Analyzing Real Traffic Stop Data for Racial Bias

Vineesh Adala
Duke University

Tim Jun
Duke University

Jason Pereira
University of North Carolina
Chapel Hill

Vivaan Sandwar
University of Michigan

Anthony Fernandes
University of North Carolina Charlotte

Abstract

This student-authored paper describes our—four high school students'—statistical investigation with a sample of 5000 traffic stops from our city. In addition to sharing our analysis to show that Black drivers were disproportionately stopped, we seek to highlight the possibilities that analysis of large datasets open up about talking about race and racism, especially institutional racism. We share this work to inspire teachers to engage their students in developing their critical statistical literacy through investigating real-world data.

Discussion And Reflection Enhancement (DARE) Pre-Reading Questions

1. Have you, a friend, or family members, experienced a traffic stop? Did you feel like there was racial bias in the traffic stop?

2. What is race? [Note that this is challenging to answer and the Supreme Court has ruled this in two different ways in *Ozawa vs U.S.* (1922) and *Thind vs U.S.* (1923)]

3. How can you conceptualize racism as a system?

4. If you were examining a traffic stop dataset to determine if there was bias or not, what variables would you like to use in your analysis? Why?

Vineesh Adala (va138@duke.edu) is a student at Duke University studying Public Policy and Economics. He is interested in data analysis and its role in shaping effective public policy. He plans to pursue a career in law, where he hopes to apply quantitative methods to issues in criminal justice.

Tim Jun (tj144@duke.edu) is a student at Duke University who is currently studying Mathematics and Economics, and is involved in the Applied Machine Learning and Quantitative Finance club. He is interested in teaching younger students and using math to model and better understand issues in the humanities.

Jason Pereira (jasonpereira.518@gmail.com) is an Honors Carolina student at the University of North Carolina at Chapel Hill studying Computer Science and Statistics & Analytics. He is actively involved in CS+Social Good and Velocity Labs, where he applies artificial intelligence and machine learning to real-world challenges. He is an Eagle Scout and TEDx Speaker, and will be interning at Amazon Web Services (AWS). His research interests lie in AI and machine learning, particularly in developing data-driven solutions for society.

Vivaan Sandwar (vsandwar@umich.edu) is a student at the University of Michigan - Ann Arbor, studying Mathematics and Economics. On campus, Vivaan is involved in Michigan Venture Club, Michigan Engineering Consulting Group, and Wolverine Sports Analytics. Passionate about the intersection of humanities and financial mathematics, Vivaan is interested in learning more about how to create societal change through the quantitative lens of mathematics.

Anthony Fernandes (anthony.fernandes@charlotte.edu) is Professor in the Department of Mathematics and Statistics at University of North Carolina at Charlotte. He is interested in how statistical analysis of large datasets on social issues can promote critical consciousness of future mathematics and statistics teachers.

Acknowledgements: This article is based upon work partially supported by the National Science Foundation (NSF) under Grant No. 2121364. Any opinions, findings, and conclusions or recommendations expressed in this article are those of the authors and do not necessarily reflect the views of the NSF.

Analyzing Real Traffic Stop Data for Racial Bias

Vineesh Adala, Tim Jun, Jason Pereira, Vivaan Sandwar, & Anthony Fernandes

We are four high school students at Ardrey Kell High School in Charlotte, North Carolina, who came together around our interest in data science. We wanted to get exposure to data science and got in touch with the last author who had designed data investigation modules for mathematics pre-service teachers, as part of their content courses. We were intrigued by the traffic stops module that used real data from the city that we lived in. Though our main interest was gaining experience with data investigations, we learned the importance of understanding the context to interpret the statistical results. The traffic stop investigation was related to racial justice. Based on our own experiences, and our friends', we were familiar with traffic stops and were aware of the issue of 'Driving While Black' where Black drivers are more likely to be stopped. However, we had not seen if the same was true with local data. In this article we describe our experience working with the traffic stops data. We believe that our experience and insights will help teachers who are interested in working with real data to draw insight into a particular social issue.

Dataset and Work

The City of Charlotte regularly publishes traffic stop data. We worked with a random sample of 5000 traffic stops, which were taken from about 124,000 traffic stops between January 2020 and October 2021. We used the Common Online Data Analysis Platform (CODAP) to complete the analysis. The dataset included variables such as the month of the stop, race and gender of the driver and officer, driver ethnicity, was a search conducted, police division, officer age, the reason for the stop, and the result of stop. We met online regularly as a group for about 8 months to reflect on our work. In the initial meetings, we learned how to use CODAP and then attempted to answer the overarching research question: Is there racial bias in the policing of traffic stops in Charlotte? Based on the research question we worked individually, exploring statistical questions that related to the overarching research question, then shared our work and had discussions during the meetings. Below we share samples of explorations using descriptive statistics and informal inferential statistics. Given that Black[1] and white drivers were a majority of the stops conducted, we decided to focus on these races in the questions we asked.

[1] Our motivation to capitalize Black, and not white is guided by Laws (2020) [https://www.cjr.org/analysis/capital-b-black-styleguide.php]

Are Black Drivers Searched More Often Than White Drivers After a Stop?

We created a two-way table/graph between driver race and whether a search was conducted, and filtered this by the race of the driver. Black drivers formed 76% of the drivers who were searched, compared to white drivers who were 21% of those searched (Figure 1). This seemed excessive given that Black drivers were 56% of all the stops in the sample.

Figure 1

Searches by driver race (Percents by column)

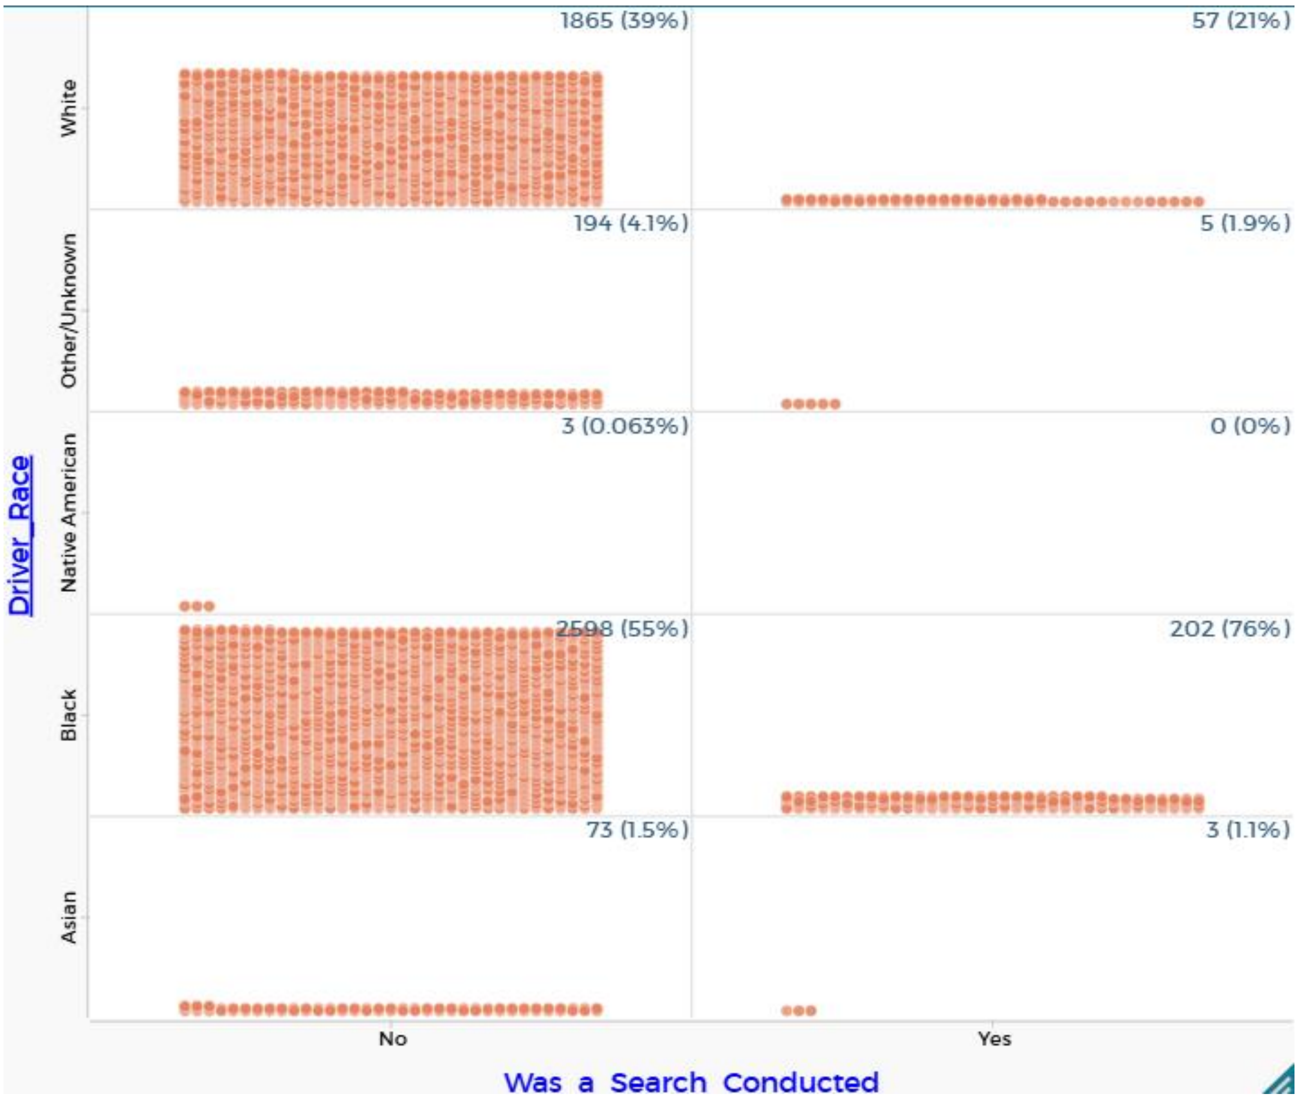

Further, looking across by race (Figure 2), 7.2% of Black drivers were searched, compared to 3% of white drivers. Thus, Black drivers were 2.4 times more likely to be searched.

Figure 2

Searches by driver race (Percents by rows)

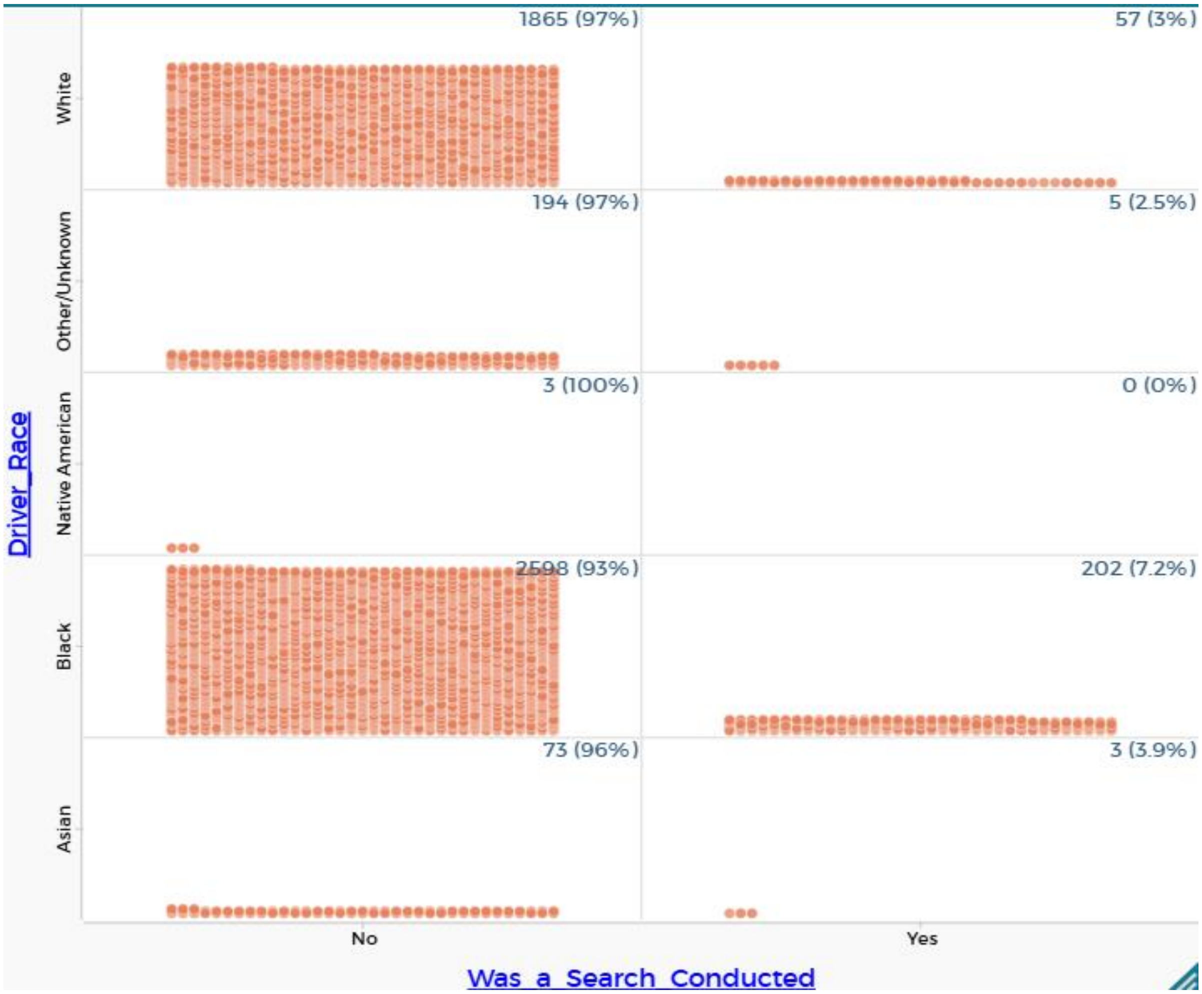

Figures 1 and 2 provided some indication that there could be bias on the part of the officers as they decided who to search. However, we could not be quite sure so we decided to look at the consequences after the search. We conjectured that if the officers' suspicion was not justified, then the driver would have a low consequence. This led to the next question.

Are There Any Differences in the Severity of Consequences for Black Drivers and White Drivers Who Are Searched?

After a stop the drivers get one of the following consequences: no action taken, verbal warning, written warning, citation, or arrest. To answer our question, we created two categories: severe and nonsevere. Severe consequences included citations issued and arrests, and nonsevere consequences included no action taken, verbal warnings, and written warnings. We dropped the severity variable in the middle of Figure 2 (a feature of CODAP) and obtained Figure 3 which allowed us to filter the drivers, who were searched or not, by severity. We also looked at each cell individually and selected the nonsevere consequences. From Figure 3, we see that 58% of Black drivers who were searched received a nonsevere consequence (118 out of 202), compared to 44% for white drivers (25 out of 57). Thus, we see that (a) the suspicion of the police officer seems to play a role in searches, and (b) for most of the Black drivers the search does not yield a severe consequence.

Figure 3

Severity of consequence after a search (Totals by cell)

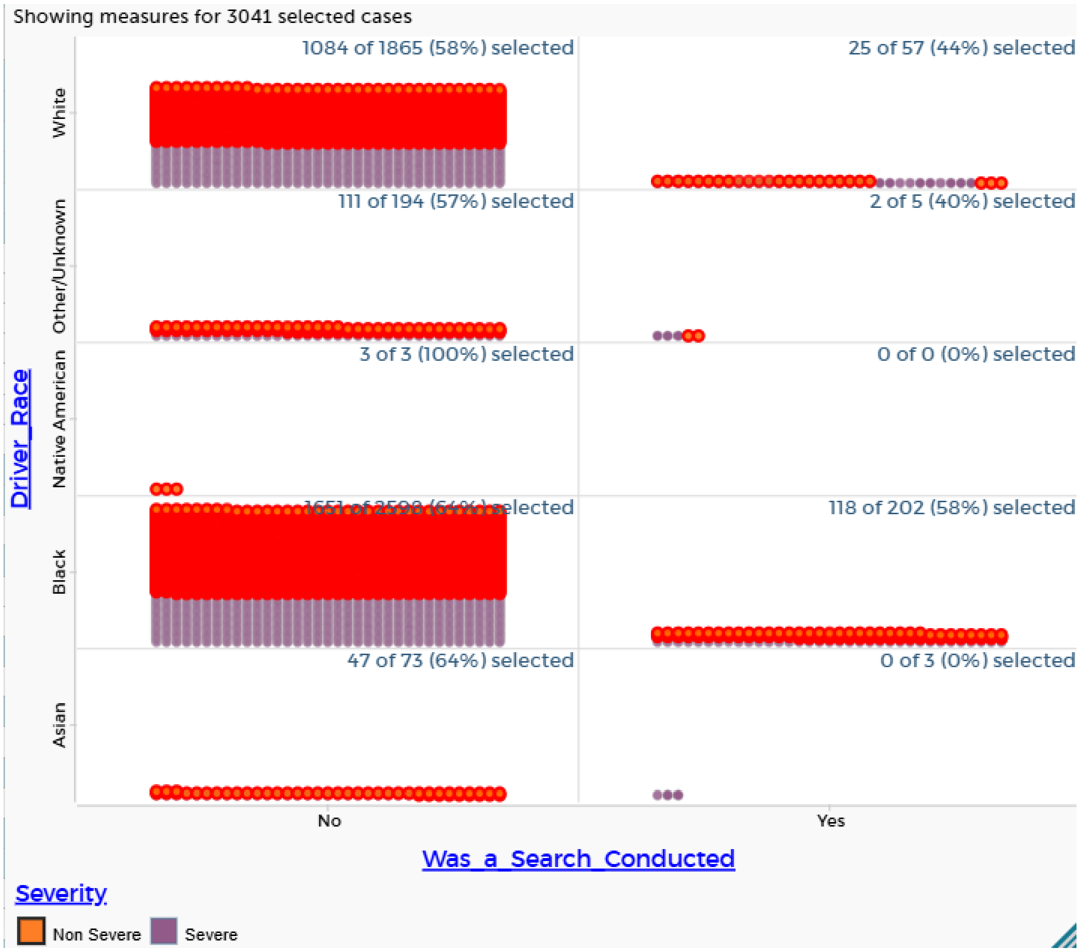

Are There Differences Between the Race of the Driver That More Experienced Officers Stop, Compared to Less Experienced Officers?

Another exploration was related to the officers' years of service. We believed that officers with fewer years on the force would most likely be younger,would have had interactions with people from other races, and would be more receptive to diversity training. On the other hand, we assumed that officers with more years of service would be older and less likely to be open to diversity training. Looking at the officer years of service partitioned by driver race (Figure 4), we observed that Black drivers were stopped by officers who, on average, spent less time in the force (median 6 years) compared to white drivers, who were stopped by officers with more time on the force (median 12 years). We noted that our assumption about the officers' years of service and the drivers they stopped were not supported by the data. Instead, we wondered if newer officers were being sent to areas that were predominantly Black?

Figure 4

Differences in stops made by race and years of service

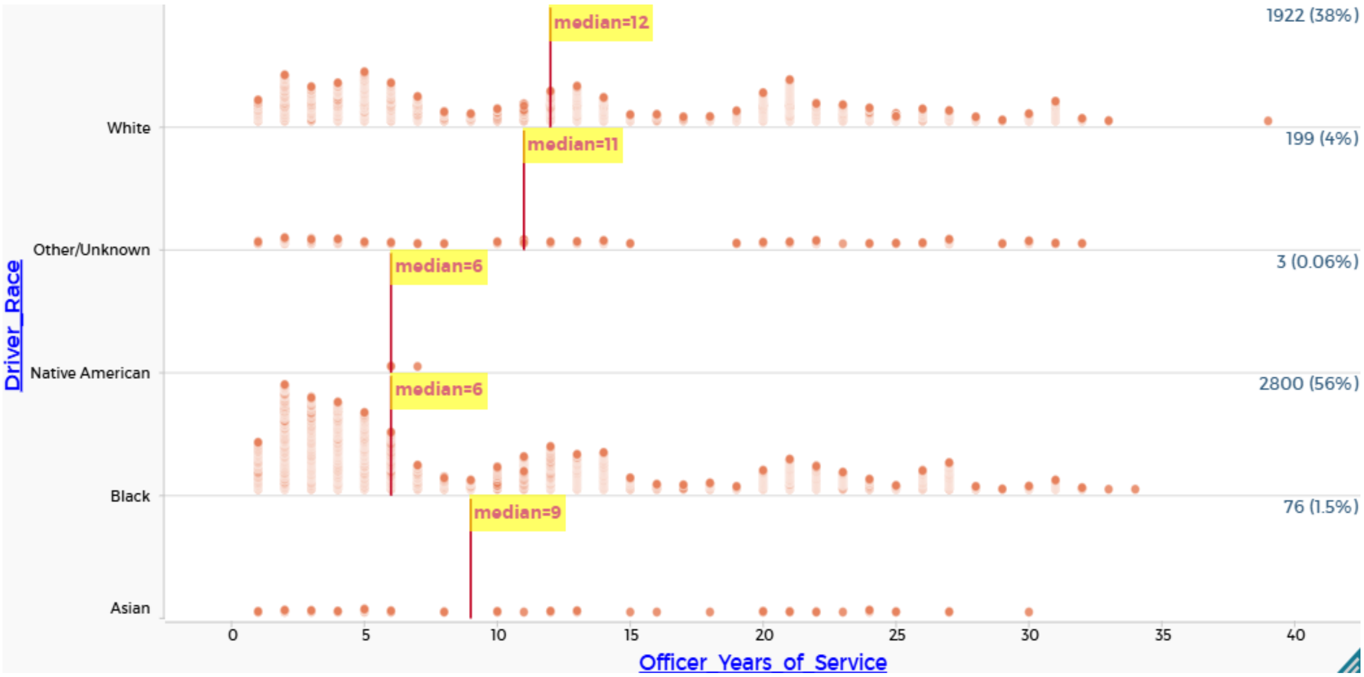

We decided to filter the data by the race of the officer. From Figure 5 we see that the median years of service of Black officers who stop white drivers is 6.5 years and Black officers who stop Black drivers is 6 years. When we filter for white officers (Figure 6), the median years of service is 13 years for officers stopping white drivers, and 6 years for officers stopping Black drivers. From these results we inferred that there are racially segregated neighborhoods in the city that are mostly patrolled by newer officers, both Black and white. We wondered if newer officers were being sent to predominantly Black neighborhoods that were perceived as being more dangerous? Note that we cannot draw causal conclusions from this data and it is possible that deployment practices varied by department and policy.

Are There Differences Between the Proportion of Black Residents in the City and the Proportion of Black Drivers Who Are Stopped?

After the initial exploration, a major question that remained was how we could determine if there was a significant difference between the proportion of Black drivers who were stopped and the proportion of Black residents in the city. One graph that we examined at the beginning (Figure 7) showed that 56% of the drivers stopped were Black. Could it be that the proportion of Black residents in the city was close to 56%?

Figure 5

Differences in stops made by race and years of service, filtered by Black officers

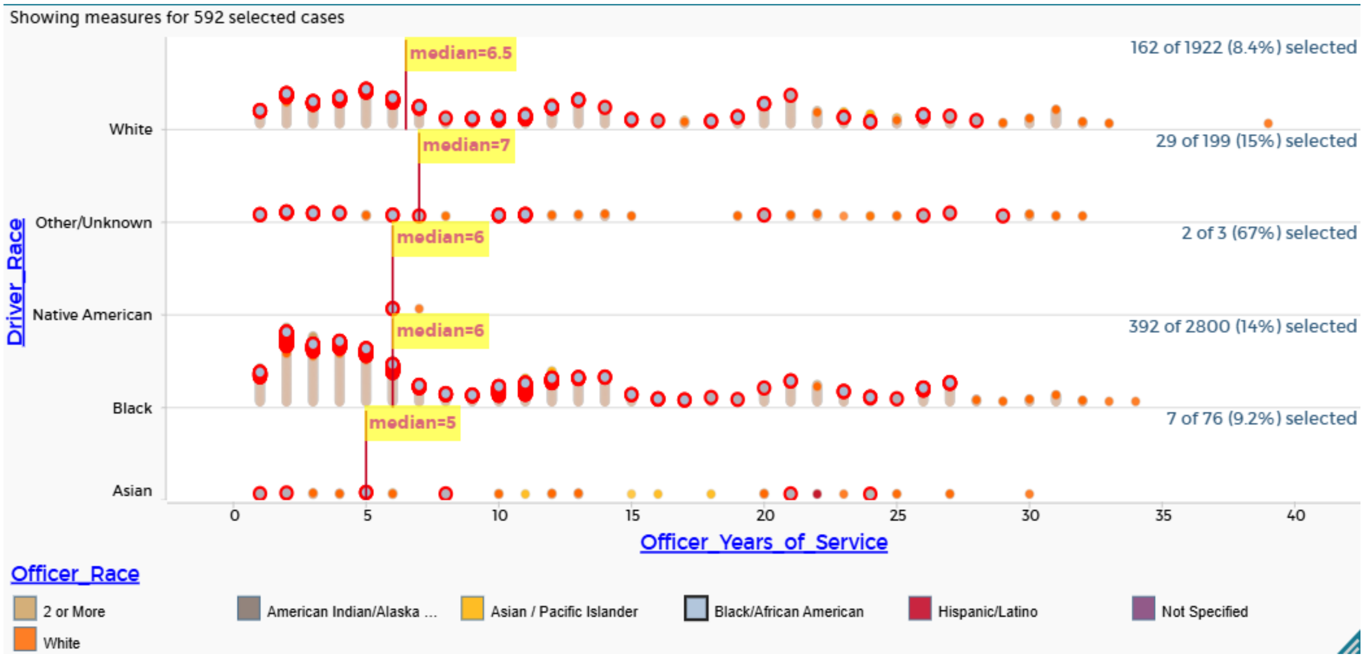

Figure 6

Differences in stops made by race and years of service, filtered by white officers

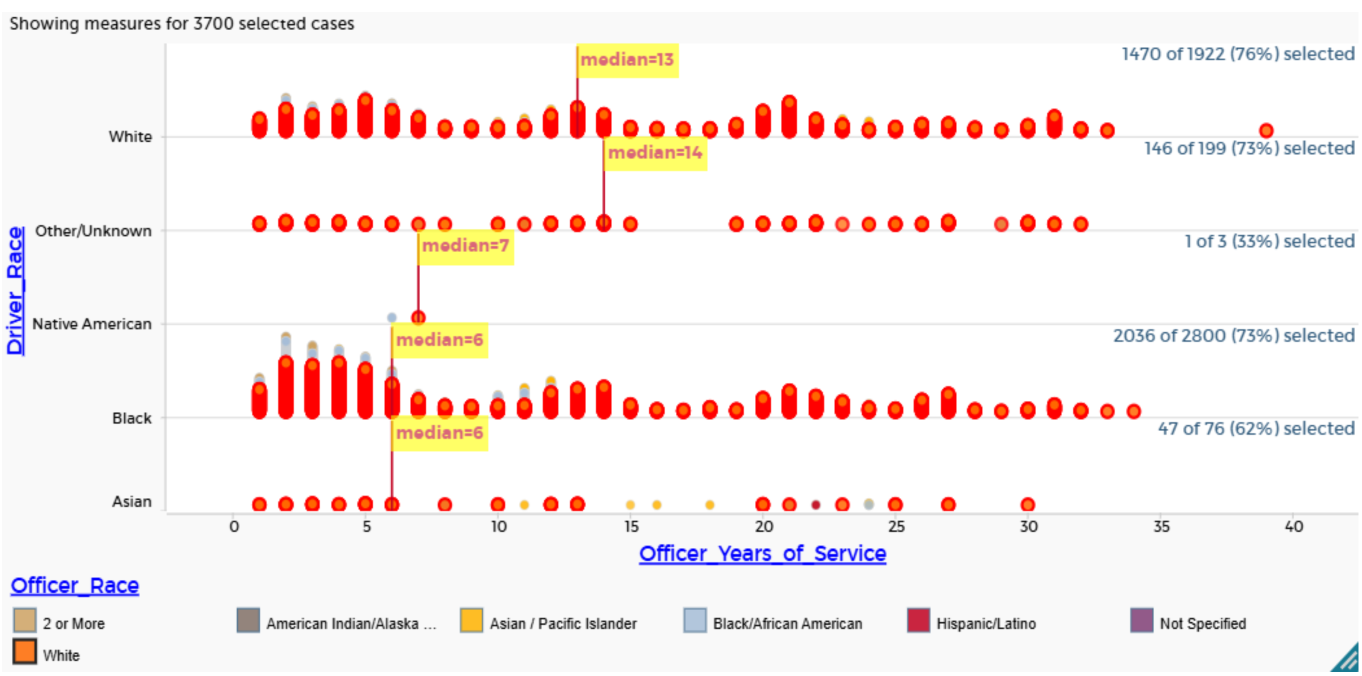

Figure 7

Drivers stopped in the city by race (Asian - 1.3%, Black - 56%, Native American - 0.066%, Other/Unknown - 4%, white - 38%)

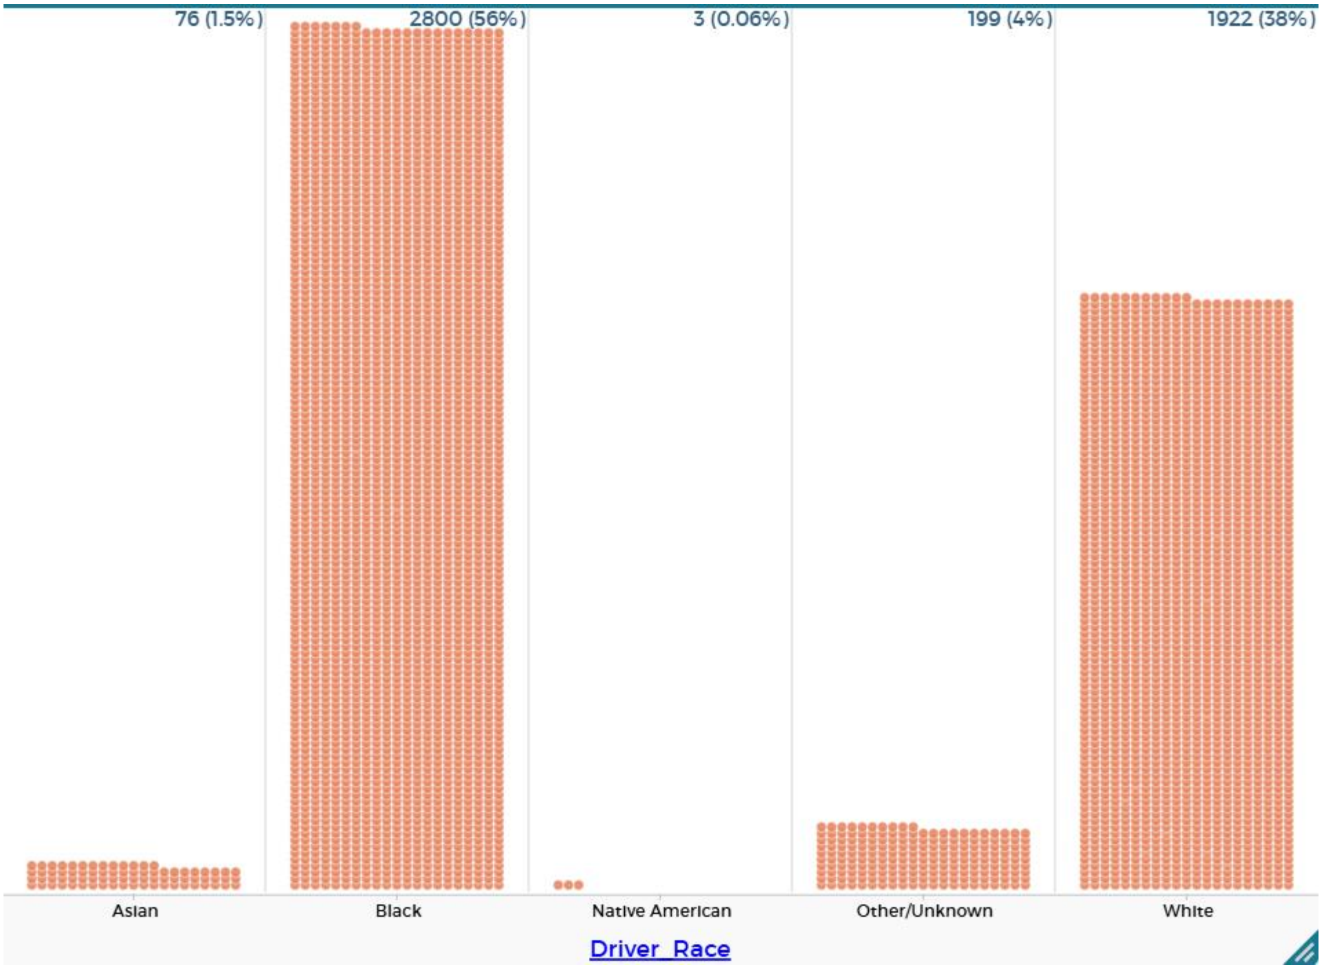

Using the 2020 Census, we found that 33% of residents were Black. Thus, there seemed to be a notable disparity in the proportions of Black residents to Black drivers that were stopped. To investigate our question, we conducted an informal inference by designing a simulation (Zieffler et al., 2008). We imagined 5000 residents that reflected the racial demographics of the city and drew random samples of size 50. We used n=50 based on the assumption that the police stopped about 120,000 drivers in a city with a population of about 800,000. This would lead to using a sample of about 750 with 5000 stops. However, using a much smaller number like 50 would account for more variation in the sampling distribution and thus highlight the unusual nature of the observed result. Note that n=50 was greater than the minimum sample size of 30, yielding a distribution of proportions that could be approximated using the normal distribution according to the Central Limit Theorem. After drawing 100 samples, we approximated the 95% confidence interval by bounding 95 of the 100 dots. We observed that none of the random samples had a proportion as high as 56% (Figure 8). Thus, we knew that we were observing something that was highly unlikely to happen by chance and that there was an underlying reason for the difference.

Figure 8

Results of the simulation of Black drivers stopped in the city. Includes the approximate 95% confidence interval estimate of the proportion of Black drivers.

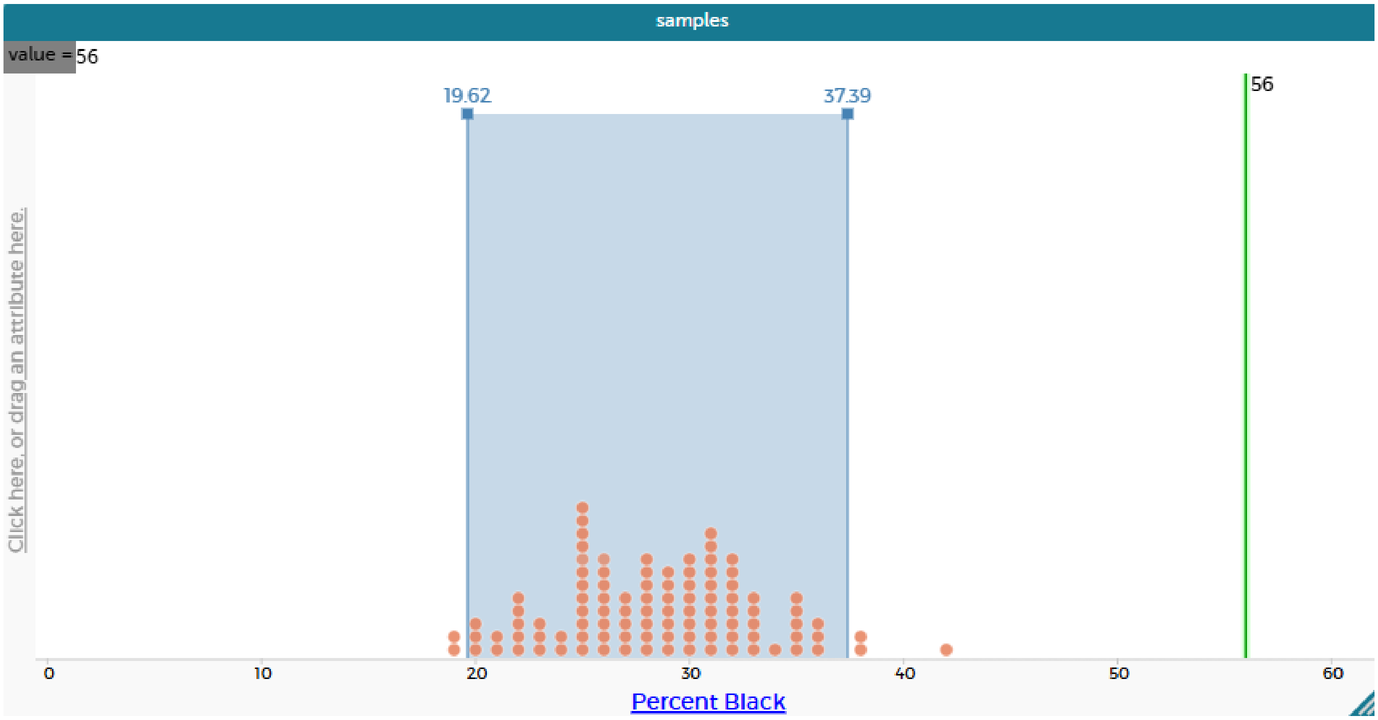

We examined two reasons for the difference in proportions: algorithms and laws. The city used PredPol, a predictive algorithm that was used to guide the placement of officers in various areas of the city based on the street crime data (O'Neil, 2017). We were unclear about the connection between street crime and traffic stops. Further, the presence of officers in over-policed areas meant that they would report more crime in these areas compared to other (predominantly white) areas that were not patrolled as often or policed with officers intent on finding crime (Gordon, 2022). The algorithm would add the new data and decide that the officers patrol the same areas due to a predominance of crime, thus creating a feedback loop (O'Neil, 2017). This could explain the disproportionality that we observed. We also examined two court cases that impact the policing of traffic stops - *Terry vs Ohio* (1968) and *Whren vs U.S.* (1996). According to Baumgartner et al. (2018), *Terry* lowered the bar of policing from 'probable cause' to 'reasonable suspicion', which could let bias enter the officer's decision. Further, *Whren* opened the door for pretextual stops based on traffic laws. Together the two cases opened the possibility for more stops to be made based on suspicion. We believe that the lighter consequences for Black drivers, including when their vehicles are searched, points to police bias.

In the end, we concluded that there was racial bias in policing, and it was because of a combination of factors that went beyond just individual officers. Historically, segregation caused a separation in the population by race that still seemed to persist over time. Amplifying this, were the laws and policies in policing that allowed officers to act on suspicion. Finally, the use of algorithms caused a feedback loop that magnified the problem.

Learning

This activity was the first time that we had worked with a large real-world dataset. Working with data from our city was motivating. Usually, we worked with small made up datasets where there was always a 'right' answer. We also learned that data exploration with real data was not as cut and dry as we initially believed, with only one answer. We had to construct various graphs to answer the overarching research question. Further, we had to draw on an extensive array of statistical techniques. Given the real-world nature of the investigation, we were also not sure of the path to take initially, so there were some ‘blind alley’ explorations. However, this indirectly pointed to more promising ways to explore the data. Throughout the process we had to discuss our ideas, helping us to gain a better understanding of statistical concepts and how to apply them in real-world scenarios. For example, we had learned about hypothesis testing, but we had not taken a simulation approach, which gave us a better understanding of how hypothesis tests work.

In addition to getting a better grasp on the statistics, we also learned about race and racism in the context of policing traffic stops. In this activity, the context was crucial to making sense of the statistical analysis that we were doing. For example, we were forced to think about the context when we had to explain the gap between the proportion of Black residents in the city and the proportion of Black drivers who were stopped. This gap was not by chance, and in our quest to understand the gaps, we learned about the broader factors affecting these disparities like the court cases, segregation, and the policing algorithms. In the process, we broadened our understanding of racism and how it operates in the system. For instance, that racism does not just happen between individuals. Rather, the disparate impact of laws and policies on certain groups also constitutes a form of racism. We learned that this is referred to as institutional racism, which according to Racial Equity Tools (2021), “refers specifically to the ways in which policies and practices of organizations or parts of systems (schools, courts, transportation authorities, etc.) create different outcomes for different racial groups” (Racial Equity Tools, 2021, Glossary).

Recommendations

As students ourselves, we have not experienced using real-world and civically important data in the classroom. However, after engaging in this project, we noticed that the scope of mathematics is larger than numbers; applying mathematics to real world issues like racism can improve students’ understanding of the society around them. For example, we found that racism is a much deeper issue than the surface-level knowledge that most students have and penetrating the systems that we operate within. By digging deeper into these concepts, the students of the future have the potential to be more educated and aware of the problems at hand, enabling them to mitigate issues of racism. Though we recognize that race-based e discussions can be controversial, many of our peers expressed interest in engaging in these types of activities. They felt that if the teacher set up an open environment for discussion, students could learn a lot from each other. We spoke with an AP Statistics teacher at our school, and he agreed that it would be beneficial for the students if they worked with more civically relevant datasets. However, we also understand that conversations about policing and traffic stops can be emotionally triggering, particularly for Black and other students of color. To mitigate this, teachers can normalize emotional reactions in the class, offer students different options in how they choose to participate, and set up norms for participation, to name a few teaching approaches for potentially controversial topics (see Appendix for more resources). Further, we do not need the teachers to be “experts” on racism to do this work. Instead, teachers can be reflective and thoughtful about how they frame and facilitate the investigation, emphasizing humility, preparation, and co-learning with students.

Allowing students to work through data sets, figuring out what works and what does not work, gives students the opportunity to explore different conclusions from data. Students are usually limited to worksheets and homework, but more freedom in the classroom to explore personally relevant data sets would allow them to better understand what they are being taught. For instance, the data from this study connected to our personal experiences of seeing our adult family members stopped by police and a number of our peers drive, so we were better able to reach the goals of the learning experience.

To set up data analysis projects, we believe teachers should provide students with a pre-existing dataset. The dataset can originate from local data or be something that pertains to the students. Using software like CODAP can allow students to handle larger datasets. Students will need some guidance initially, so it is important to have examples of graphs and tables that they can take inspiration from. However, when students start to encounter more advanced concepts in the software like statistical inference, it is helpful to provide some direct instruction on the concept before allowing students to explore it in the software. For example, educators should explain what inference is and its various types before allowing students to apply inferential methods to the dataset. A project like this is a hands-on experience, so it would be nice to allow students to produce their own end products and conclusions.

Students will come to class with their own worldviews and biases that will come through in their investigation. Instead of minimizing these biases, teachers could encourage them to use statistics to justify their point of view. For example, it was challenging for us to understand why the gap between the proportion of Black drivers stopped was so much more than the proportion of Black residents in the city. Even if we came into the class with a pro police bias, we would still need to explain the gap we observed in the analysis.

Critical court cases and the history behind the data gave us a more open-minded perspective on our data exploration. For example, when we didn't see any discrepancies between Black and white drivers, our first instinct wasn't to reject ideas of institutional racism; instead, we referred back to the court cases to hypothesize whether they had an influence on the results we were observing. To help students learn about these concepts, it is important to provide them with the historical and legal influences that have shaped them. This will inform their analysis.

References

Baumgartner, F. R., Epp, D. A., & Shoub, K. (2018). *Suspect citizens: What 20 million traffic stops tell us about policing and race*. Cambridge University Press.

Gordon, D. (2022). *Policing the racial divide: Urban growth politics and the remaking of segregation*. New York University Press.

O'Neil, C. (2017). *Weapons of math destruction*. Penguin Books.

Ozawa v. United States, 260 U.S. 178 (1922). https://supreme.justia.com/cases/federal/us/260/178/

Racial Equity Tools (2020). *Racial Equity Tools glossary*. Available at https://www.racialequitytools.org/glossary.

Terry v. Ohio, 392 U.S. 1 (1968). https://supreme.justia.com/cases/federal/us/392/1/

United States v. Bhagat Singh Thind, 261 U.S. 204 (1923). https://supreme.justia.com/cases/federal/us/261/204/

Whren v. United States, 517 U.S. 806 (1996). https://supreme.justia.com/cases/federal/us/517/806/

Zieffler, A., Garfield, J., Delmas, R., & Reading, C. (2008). A framework to support research on informal inferential reasoning. *Statistics Education Research Journal*, *7*(2), 40–58. https://doi.org/10.52041/serj.v7i2.469

Discussion And Reflection Enhancement (DARE) Post-Reading Questions

1. Based on your reading of the article, what recommendations would you give the police department to lower racial bias?

2. What is your understanding of institutional racism now?

3. What other datasets can you use for statistical investigations around race, racism, and racial bias?

4. Do you think that the context interferes with learning statistics in investigations like the one in this article?

Appendix

- Project website: https://pages.charlotte.edu/datamodulesandsystemicracism/
  - Includes slides, instructor guides, and resources, for two modules: traffic stops and school discipline
- City of Charlotte Open Data Portal (for Charlotte traffic stop data): https://data.charlottenc.gov/
- The Stanford Open Policing Project (for traffic stop data of other cities across the U.S.): https://openpolicing.stanford.edu/

Resources for Conversations about Race and Racism

- Chapter 2 from *High School Mathematics Lessons to Explore, Understand, and Respond to Social Injustice*
- Norm Setting Activity from the A[3]IMS project
- Arao, B., & Clemens, K. (2013). Title of chapter. In L. M. Landreman (Ed.), *The art of effective facilitation: Reflections from social justice educators* (pp. 135 - 150). Routledge. https://doi.org/10.4324/9781003447580
- Let's Talk: Discussing Race, Racism and Other Difficult Topics with Students (from *Learning for Justice*)
- Singleton, G. E. (2021). *Courageous conversations about race: A field guide for achieving equity in schools* (3rd ed.). Corwin Press.
- The four agreements of courageous conversations about race (5:25 minute video)

Resources for Using CODAP

- Awash in Data. An ever-evolving open book on teaching lessons on data science using CODAP
- For Educators: Teaching with CODAP. The creators of CODAP created this page of resources for educators using CODAP in their teaching