

Avoiding bias when inferring race using name-based approaches

Diego Kozlowski¹, Dakota S. Murray², Alexis Bell³, Will Hulsey³, Vincent Larivière⁴, Thema Monroe-White³, Cassidy R. Sugimoto²

¹ diego.kozlowski@uni.lu

DRIVEN DTU-FSTM, University of Luxembourg, Luxembourg

² dakmurr@iu.edu; sugimoto@indiana.edu

School of Informatics, Computing, and Engineering, Indiana University Bloomington, IN, USA

³ alexis.bell@vikings.berry.edu; william.hulsey@vikings.berry.edu; tmonroewhite@berry.edu

Campbell School of Business, Berry College, GA, USA

⁴ vincent.lariviere@umontreal.ca

École de bibliothéconomie et des sciences de l'information, Université de Montréal, Montréal, QC, Canada

Abstract

Racial disparity in academia is a widely acknowledged problem. The quantitative understanding of racial-based systemic inequalities is an important step towards a more equitable research system. However, few large-scale analyses have been performed on this topic, mostly because of the lack of robust race-disambiguation algorithms. Identifying author information does not generally include the author's race. Therefore, an algorithm needs to be employed, using known information about authors, i.e., their names, to infer their perceived race. Nevertheless, as any other algorithm, the process of racial inference can generate biases if it is not carefully considered. When the research is focused on the understanding of racial-based inequalities, such biases undermine the objectives of the investigation and may perpetuate inequities. The goal of this article is to assess the biases introduced by the different approaches used name-based racial inference. We use information from US census and mortgage applications to infer the race of US author names in the Web of Science. We estimate the effects of using given and family names, thresholds or continuous distributions, and imputation. Our results demonstrate that the validity of name-based inference varies by race and ethnicity and that threshold approaches underestimate Black authors and overestimate White authors. We conclude with recommendations to avoid potential biases. This article fills an important research gap that will allow more systematic and unbiased studies on racial disparity in science.

Introduction

The use of racial categories in the quantitative study of science is so widely extended that it intertwined with the controversial origins of statistical analysis itself (Galton, 1891; Godin, 2007). While Galton and the eugenics movement reinforced the racial stratification of society, racial categories have also been used to acknowledge and mitigate racial discrimination. As Zuberi (2001) explains: "The racialization of data is an artifact of both the struggles to preserve and to destroy racial stratification." This places the use of race as a statistical category in a precarious position: one that both reinforces the social processes that segregate and disempower parts of the population, while simultaneously providing an empirical basis for understanding and mitigating inequities.

Science is not immune from these inequities (Hoppe, et al., 2019; Prescod-Weinstein, 2020; Stevens et al, 2021; Ginther et al., 2011). Early research on racial disparities in scientific publishing relied primarily on self-reported data in surveys (e.g., Hopkins et al., 2013), geocoding (e.g., Fiscella & Fremont, 2006) and directories (e.g., Cook, 2014). However, there is an increasing use of large-scale inference of race based on names (Freeman & Huang, 2014), similar to the approaches used for gender-disambiguation (e.g., Larivière et al., 2013). Algorithms, however, are known to encode human biases (Buolamwini & Gebru, 2018; Caliskan, 2017): there is no such thing as *algorithmic neutrality*. The automatic racial inferences of authors based on their features in bibliographic databases is itself an algorithmic

process that needs to be scrutinized, as it can fall into implicitly encoded bias, with major impact in the over and under representation of racial groups.

In this study, we use the self-declared race/ethnicity from the 2010 U.S. census and mortgage applications as the basis for inferring race from author names on scientific publications indexed in the Web of Science (WoS) database. Bibliometric databases do not include self-declared race by authors, as they are based on the information provided in publications, such as given and family names. Given that the U.S. census provides the proportion of self-declared race by family name, this information can be used to infer US authors' race given their family names. We present several different approaches for inferring race and examine the bias generated in each case. The goal of the research is to provide an empirical critique of name-based race inference and recommendations for approaches that minimize bias. Even if a perfect inference is not achievable, the conclusions that arise from this study will allow researchers to conduct more careful analyses on racial and ethnic disparities in science. Although the categories analysed are only valid in the U.S. context, the general recommendation can be extended to any other country in which the Census (or similar data collection mechanism) includes self-reported race.

Racial categories in the US census

The U.S. census is a rich and long-running dataset, but also deeply flawed and criticized. Currently it is a decennial counting of all U.S. residents, whether citizens or non-citizens, in which several characteristics of the population are gathered, including self-declare race/ethnicity. The classification of race in the U.S. census is value-laden with the agendas and priorities of its creators, namely 18th century White men who Wilkerson (2020) refers to as “the dominant caste.” The first U.S. census was conducted in 1790 and founded on the principles of racial stratification and White superiority. Categories included: “Free White males of 16 years and upward,” “Free White males under 16 years;” “Free White females,” “All other free persons,” and “Slaves” (U.S. Bureau of the Census, 1975). At that time, each member of a household was classified into one of these five categories based on the observation of the census-taker, such that an individual of “mixed white and other parentage” was classified into “All other free persons” in order to preserve the “Free White...” privileged status. To date, anyone classifying themselves as other than “non-Hispanic White” is considered a “minority.” The shared ground across the centuries of census survey design and classification strata reflects the sustained prioritization of the White male caste (D’Ignasio & Klein, 2020, Zuberi, 2001).

Today, self-identification is used to assign individuals to their respective race/ethnicity classifications (Locke, Blank, & Groves, 2011), per the U.S. Office of Management and Budget (OMB) guidelines. However, the concept of race and/or ethnicity remains poorly understood. For example, in 2000 the category “Some other race” was the third largest racial group, consisting primarily of individuals who in 2010 identified as Hispanic or Latino (which according to the 2010 census definition refers to a person of Cuban, Mexican, Puerto Rican, South or Central American, or other Spanish culture or origin regardless of race). Instructions and questions which facilitated the distinction between race and ethnicity began with the 2010 census which stated that “[f]or this census, Hispanic origins are not races” and to-date, in the U.S. federal statistical system, Hispanic origin is considered to be a separate concept from race. However, this did not preclude individuals from self-identifying their race as “Latino,” “Mexican,” “Puerto Rican,” “Salvadoran,” or other national origins or ethnicities (Humes, Jones, & Ramirez, 2011, p. 3). Furthermore, 6.1% of the US population changed their self-identification of both race and ethnicity between the 2000 and 2010 censuses (Liebler, et al., 2017), demonstrating the dynamicity of the classification. The inclusion of certain categories

has also been the focus of considerable political debate. For example, the inclusion of citizenship generated significant debates in the preparation of the 2020 Census, as it can generates a larger nonresponse rate from the Hispanic community (Baum, 2019). For the aim of this article, however, is to represent the full extent of possible US-based authors. Therefore, we consider both citizens and non-citizens.

In this paper, we design and compare multiple approaches for inferring race based on the given and family names of authors in the WoS database. The social function of the concept of race (i.e., the building of racialized groups) underpins its definition more than any physical traits of the population. For example, "Hispanic" as a category arises from this conceptualization, even though in the 2010 U.S. census the question about Hispanic origin is different from the one on self-perceived race. While Hispanic origin does not relate to any physical attribute, it is still considered a socially racialised group, and this is also how the aggregated data is presented by the Census Bureau. Therefore, in this paper, we will utilize the term race to refer to these social constructions, acknowledging the complex relation between conceptions of race and ethnicity. But even more important, this conceptualization of race also determines what can be done with the results of the proposed models. Given that race is a social construct, inferred racial categories should only be used in the study group-level social dynamics underlying these categories, and not as individual-level traits. Census classifications are founded upon the social construction of race and reality of racism in the US, which serves as "a multi-level and multi-dimensional system of dominant group oppression that scapegoats the race and/or ethnicity of one or more subordinate groups" (Horton & Sykes, 2001, p. 209). Self-identification of racial categories continue to reflect broader definitional challenges, along with issues of interpretation, and above all the amorphous power dynamics surrounding race, politics, and science in the U.S. In this study, we are keenly aware of these challenges, and our operationalization of race categories are shaped in part by these tensions.

Data

This project uses several data sources to test the different approaches for race inference based on the author's name. First, to test the interaction between given and family names distributions, we simulate a dataset that covers most of the possible combinations. Using a Dirichlet process (Teh, 2010), we randomly generate 500 multinomial distributions for given names, and another 500 random multinomial distributions for family names. After this, we build a grid of all the possible combinations of given and family names random distributions (250,000 combinations).

In addition to the simulation, we use two datasets with real given and family names and an assigned probability for each racial group. The data from the given names is from Tzioumis (2018), who builds a list of 4,250 given names based on mortgage applications, with self-reported race. Family name data is based on the 2010 US census (US Census Bureau, 2016), which includes all family names with more than 100 appearances in the census, with a total of 162,253 surnames that covers more than 90% of the population. For confidentiality, this list removes counts for those racial categories with fewer than five cases, as it would be possible to exactly identify individuals and their self-reported race. In those cases, we replace with zero and renormalize. As explained previously, changes were introduced in the 2010 US census racial categories. Questions now include both racial and ethnic origin, placing "Hispanic" outside the racial categories. The racial categories used in both datasets include Hispanic as a category, and all other racial categories excluding people with Hispanic origin, therefore the category "White" becomes "Non-Hispanic White Alone", and "Black or African American"

becomes "Non-Hispanic Black or African American Alone", and so on. The final categories used in both datasets are:

- Non-Hispanic White Alone (*White*)
- Non-Hispanic Black or African American Alone (*Black*)
- Non-Hispanic Asian and Native Hawaiian and Other Pacific Islander Alone (*Asian*)
- Non-Hispanic American Indian and Alaska Native Alone (*AIAN*)
- Non-Hispanic Two or More Races (*Two or more*)
- Hispanic or Latino origin (*Hispanic*)

We test these data on WoS to study how name-based racial inference performs on the population of U.S. scientific authors. WoS did not regularly provide first names in articles before 2008; therefore, the data includes all articles published between 2008 and 2019. This results in 21,295,333 articles, 1,609,107 unique first authors, 152,835 distinct given names and 288,663 distinct family names for first authors. Given that in this database, 'AIAN' and 'Two or more' account for only 0.69% and 1.76% of authors respectively, we remove these and renormalize the distribution with the remaining categories. Therefore, in what follows we will refer exclusively to categories *Asian*, *Black*, *Hispanic*, and *White*.

Methods

Manual validation

The data is presented as a series of distributions of names across race (Table 1). In name-based inference methods, it is not uncommon to use a threshold to create a categorical distinction: e.g., using a 90% threshold, one would assume that all instances of Juan as first name should be categorized as Hispanic and all instances of Washington as a given name should be categorized as Black. In such a situation, any name not reaching this threshold would be excluded (e.g., those with the last name of "Lee" would be removed from the analysis). This approach, however, assumes that the distinctiveness of names across races does not significantly differ.

Table 1. Sample of family names (US census) and given names (mortgage data).

Type	Name	Asian	Black	Hispanic	White	Count
Given	Juan	1.5%	0.5%	93.4%	4.5%	4,019
	Doris	3.4%	13.5%	6.3%	76.7%	1,332
	Andy	38.8%	1.6%	6.4%	53.2%	555
Family	Rodriguez	0.6%	0.5%	94.1%	4.8%	1,094,924
	Lee	43.8%	16.9%	2.0%	37.3%	693,023
	Washington	0.3%	91.6%	2.7%	5.4%	177,386

To test this, we began our analysis by manually validating name-based inference at three threshold ranges: 70-79%, 80-89%, and 90-100%. We sampled 300 authors from the WoS database, 25 randomly sampled for every combination of racial category and inference threshold. Two coders manually queried a search engine for the name and affiliation of each

author and attempted to infer a racial category through visual perception and information listed on their websites and CVs (e.g., affiliation with racialized organizations such as *Black in AI*, *SACNAS* etc.).

Figure 1 shows the number of valid and invalid inferences, as well as those for whom a category could not be manually identified, and those for whom no information was found. Asian authors were valid at every threshold considered. The inference of Black authors, in contrast, was invalid or uncertain at the 70-80% threshold, but high at the 90% threshold. Similarly, inferring Hispanic authors was only accurate after the 80% threshold. Inference of White authors was valid at all thresholds but improved above 90%. This suggests that a simple threshold-based approach is not equally valid across all races. We thereby move to a weighting-based scheme for analysis that takes does not provide an exclusive categorization but uses the full information of the distribution.

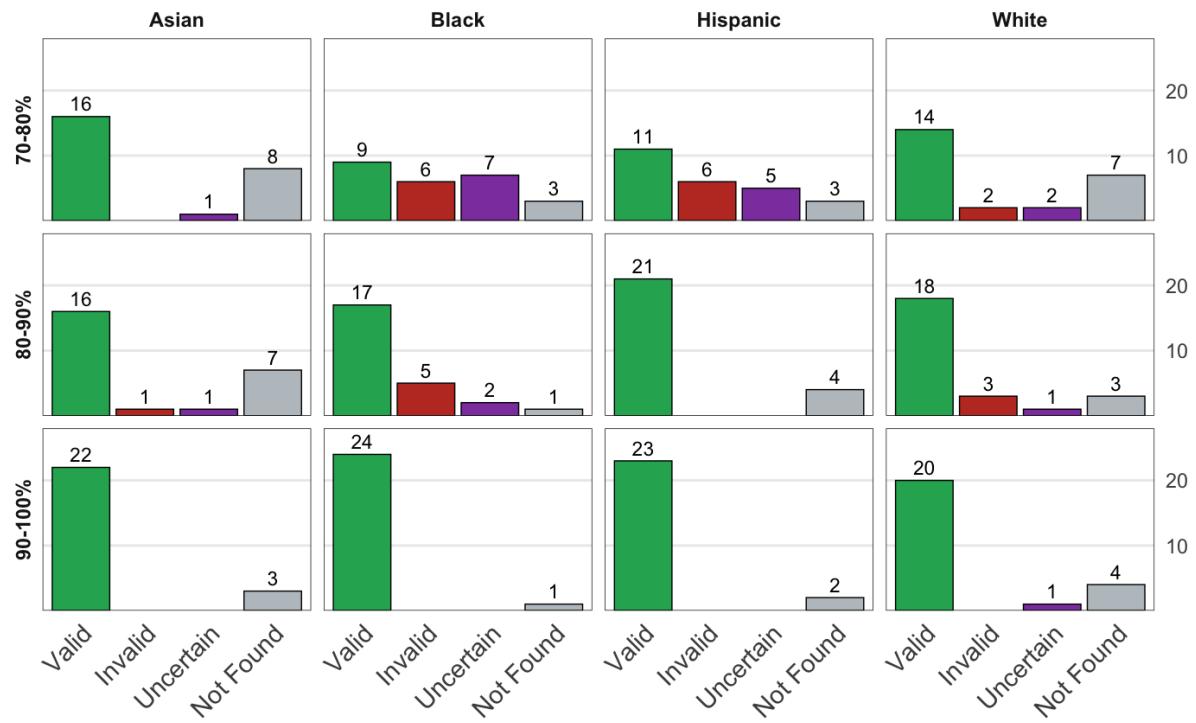


Figure 1. Manual validation of racial categories

Weighting scheme

We assess three strategies for inferring race from an author's name using a combination of their given and family name distributions across racial categories (Table 1). The first two aim to build a new distribution as a weighted average from both the given and family name racial distributions, and the third uses both distributions sequentially. In this section we explain these three approaches and compare them to alternatives that use only given or only family name racial distributions.

The weighting scheme should account for the intuition that if the given (family) name is highly informative while the family (given) name is not, the resulting average distribution should prioritize the information on the given (family) name distribution. For example, 94% of the people with Rodriguez as a family name identify themselves as Hispanic, whereas 39% of the people with the given name Andy identify as Asian, and 53% as White (see Table 1). For an author called Andy Rodriguez, we would like to build a distribution that encodes the informativeness of their family name, Rodriguez, rather than the relatively uninformative given

name, Andy. The first weighting scheme proposed is based on the standard deviation of the distribution:

$$sd = \sqrt{\frac{1}{N-1} \sum_{i=1}^N (x_i - \bar{x})^2}$$

With four racial categories, the standard deviation moves between 0, for perfect uniformity, and 0.5 when one category has a probability of 1. The second weighting scheme is based on entropy, a measure that is designed to capture the informativeness of a distribution:

$$e = - \sum_{i=1}^c P(x_i) \log P(x_i)$$

Using this, we propose the following weight for given names:

$$x_{weight} = \frac{f(x)^{exp}}{f(x)^{exp} + f(y)^{exp}}$$

with x and y as the given and family name. exp is the exponent applied to the function, and is a tuneable parameter. For the standard deviation, using the square function means we use the variance of the distribution. In general, the higher the exp is set, the more skewed the weighting is towards the most informative name distribution.

Figure 2 shows the weighting to the given and family names based on their informativeness, and for different values of the exponent. The horizontal and vertical axes show the highest value on the given and family name distribution, respectively. This means that a higher value on any axis corresponds with a more informative given/family name. The color palette shows how much weight is given to the given names. When the exponent is set to two, both the entropy and standard deviation-based models skew towards the most informative feature, a desirable property. Compared to other models, the variance gives the most extreme values to cases where only one name is informative, whereas the entropy-based model is the most uniform.

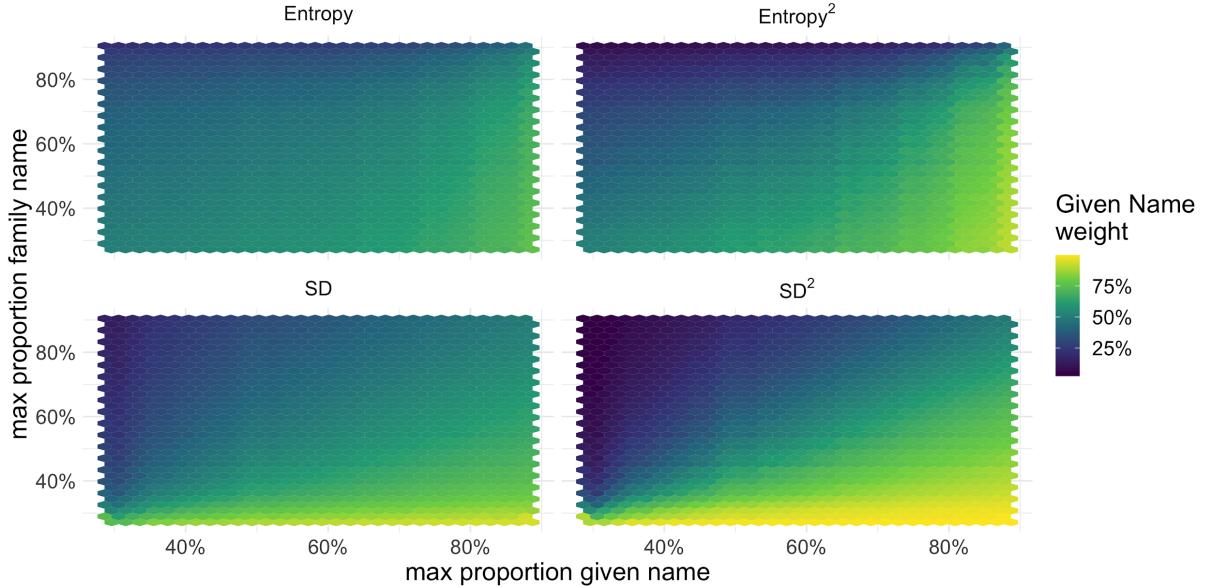


Figure 2. Given names weight distribution by given and family name skewness. Simulated data

Information retrieval

The above weighting schemes result in a single probability distribution of an author belonging to each of the racial categories, from which a race can be inferred. One strategy for inferring race from this distribution is to select the racial category above a certain threshold, if any. A second strategy is to make use of the full distribution to weight the author across different racial categories, rather than assigning any specific category. Third, we implement a two-step strategy that sequentially uses family and then given names. Given that family names are from census data, we first retrieve all authors which have a family name with a probability of belonging to a specific racial group greater than a given threshold. This retrieves N authors. Second, we retrieve the same amount, N , of authors as in the first step, but using the given names. Finally, we define a unique set of authors, removing those captured by both steps. This leaves a set between N and $2N$ authors, depending on the number of authors retrieved in both steps. Naturally, there are multiple possible changes on this two-step method, for example, a percentage threshold could be used in both steps, instead of a given number. It could be possible also to have a specific percentage threshold for each step. There is not a preferred method in principle, but the proposed configuration relates with our belief that the census data (family names) is of better quality, and therefore we prefer to use it to define the number observations to extract with the lower quality data.

Results

The effect of underlying skewness

Before comparing the results of the proposed strategies for using both given and family names, we present characteristics of these two distributions on the real data, and in relation to the WoS dataset. Table 2 shows the population distribution on the family names, based on the US census, and on the given names, based on the mortgage applications. Considering the US census data as ground truth, we see that the mortgage data highly over-represents the White population, particularly over-represents Asians, and underrepresents Black and Hispanic populations; this likely stems from the structural factors (i.e., economic inequality, redlining, etc.) that prevent marginalized groups from applying for mortgages in the US. People may also choose to self-report a different racial category when responding anonymously to the census bureau than when applying for a mortgage loan. Due to this bias in the distribution of given names, we normalize the given name distribution by applying an expansion factor to each racial group, so that the expanded given name count preserves the racial distribution in the US census.

Table 2. Racial representation of family names (US census) and given names (mortgage data)

<i>Racial group</i>	<i>Family names</i>	<i>Given names</i>
White	66.1%	82.6%
Hispanic	16.5%	6.9%
Black	12.4%	4.2%
Asian	5.0%	6.3%

Both given and family names share a characteristic not considered in our simulated data: the informativeness of names varies across racial groups. Inferring racial categories based on a set threshold will, then, produce biased results as typical names of one racial category are more informative, and thus more easily meet the threshold, than another. Figure 3 shows how the

representation of inferred races changes based on the assignment threshold used. Increasing the threshold results in fewer total individuals returned (top), as some names are not very sufficiently informative. For family names, only a small proportion of the population remains at a 90% threshold. The Asian population is highly over-represented between the 90% and 96% threshold, after which they suddenly become under-represented. The White population is systematically over-represented for any threshold, whereas the Black population is systematically underrepresented. The Hispanic population is overrepresented between the 65% and 92% threshold and underrepresented after. With given names, the White population is systematically overestimated for every threshold until 96%, where the Asian population is suddenly overestimated to a high degree. Similar results are observed based on given names. Again, the Asian population is highly overrepresented after the 96% threshold, whereas the White population is over-represented across nearly all thresholds and the Black and Hispanic population were underrepresented across all thresholds. The fact that Asian, and to some degree Hispanic, populations have more informative given and family names is related with higher differentiation of these two communities; White and Black populations in the United States, in contrast, tend to have more similar names (as verified in Elliott et al., 2009). Given that the White population is larger than the Black population, the use of a threshold (and assigning all people with that name to a single category), generates a Type I error on Black authors, and Type II error on White authors, thereby overestimating the proportion of White authors.

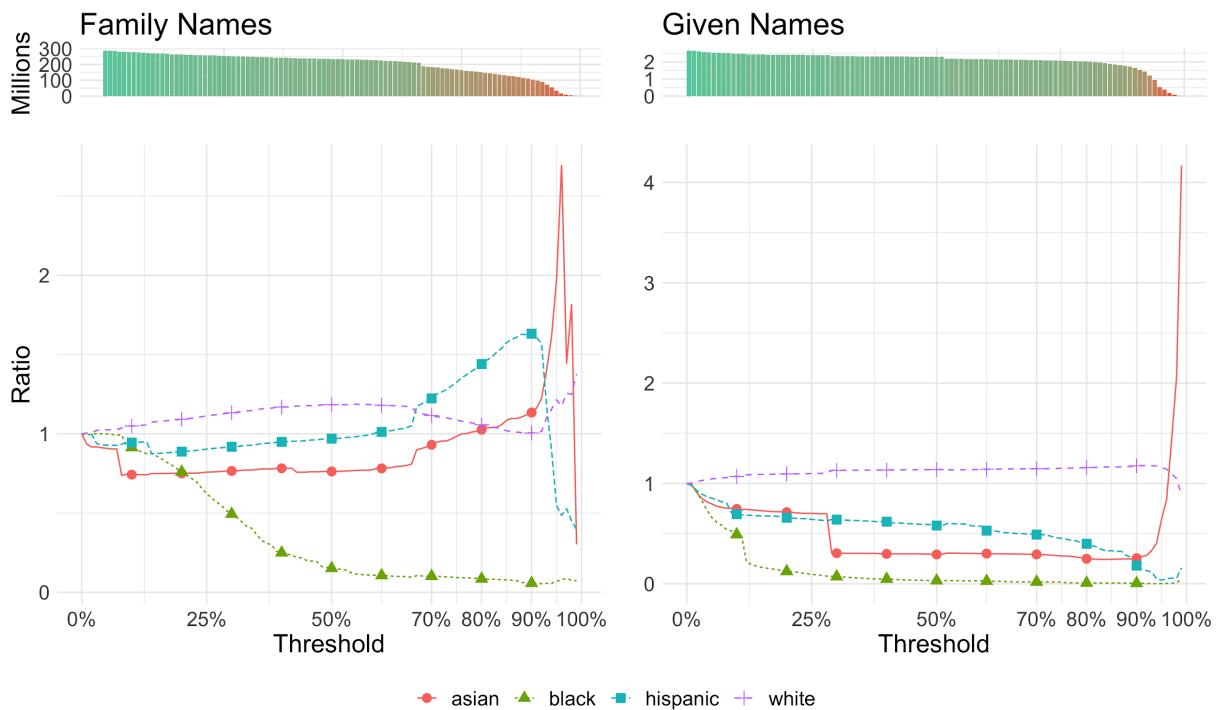


Figure 3. Changes in groups share, and people retrieved, by threshold. Census (Family names) and mortgage (Given names) datasets.

The effect of thresholding

Figure 4 shows the effect of using a 90% threshold on the WoS dataset of unique authors. The first column corresponds to each author counting fractionally towards each racial category in proportion to the probabilities of their name distribution, using family names from the census, i.e., this is the closest we can get to ground truths with the available information. The remaining columns represent (1) inference based on given and family names alone; (2) combination distributions; and (3) the two-steps strategy, using both normalized and unnormalized given names, always with a 90% threshold. All models highly underrepresented the Black population

of authors. All models except normalized given names under-represent the Hispanic population. The unnormalized given name, either alone or in the variance model, underrepresents the Asian population. Finally, the White population is overrepresented by all models except family names and the variance with normalized given names.

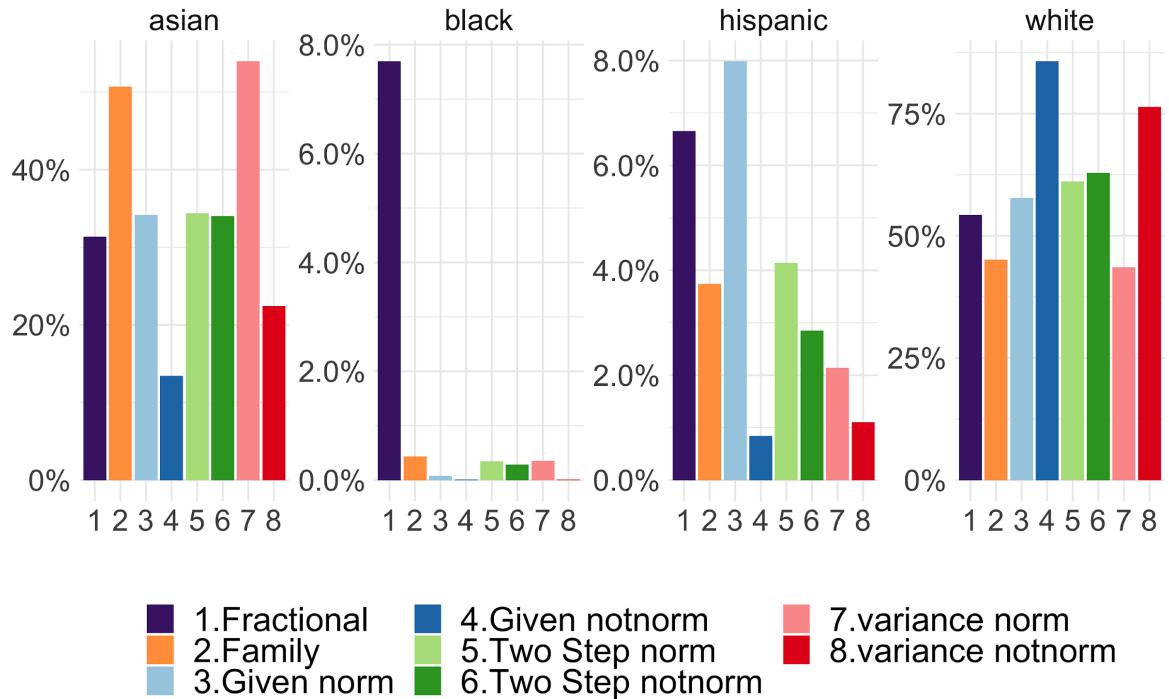


Figure 4. Resulting distribution on different models with 90% threshold. Fractional counting on family names for comparison.

Figure 5 demonstrates how the number of authors in each racial group drops when increasing the threshold. The convexity in the Black authors curve, distinct from the concave curves for other racial groups, explains the underrepresentation of this group under high thresholds. Even at less extreme thresholds, the proportion of Black authors in the dataset is far smaller than the fractionalized counting of White authors (851,843.60). The consequence of this is that a low threshold would grossly over represent White authors, while a high threshold would exclude nearly all Black authors. We conclude from this that a threshold-based approach, while intuitive and straightforward, should not be used for racial inference. Rather, analysis should be adapted to consider each author as a distribution over every racial category; in this way, even though an individual cannot be assigned into a category, aggregate results will remain unbiased.

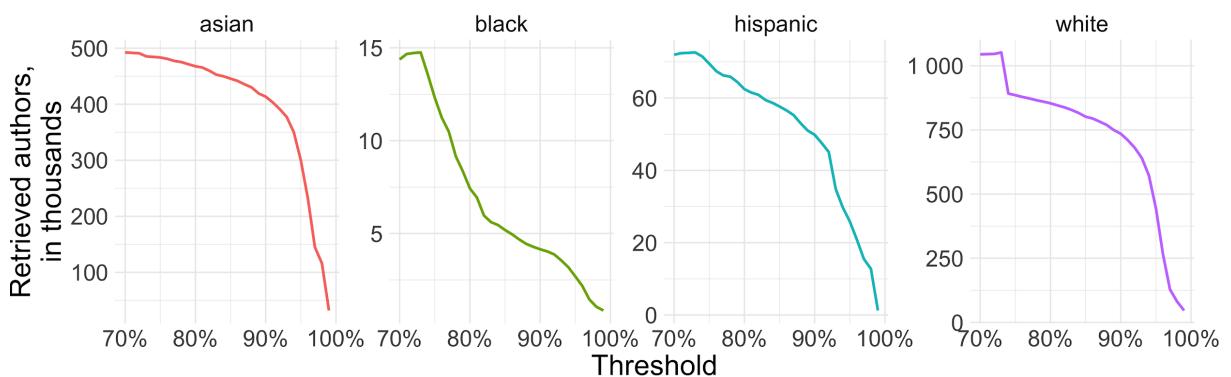


Figure 5. Two-steps strategy retrieval

The effect of imputation

Another consideration is how to deal with unknown names. As mentioned in the Data section, the family names dataset provided by the census bureau covers 90% of the US population. The remaining 10%, as well as author names not represented in the census, generates 774,381 articles, or 18.75% of the dataset, for which the family name of the first authors has an unknown distribution over racial categories.

An intuitive solution would be to impute missing names with a default distribution based on the racial composition of the entire census. Table 3 shows the distribution among racial groups in the US census and in WoS (for first authors with family names included in the US census data). The Asian population is highly overrepresented among WoS authors, whereas Hispanic and Black authors are highly underrepresented, with respect to the total US population. Imputing with the census-wide racial distribution is, therefore, equivalent to skewing the distribution towards Hispanic and Black authors and underrepresenting Asian authors. Since the ground truth is contingent to the specific dataset in use, the natural imputation would instead be the mean of the population most representative of an individual. For example, in the case of a missing author name in the WoS, the racial distribution of that individual's discipline could be imputed. Our recommendation is, therefore, to first compute the aggregate distribution of racial categories with the dataset in which the inference is intended, and then use this aggregate distribution to impute in those family names missing from the census dataset. Statistically, this preserves the aggregate distribution on this dataset.

Table 3. Racial distribution in US census and WOS US Authors with known family names.

<i>Racial group</i>	<i>US census</i>	<i>US WoS</i>
White	66.1%	59.4%
Hispanic	16.5%	5.4%
Black	12.4%	7.2%
Asian	5.0%	24.5%

Conclusion

Race scholars (Emirbayer & Desmond, 2011) have advocated for a renewal of Bourdieu's (2001) call for reflexivity in science of science. We pursue this through empirical reflexivity: challenging the instrumentation used to collect and code data for large-scale race analysis. In this paper we manually validate and propose several approaches for name-based racial inference of US authors. We demonstrated the behaviour of the different methods on simulated data, across the population, and on authors in the WoS database. We showed the risks of underestimating highly minoritized groups (e.g., Black authors) in the data when using a threshold, and the overestimation of White authors introduced by given names when they are based on mortgage data. A similar result was identified by Cook (2014), in her attempt to infer race of patent data based on the US census: she found that the approach "significantly underpredicted matches to black inventors and overpredicted matches to white inventors" and concludes that the name-based inference approach was not suitable for historical analyses.

From our analysis, we come away with three major lessons that are generally applicable to the use of name-based inference of race in the US, shown in table 4.

Table 4. General recommendations for implementing a name-based inference of race.

	<i>Do's</i>	<i>Don'ts</i>
<i>Given Names</i>	Use only family names from US census to avoid bias.	Do not use given names, except the underlying distribution of your dataset matches that of mortgage data.
<i>Thresholding</i>	Consider each person in your data as a distribution and adapt your summary statistics.	Do not use a threshold for categorical classification of each person, as this underrepresents Black population, due to the correlation between racial groups and name informativeness.
<i>Imputation</i>	Calculate first the aggregated distribution on your dataset, imputation, except your target population and use this for imputation of missing cases.	Do not use the census aggregate distribution for matches the US population.

Inferring race based on name is an imperfect, but often necessary approach to studying inequities and prejudice in bibliometric data (e.g., Freeman & Huang, 2014), and in other areas where self-reported race is not provided. However, the lessons shown here demonstrate that care must be taken when making such inferences in order to avoid bias in our datasets and studies.

It has been argued that science and technology serve as regressive factors in the economy, by reinforcing and exacerbating inequality (Bozeman, 2020). As Bozeman (2020) argued, “it is time to rethink the economic equation justifying government support for science not just in terms of why and how much, but also in terms of who.” Studies of the scientific workforce that examine race are essential for identifying who is contributing to science and how those contributions change the portfolio of what is known. To do this at scale requires algorithmic approaches; however, using biased instruments to study bias often replicates the very inequities they hope to address.

In this study, we attempt to problematize the use of race from a methodological and variable operationalization perspective in the US context. However, any extension of this work across country lines will necessarily require tailoring to meet the unique contextual needs of the country or region in question. Ultimately, scientometric researchers utilizing race data are responsible for preserving the integrity of their inferences by situating their interpretations within the broader socio-historical context of the people, place, and publications under investigation. In this way, they can avoid preserving unequal systems of race stratification and instead contribute to the rigorous examination of race and science intersections toward a better understanding of the science of science as a discipline. As Zuberi (2001) remarked: “The racialization of data is an artifact of both the struggles to preserve and to destroy racial stratification.”

References

Bozeman, B. (2020). Public Value Science. *Issues in Science and Technology*, 34-41.

Bourdieu, P. (2001). *Science of Science and Reflexivity*. Chicago, IL: University of Chicago Press.

Buolamwini, J., & Gebru, T. (2018, January). Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency* (pp. 77-91). PMLR.

Caliskan, A., Bryson, J. J., & Narayanan, A. (2017). Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334), 183-186.

Cook, L.D. (2014). Violence and economic activity: evidence from African American patents, 1870-1940. *Journal of Economic Growth*, 19, 221-257.

D'Ignazio, C., & Klein, L. F. (2020). *Data feminism*. MIT Press.

Elliott, M.N., Morrison, P.A., Fremont, A., McCaffrey, D.F., Pantoja, P., & Lurie, N. (2009). Using the Census Bureau's surname list to improve estimates of race/ethnicity and associated disparities. *Health Serv Outcomes Res Method*, 9, 69-83.

Emrbayer, M., & Desmond, M. (2011). Race and reflexivity. *Ethnic and Racial Studies*, 35(4), 574-599.

Fiscella, K., & Fremont, A.M. (2006). Use of geocoding and surname analysis to estimate race and ethnicity. *Health Services Research*, 41(1), 1482-1500.

Freeman, R.B., & Huang, W. (2014). Collaborating with people like me: Ethnic co-authorship within the US. NBER working paper 19905.

Galton, F. (1891). *Hereditary genius*. D. Appleton.

Ginther, D.K., Basner, J., Jensen, U., Schnell, J., Kington, R., & Schaffer, W.T. (2018). Publications as predictors of racial and ethnic differences in NIH research awards. *PLOS ONE*, 13(11), e0205929.

Ginther, D.K., Schaffer, W.T., Schnell, J., Masimore, B., Liu, F., Haak, L.L., & Kington, R. (2011). Race, ethnicity, and NIH research awards. *Science*, 333(6045), 1015-1019.

Godin, B. (2007). From eugenics to scientometrics: Galton, Cattell, and men of science. *Social studies of science*, 37(5), 691-728.

Hopkins, A.L., Jawitz, J.W., McCarty, C., Goldman, A., & Basu, N.B. (2013). Disparities in publication patterns by gender, race and ethnicity based on a survey of a random sample of authors. *Scientometrics*, 96, 515-534.

Hoppe, T.A., Litovitz, A., Willis, K.A., Meseroll, R.A., Perkins, M.J., Hutchins, B.A., Davis, A.F., Lauer, M.S., Valantine, H.A., Anderson, J.M., & Santangelo, G.M. (2019). Topic choice contributes to the lower rate of NIH awards to African-American/black scientists. *Science Advances*, 5: eea7238.

Horton, H. D. (1998). Toward a critical demography of race and ethnicity: Introduction of the “R” word.

Horton, H. D., & Sykes, L. L. (2001). Reconsidering wealth, status, and power: Critical Demography and the measurement of racism. *Race and society*, 4(2), 207-217.

Horton, H. D. (2002). Rethinking American diversity: Conceptual and theoretical challenges for racial and ethnic demography. In N. Denton & S. Tolnay (Eds.), *American diversity: A Demographic challenge for the twenty-first century* (p. 261-278). New York: State University of New York Press.

Humes, K., Jones, N. , Ramirez, R. (2011). Overview of Race and Hispanic Origin: 2010. U.S. https://10.4.51awww.atlantic.org/images/publications/Democratic_Defense_Against_Disinformation_FINAL.pdf

Lariviere, V., Ni, C., Gingras, Y., Cronin, B., & Sugimoto, C.R. (2013). Global gender disparities in science. *Nature*, 504, 211-213.

Liebler, C.A., Prter, S.R., Fernandez, L.E., Noon, J.M., & Ennis, S.R. (2017). *Demography*, 54(1), 259-284.

Locke, G., Blank, R., Groves, R. (2011). 2010 Census Redistricting Data (Public Law 94-171) Summary File. <https://www.census.gov/prod/cen2010/doc/pl94-171.pdf>

Prescod-Weinstein, C. (2020). Making Black women scientists under white empiricism: the racialization of epistemology in physics. *Signs: Journal of Women in Culture and Society*, 45(2), 421-447.

Stevens, K. R., Masters, K. S., Imoukhuede, P. I., Haynes, K. A., Setton, L. A., Cosgriff-Hernandez, E., ... & Eniola-Adefeso, O. (2021). Fund Black scientists. *Cell*.

Teh, Y. W. (2010). Dirichlet Process. <https://www.stats.ox.ac.uk/~teh/research/npbayes/Teh2010a.pdf>

Tzioumis, K. (2018). Demographic aspects of first names. *Scientific data*, 5, 180025.

U.S. Bureau of the Census. (1975). Historical Statistics of the United States, Colonial Times to 1970, Bicentennial Edition, Part 1). <https://www.census.gov/history/pdf/histstats-colonial-1970.pdf>

US Census Bureau. (2016). Frequently Occurring Surnames from the 2010 Census. The United States Census Bureau. https://www.census.gov/topics/population/genealogy/data/2010_surnames.html

Wilkerson, I. (2020). Caste: The Origins of Our Discontents. Random House.

Zuberi, T. (2001). Thicker than blood: How racial statistics lie. U of Minnesota Press.