To that end, the purpose of this paper is to introduce one such workshop lesson plan that integrates geospatial education with an exploration of systemic racism. In particular, the proposed workshop lesson plan makes use of data that has been made publicly available by the Open Policing Project at Stanford University[3], an online archive of traffic police stop data collected and organized by Pierson et al (2020). This workshop guides learners through the process of downloading a state-level selection of this data, and reshaping and processing the data with a view towards developing a simple indicator of county-level racial bias in traffic police stops; it then instructs workshop participants on how to create a map that displays county-level variation in this indicator on a state map. The workshop highlights one of the great virtues of GIS as a methodological tool, namely, its ability to help researchers explore social phenomena at granular scales. It also invites learners to "think spatially" by reflecting on the origins and implications of the spatial patterns that the workshop helps to uncover.

[3] https://openpolicing.stanford.edu/

In what follows, I will briefly introduce the Stanford Open Policing Project data archive, and suggest considerations for geospatial librarians to keep in mind when selecting datasets from this archive for use in potential GIS workshops. I then provide an overview of a sample introductory GIS workshop that uses data from the Open Policing Project. A detailed lesson plan for this sample workshop, which includes relevant R code, code comments, explanations, and code outputs, is available as supplemental material at the following link: https://doi.org/10.25810/x6yz-6g18. This lesson plan can be used by readers as a template to develop their own GIS workshops that use data from the Stanford Open Policing Project.[4] Finally, I provide some general content and delivery-related suggestions for a prospective workshop that uses Open Policing Project data, before concluding.

Background and Data

A prominent manifestation of systemic racism in American society is the effective criminalization of "Driving while Black", a phrase that has entered the lexicon as shorthand for a punitive traffic policing system that subjects Black motorists to greater scrutiny than their white peers. The criminalization of "Driving while Black", is a demeaning and dangerous aspect of everyday life for America's Black population. Moreover, the tendency of the traffic police to stop Black drivers at higher rates than other races for pretextual reasons that are not tied to actual traffic violations has in turn been extensively studied and documented by social scientists using large datasets and sophisticated analytic methods (Harris 2010; Braga, Brundson, and Drakulich 2019).

[4] The lesson plan in the supplemental material appendix uses the R programming language, but does not presuppose a background in R, and is designed to be accessible to geospatial librarians from diverse intellectual backgrounds. Moreover, geospatial librarians who are not R users, but who would like to teach the workshop using a different GIS software platform, will be able to adapt the lesson plan to the GIS platform of their choice after reading through the supplemental materials. In other words, while the code which implements the analysis and visualization tasks in the sample lesson plan is written using R, these tasks are platform agnostic, and can be implemented using a variety of GIS software applications.

Pierson et al (2020) make one of the most important and comprehensive recent contributions to this social scientific literature on the criminalization of "Driving while Black" and the inequitable treatment of Black motorists. The heart of this contribution is the collection of a large dataset (>100 million observations) of traffic stops across the United States, which was compiled by the researchers through public records requests across all 50 states. This data has since been cleaned, standardized, and publicly released (under an Open Data Commons license) by the researchers for use by the public and scholarly community as part of the Stanford Open Policing Project.[5] The workshop introduced in the next section (and presented in detail in the supplemental materials) uses a selection from this broader collection of traffic stop data and demonstrates how it can be analyzed in a spatial context.

The Open Policing Project data is provided for several states, often at multiple geographic scales, and in multiple file formats. In some cases, the Open Policing Project data archive provides explicitly geographic data (i.e. in standard file formats for geospatial vector data, such as shapefiles). However, even in cases where only tabular datasets are provided, the datasets usually contain geographic information related to the traffic stops. Sometimes, this geographic information is explicit, coming in the form of latitude/longitude coordinates that allow users to easily generate a point vector layer using standard GIS techniques. Other times, explicit geographic coordinates are not available, but more general geographic information (such as information about the county in which a given stop occurred) can be used to join tabular datasets from the Open Policing Project to spatial datasets of geographic boundaries, which subsequently allows for a GIS analysis or spatial visualization.

The variety of datasets provided through the Open Policing Project requires prospective workshop instructors to review the options with care, and think critically about

[5] The Open Policing Project website provides a list of various publications that use the data, which may be useful for prospective workshop instructors to explore if they plan to develop a workshop that uses this data: https://openpolicing.stanford.edu/publications/

which dataset will best serve their pedagogical goals and effectively engage the interest of workshop participants. This can be a time-consuming aspect of preparing a workshop that uses this data archive, but the diversity of the Open Policing Project data is one of the reasons why it is such a rich and fertile resource for librarian instructors. Indeed, it offers prospective librarian instructors the opportunity to develop workshops that explore the racial dynamics of traffic policing in local communities with which learners are familiar (i.e. areas that are close to campus). In doing so, instructors can more effectively connect the data and GIS analyses they demonstrate in the workshop to the lived experiences of real people in a concrete and compelling way, and thereby encourage learners to empathize with the human beings that are behind the data points. Moreover, the prospect of adapting the Open Policing Project data to relevant local geographies offers librarian instructors the opportunity to use the workshop to pursue a pedagogy that is informed by what Brunner et al. (2022) call "community engaged digital scholarship", wherein data and information literacy instruction engages with the "layered communities" in which learners are embedded (4).

In addition to selecting an Open Policing dataset that is relevant to a campus's geographic locale, there are also other factors that might inform the choice of a workshop dataset, which is a choice that in turn may affect the workshop's content. For example, several of the Open Policing Project's datasets are tabular datasets that contain latitude/longitude information, which can be used to generate a GIS point layer that can be mapped. However, this latitude/longitude information is rarely comprehensive, and there are often observations for which this spatial information is missing. As a result, instructors will want to ensure that a given dataset's spatial information is sufficiently robust for workshop use. Inevitably, even datasets that are fit for use in a geospatial workshop will contain missing data (after all, missing location data is a common problem in spatially explicit datasets); it is therefore important to be transparent about such shortcomings in the data with

learners, and discuss ways in which such shortcomings might affect the visualization and analysis. Indeed, such discussions about data limitations can be valuable pedagogical opportunities; for instance, participants might discuss whether they think the missing spatial data is randomly distributed or systematically patterned, and the possible implications of either scenario for the conclusions they might reasonably draw from the data. Additional topics relevant to GIS data, such as geocoding algorithms, and the accuracy (or inaccuracy) of geocodes, might also be usefully discussed in this context. Ideally, learners should come away from such a discussion with an appreciation for the fact that any conclusions drawn from the workshop exercise must be considered preliminary and provisional (but it should be emphasized that such caution is warranted in any exploratory data analysis).

It is also important to explicitly consider the relationship between a given Open Policing dataset and the geographic unit of analysis to which the data is aggregated for visualization purposes. In many Open Policing Project datasets, data is provided in the form of a tabular dataset that does not contain explicit spatial information, but which does contain location information that can be used to implement a table join with a spatial dataset, and subsequently carry out a GIS analysis. This is the case for the dataset used in the sample workshop discussed below, in which county-level information on traffic stops is used to generate a county-level map of racial bias in traffic stops. Of course, as experienced GIS users are aware, the choice of aggregation units directly conditions the conclusions that are likely to be drawn from a given spatial analysis or visualization; this is the essence of the modifiable aerial unit problem (MAUP), and it is worthwhile to invite learners to reflect on its implications at some point in a workshop. In some cases, there is no flexibility about the geographic aggregation units that are used, and this discussion will necessarily be somewhat abstract; after all, if the original data is provided at the county level (without more explicit information), one is effectively constrained to develop a county-level analysis and

visualization. However, if one is working with a spatially explicit dataset (either a point layer or a tabular dataset with lat/long information about where traffic stops occurred), it is possible to aggregate the data to different geographic units (using the standard spatial join procedure), and directly demonstrate how the choice of geographic units might shape our interpretations of a dataset's spatial patterns. This would also demonstrate the value of sensitivity analysis in a spatial analysis or visualization. Different datasets, in short, might lend themselves to different learning objectives and points of emphasis (i.e. if the modifiable aerial unit problem is to be a point of emphasis in the lesson, one might want to use a point-layer and explore different aggregation schemes rather than a dataset in which the geographic unit of analysis is predetermined).

An Overview of a Sample Workshop

The previous section briefly introduced the Stanford Open Policing Project data and the rich possibilities it offers for GIS and data literacy education in academic libraries. It also suggested some general guidelines for selecting specific Project datasets for use in a potential GIS workshop. This section turns to a discussion of a sample workshop that I designed using a dataset from the Open Policing Project. The workshop lesson plan itself is presented in the Supplemental Materials Appendix, and can be used
as a template by geospatial librarians interested in developing such a workshop for their own teaching portfolios.

The workshop presented in the Supplemental Materials uses traffic stop data from the state of Colorado; in particular, it leverages an Open Policing Project dataset of traffic stops by the Colorado State Patrol, which includes information for traffic stops carried out throughout the state between December 2009 and December 2017. The dataset consists of 3,112,853 observations, where each observation (which occupies a distinct row in the dataset) represents the record for a distinct traffic stop. The dataset has several columns that provide

information about each traffic stop. These columns contain information about driver attributes, such as sex and age, as well as outcomes of the police/driver interaction, such as whether the police conducted a search or issued the driver a warning. Most importantly, for our purposes, the dataset also contains information about the motorist's race, and the county in which the stop occurred.

The purpose of the sample workshop is to use the information contained in this dataset to develop a simple county-level indicator of the magnitude of anti-Black racial bias in traffic stops during the year 2010, and display this indicator on a map of Colorado counties. By translating the information in a large dataset into a map that displays potential "problem counties" where Black drivers appeared to face disproportionately aggressive policing, the workshop helps learners explore the spatial dimensions of systemic racism, and recognize the power of GIS as a tool for disaggregating and visualizing data in ways that help identify research questions and hypotheses that merit further exploration.

At a conceptual level, the most challenging question in the sample workshop concerns how we might define and operationalize an index of racial bias in traffic police stops. While the social science literature has developed rigorous ways to measure police discrimination that effectively control for various confounding factors, these techniques may not be suitable to present in an introductory workshop, where learners are often new to data analysis. Given the introductory nature of the workshop, it is important to prioritize intuition and ease of exposition, rather than methodological rigor, in the development of a "bias index". To that end, the workshop adopts an index used by Stelter et al (2021), who operationalize a summary measure of racial bias in police stops by "[subtracting] the percentage of Black residents in each county from the percentage of Black drivers stopped in each county" (7). This yields "a score of disproportionate stopping of Black drivers", in which values greater than zero might be construed as prima facie evidence for anti-Black bias in traffic policing

practices within a given county (7). The workshop makes only one minor modification to the Stelter et al index; in particular, it calculates the index with respect to the adult (over-17) population, since this better approximates the driving population than the overall population. Specifically, the index derived in the workshop subtracts the percentage of adult (over 17 years old) Black residents of a county (relative to the county's over-17 population as a whole) from the percentage of county traffic stops that involved Black drivers. The simplicity of this index means that it is intuitive to explain, and practical to calculate in a real-time workshop setting.

Because this index must be calculated using traffic stop data from Stanford Open Policing Project dataset, as well as demographic information that is measured at the same geographic scale as the traffic stop data (in this case, counties), an important part of an instructor's pre-workshop preparation is to collect this demographic information from the US census, which can then be provided to learners at the beginning of the workshop. The Appendix to the sample workshop provided in the supplemental materials package provides a reproducible R script that shows how the R package *tidycensus* (Walker and Herman 2021) can be used to extract the county-level demographic information used in the workshop; as with the workshop tutorial itself, prospective workshop instructors can adapt this script to generate demographic data that is relevant to their instructional needs.

The sample workshop begins by showing learners how to load the tabular dataset of Colorado police stops into R Studio, and then generate a new variable that extracts the year of a stop from an existing date field that is in "YYYY/MM/DD" format. It then uses this new "date" field to extract observations from the year 2010, and then deploys a variety of functions from R's *dplyr* package[6] to reshape and summarize a new dataset of 2010

[6] *dplyr* is part of the broader suite of packages known as the *tidyverse* (Wickham et al 2019), which figures prominently in the sample tutorial.

observations; the end result is a new data frame that contains information about the percentage of each county's total police stops (in the year 2010) which involved a Black motorist. This data frame is then joined with a dataset of census demographic data which contains information on each county's adult Black population, and overall adult population. The information in this joined dataset is then used to create the bias index described above, by subtracting the county's Black percentage of the adult population from the percentage of its traffic stops that involved Black drivers.

The dataset containing the bias index is then joined to a geospatial vector dataset of Colorado counties (which is provided to workshop participants at the beginning of the lesson). At this point, with the bias index data ready to be visualized, learners are taught how to use the R package *tmap* (Tennekes 2018) to visualize the bias index on a county-level map, and make both static maps and web maps that are suitable for use in publications. The workshop concludes by offering learners the opportunity to reflect on the patterns identified through the mapping exercise, as well as the nature of the basic GIS and data analytic methods that they used.

Advice and Suggestions on Teaching the Workshop

The workshop is designed to be completed in 60 minutes (if the session is demonstration-only, where learners are not expected to implement the workflow on their own machines) to 90 minutes (if it is a hands-on workshop in which learners replicate the code presented by the instructor). The session can also be taught as a "lab session", in which the instructor does not demonstrate the analysis, but rather gives learners a well-documented written lesson (again, the lesson plan in the supplemental materials might be used as a template for this), and asks them to work in small groups to replicate the analysis while the instructor monitors their progress and provides assistance where necessary. A lab session of this kind lends itself to active learning, and encourages learners to think deeply about each step in the analysis. A

more challenging version of this approach would be to supply learners with a lesson plan that implements the analysis in one context, but requires learners to replicate it in another; for example, an instructor based in California could use the lesson plan for Colorado as a teaching tool, and then ask learners to replicate the analysis for California in a lab session.

Just as the lesson plan lends itself to various instructional modalities, it also lends itself to different potential areas of emphasis. For instance, much of the data cleaning and pre-processing work required to make the traffic stop data usable for applied researchers has already been completed by the Open Policing Project itself, but it is likely that in many cases (as with the sample workshop) additional processing and cleaning steps are required before the data is ready to be mapped. If an instructor wanted to emphasize the map-making parts of the lesson, and explore the cartographic capabilities of a given software package in greater detail, they could implement the various data cleaning and preprocessing steps on their own, and begin the lesson at a later stage in the lesson plan. On the other hand, an instructor may want to emphasize the data processing and cleaning steps, which are seen by data scientists as a fundamental part of any empirical data science project (Wickham 2014). To the extent that data processing and cleaning do figure prominently in a particular geospatial workshop that uses Open Policing Project data, it would be a good opportunity to not only emphasize the practical steps involved in implementing these tasks, but to also encourage critical reflection on the data cleaning process itself. As D'Ignazio and Klein (2020) emphasize (in Chapter 5 of their book), data "cleaning" often discards valuable information, which in turn forecloses certain lines of inquiry. By reshaping and reorganizing individual-level data in ways that allow them to be aggregated to geographic units, for instance, we are able to create a choropleth map; however, this process effectively discards the original individual-level data and therefore obscures more granular patterns in the traffic-stop data. Something is gained in this process, but learners should appreciate that something is lost as well; one way to

underscore this point might be to ask them to compare the "cleaned" dataset to the "original" dataset, and identify a research question that could be addressed with the latter but not the former.[7]

The topic of geospatial ethics is extremely broad, and one that merits a dedicated workshop of its own. Nonetheless, geospatial librarians will likely want to briefly discuss geospatial ethics in an introductory workshop, and the one proposed in the sample lesson plan offers meaningful avenues for discussion and exploration along these lines. One possible way into this discussion is to explore the concept of "lying with maps", and show learners how it is possible to manipulate the visual parameters of a map (such as the structure of interval breaks or map colors) to shape the conclusions an audience may draw from it (Monmonier 2018). It is relatively straightforward to adapt the sample workshop's code to demonstrate the implications of such design choices, which will allow learners to appreciate the importance of making choices that minimize the possibility of misleading an audience. Such a demonstration will also empower them to become more critical consumers of the maps and visualizations they see in the media and in scholarly publications.

It is worth noting that the lesson plan lends itself to collaboration between different librarians or instructional professionals. For instance, the lesson plan could lend itself to a division of labor between a data or data science librarian (who could present sections related to data cleaning and processing) and a geospatial librarian (who could discuss the distinctive properties of spatial data and present the sections on mapping, cartography, and geospatial ethics). Geospatial librarian instructors might also collaborate with a subject librarian (perhaps a public policy, criminology, sociology, or political science librarian) who has subject matter expertise in systemic racism and policing. A subject matter expert could help

[7] For a thought-provoking discussion of how a pluralistic and ethically aware conception of data science considers "data cleaning" (which might inform this section of a potential workshop), see Chapter 5 of D'Ignazio and Klein (2020).

to provide useful background information on the workshop's substantive theme of systemic racism, and help lead a synthetic discussion that connects the mapping exercise (taught by the geospatial librarian) to this theme.[8] To be clear, the lesson can be taught by a single person, but such collaborations could help geospatial librarians build professional and intellectual bridges to other areas of digital scholarship and librarianship.

Finally, it is worth noting some of the potential logistical or operational challenges of organizing and delivering a workshop of this nature. If learners are expected to implement the analysis on their own (either by following along with the instructor, or within a lab-style instructional setting), there must be provisions to disseminate the relevant workshop data. In the case of the sample workshop, three distinct datasets are required: the relevant Open Policing Project dataset for Colorado, the demographic dataset used to help calculate the bias index, and a spatial dataset of Colorado counties to which the Open Policing dataset can be joined (since the Policing dataset used in the sample workshop does not contain explicit geographic coordinates). Geospatial librarians may already have established workflows for disseminating workshop data to participants, but it is worth noting that the file sizes for some Open Policing datasets are quite large. As a result, librarian instructors may want to preprocess the data beforehand, so that large file sizes do not cause complications for workshop participants attempting to download the data or load it into a software platform. For example, in the sample workshop, the process of taking a multiyear dataset and filtering observations by year to generate a smaller dataset of observations from the year 2010 was

[8] The sample workshop materials provided in the supplemental materials provides some possible post-workshop discussion questions. Some of these discussion questions, such as one that asks learners to reflect on the implications of the maps they made for public policy, might naturally be led by a subject-librarian expert. If a geospatial librarian is teaching the workshop without a subject-matter expert as a co-instructor, and wishes to engage in a broader discussion of systemic discrimination in the context of the workshop, several resources are available to help instructors navigate potentially fraught classroom discussions of race and racism. A useful place to start is the following reading list from the Center for Racial Justice in Education: https://centerracialjustice.org/resources/reading-lists/.

part of the lesson plan; however, it may make sense for instructors to carry out such operations beforehand, and distribute a dataset of tractable size (rather than a much larger original dataset from the Open Policing Project) at the beginning of the workshop.

Another challenge concerns software installation. It may be advisable to hold an optional "installation clinic" prior to the workshop, which could minimize the workshop time that is devoted to setup and troubleshooting matters. In addition, campus information technology departments are sometimes able to create virtual environments that can bypass the need for local installations; prospective workshop instructors may find it worthwhile to explore the availability of such options on their campuses.

Conclusion

A growing interdisciplinary literature has documented the ways in which contemporary data science and "big data" are implicated in social structures of oppression and discrimination (Noble 2018; Benjamin 2019). At the same time, however, these intellectual and social resources can potentially be reclaimed and reoriented in ways that advance the causes of social justice and racial equity. Indeed, big data and the tools of data science could clarify ways in which racial biases are hardwired into our social institutions, and help us to transform those institutions by inspiring collective action that pushes for reform and accountability. Important coalitions of scholars and activists, increasingly formalized in new organizations such as Data for Black Lives, are animated by this understanding of big data's Janus face, and are spearheading efforts to reclaim and reconfigure tools such as "statistical modeling, data visualization, and crowd-sourcings…[as] powerful instruments for fighting bias, building progressive movements, and promoting civic engagement" (Data for Black Lives, n.d.).

Geographic Information Systems and spatial approaches to data science are likely to play a critical role in such efforts. This represents a civic and pedagogical opportunity for geospatial librarians, who are well-positioned to introduce the communities they serve to the

ways in which GIS can be used to illuminate and document the spatial dimensions of systemic racism. This instructional agenda will require the development of accessible, topical, and thought-provoking lesson plans that address the substantive topic of systemic racism, while simultaneously introducing basic GIS concepts and competencies. The purpose of this essay has been to propose and discuss one such lesson plan, which geospatial librarians might use as a starting point in considering how to use their instructional opportunities to advance the broader causes of racial equity and anti-racism.

Acknowledgements

I would like to thank two anonymous reviewers for insightful suggestions that helped improve this paper, and the associated workshop lesson plan. I would also like to thank colleagues at New York University Libraries and the University of Colorado Libraries for thought-provoking conversations that inspired the development of the workshop discussed in this paper.

Declaration of Interest Statement

There are no competing interests to declare.

References

Benjamin, R. 2019. *Race After Technology: Abolitionist Tools for the New Jim Code.* Cambridge, UK: Polity Press.

Braga, A.A., R.K. Brunson, and K.M Drakulich. 2019. Race, Place, and Effective Policing. *Annual Review of Sociology 45:* 535-555. doi: https://doi.org/10.1146/annurev-soc-073018-022541

Brunner, M., M. Sisk, R. Starry, and S. Tennant. 2022. Geographies of Engaged Digital Scholarship: Remaking Space and Place in the Academic Library. In CLIR Curated Futures, ed. F. Day, J. Maclachlan, J.R. Eyre, S. Smith, C. Williford.
doi: https://doi.org/10.17613/8zza-gc10

Center for Racial Justice in Education. n.d. Reading Lists. Accessed May 20, 2022. https://centerracialjustice.org/resources/reading-lists/

Data for Black Lives. n.d. Data for Black Lives: About Us. Accessed February 15, 2022. https://d4bl.org/about.html

D'Ignazio, Catherine and Lauren F. Klein. 2020. *Data Feminism.* Cambridge, MA: The MIT Press.

Ehrman-Solberg, K., P. Petersen, M. Mills, K. Delegard, R. Mattke. 2020. Racial Covenants in Hennepin County. *University of Minnesota Data Repository.*
doi: https://doi.org/10.13020/a88t-yb14

Harris D. 2010. The stories, statistics, and the law: why "driving while black" matters. In Race, Ethnicity and Policing, ed. SK Rice, MD White, pp. 36–83. New York: NYU Press.

Monmonier, Mark. 2018. *How to Lie with Maps, 3rd Edition.* Chicago: The University of Chicago Press.

Noble, S. 2018. *Algorithms of Oppression: How Search Engines Reinforce Racism.* New York: New York University Press.

Pierson, E., C. Simoiu, J. Overgoor, S. Corbett-Davies, D.Jenson, A. Shoemaker, V. Ramachandran, P. Barghouty, C. Phillips, R. Shroff, and S. Goel. 2020. A large-scale analysis of racial disparities in police stops across the United States. *Nature Human Behaviour 4*: 736-745. doi: https://doi.org/10.1038/s41562-020-0858-1

Stelter, M., I. Essien, C. Sander, and J. Degner. 2021. Racial bias in police traffic stops: White residents' county-level prejudice and Sterotypes are related to disproportionate stopping of Black drivers." *PsyArXiv Preprints.* https://psyarxiv.com/djp8g/.

Tennekes, M. 2018. tmap: Thematic Maps in R. *Journal of Statistical Software 84*(6): 1-39. doi: https://doi.org/10.18637/jss.v084.i06.

Walker, K. and M. Herman. 2021. tidycensus: Load US Census Boundary and Attribute Data as tidyverse and sf-Read Data Frames. R Package version 1.1. https://CRAN.R-project.org/package=tidycensus

Wickham, Hadley. 2014. Tidy Data. *Journal of Statistical Software* 59(10): 1-23. doi: https://doi.org/10.18637/jss.v059.i10

Wickham, H., M. Averick, J. Bryan, W. Chang, L.D. McGowan, R. Francois, G. Grolemund et al. 2019. "Welcome to the tidyverse. Journal of Open Source Software 4(43): 1686. doi: https://doi.org/10.21105/joss.01686.