

Representing Race and Ethnicity in American Fiction, 1789-1920

Mark Algee-Hewitt, J.D. Porter, Hannah Walser

Mark Algee-Hewitt, Stanford University
J.D. Porter, The University of Pennsylvania
Hannah Walser, Harvard University
Peer-Reviewers: Hsuan Hsu, Ryan Cordell
Dataverse DOI: 10.7910/DVN/P2RD03

ABSTRACT

Our project, which aims to reconstruct racial discourse in American literature, tracks three critical aspects of the representation of race and ethnicity in a corpus of over 18,000 American novels published between 1789 and 1920. First, we provide a historically sensitive account of the ethnicities that most occupied the nation's racial imaginary, registering how different ethnic groups were perceived to be biologically, geographically, or socially linked. Second, we track the descriptive terms most associated with particular ethnicities over time as we trace the changing discursive fields surrounding particular racial groups. Finally, we explore the coherence of the discourse around each race and ethnicity represented across American literature before 1920, paying close attention to the ways in which various groups did or did not exist as semantically unified groups at specific historical moments. Taken together, our three questions show not just who was under discussion and how, but also the history—and historicity—of racialization and ethnic thinking writ large.

Our goal in this paper is to identify and surface the racialized language of American Fiction and to face the harms that it caused without eliding its historical violence and force. At the same time, while we feel that confronting such racism is important work, we do not want to perpetuate the harm that this language, including many slurs, continues to cause to oppressed peoples, particularly in the Black and Native American communities. To that end, throughout this paper, we have adopted the practice of Brigitte Fielder, among others, in representing particularly harmful terms using the following convention: n[-----].¹

Introduction

In many disciplines across the sciences and humanities, the “folk conception” of race has become an important foil for the more accurate understanding of race as a cultural construct. While scholars understand that racial categories have an arbitrary, historically determined form with no basis in

biological reality, in the popular imagination race often retains the authority and immutability of established science. Yet within the academy this folk conception is often treated with the same casual, unsubstantiated confidence that is supposed to characterize it. From the simplest issues (how many races do people think there are?) to the most complex (how does biogeographical ancestry map onto racial categories?), the presumed nature of the popular understanding of race varies from one study to the next—whether because of actual variation across cultures and eras, different research contexts, or lack of precision on the part of the researchers. The pressing question is: What *do* the folk think about race?²

The methods of cultural analytics are especially well-suited for developing an answer to this question. On the one hand, literary scholarship has developed a robust set of tools for analyzing something like a “folk conception” as it appears in cultural artifacts—as evident in the many forays into the cultural creation of race undertaken by critics ranging from Henry Louis Gates and Toni Morrison to Michael Hames-García and Robin Bernstein. On the other hand, to suggest that American fiction has been a particularly active site for racial redefinitions does not necessarily imply that authors were engaged consciously or deliberately in this activity — nor that a novel must explicitly focus on race in order to contribute to a widespread racial discourse. One particular benefit of computationally analyzing a corpus of thousands of texts is the amplification of racial “signals” that might, in an individual text and to an individual reader, be so faint as to escape notice. Recent scholarly work and anti-racist activism has shifted both empirical research agendas and political conversations away from individual prejudice or conscious bias and toward structural inequities, emphasizing the degree to which racism “lives” not in hearts or minds but in brains, bodies, and institutions. Those who study identity are alert to the need to name and challenge dominant narratives surrounding race, even (or especially) when those narratives are actually instantiated only in partial, tacit, or compromised form.

For the study of representations of identity in fiction, quantitative textual analysis offers a unique opportunity to access a kind of discursive unconscious — the background of associative biases against which any individual author constructs his characters. Our particular method of statistical analysis, which records significant collocates of ethnic and racial target terms, mimics on the level of language the implicit associations that social psychologists use to identify unconscious biases and predispositions in their subjects, determining which words are likely to show up near each other even when they are not necessarily deliberately linked.³ Insofar as our statistical analysis remains insensitive to

authorial intention, it represents, on one level, a loss of complexity—but it also makes possible a shift in scale that allows us to identify patterns of racial discourse too diffuse to be perceptible. This process not only gives us a better sense of the contours and content of “the racial unconscious” (in Eric Lott’s phrase) across a wide range of ethnic groups; it also provides crucial empirical evidence and historical depth for claims of bias in representation, revealing racism’s place in the *longue durée* of American attitudes. As in the Freudian unconscious, racial associations here assume a kind of absolute value that disregards negation: the statistical link between the words “negro” and its most negative valences for instance, which can be found in table 1 below, can be bolstered both by novels that depict Black characters in these ways and by works (like, for instance, *To Kill a Mockingbird*) that raise these associations in an ostensible attempt to undercut them. As in the “cognitive nonconscious” sketched by Katherine Hayles (2017), on the other hand, these associations represent not so much affective attitudes elaborated over the life of an individual, but epistemes emergent from unguided (but deeply unequal) systems. To the idea of unconscious racial cognition, then, we would add something like Michael Omi and Howard Winant’s concept of racial “common sense”: “a way of comprehending, explaining, and acting in the world” that, through constant “racial projects” that differentially distribute meaning and resources, constructs race as something obvious, visible, and determinative.⁴ The product of constant collective effort, but often experienced by the individual as effortless and ineluctable, this common sense would emerge, we wagered, as a kind of background noise in the nineteenth-century American novel that our statistical methods could amplify.

In this paper, we cash in on this promise of the literary digital humanities to reveal this background through an analysis of racial and ethnic language in about 18,000 novels published in the United States from 1789 to 1920. Drawn from the Gale American Fiction collection, which is based on scholarly bibliographies of that period, these novels represent the vast majority of all extant prose fiction published in the country during the first half of its history. Beginning with a list of racial and ethnic terms that we created, we track three things in this corpus: 1) The frequency of those terms over time, 2) The words that tend to show up near those terms, or their “collocates”, and 3) The coherence of the discourse around each race and ethnicity represented by the terms. The first shows us which races and ethnicities were under discussion during the long nineteenth century in America; the second shows us how they were discussed. The last is more

complicated, and leads to the major finding of this paper. Initially, we wanted to know which racial and ethnic groups would be discussed most consistently over time; for instance, would conversations about Irish Americans look the same in 1810 and 1910? Would stereotypes about Chinese immigrants change as Chinese immigration itself increased? Would the words used to describe African Americans undergo a massive change after the Civil War? In the course of testing this concept, we reached a fundamentally different and more radical conclusion. Rather than finding that the words surrounding groups of people changed or didn't change over time, we found that the groups only cohered *as groups* at specific historical moments. In other words, many of the ethnicities we set out to find did not exist as semantically unified groups throughout literary history; in a sense, the number of races and ethnicities varied over time. Our chief historical finding, then, is less about the discourse of, say, Irish Americans and more about whether "Irish Americans" existed as a salient category at all. Taken together, our three questions show not just who was under discussion and how, but also the history—and historicity—of racialization and ethnic thinking writ large.

Results, Part 1: The Racial Field

Our quantitative approach to identifying the background discourse of race in American fiction rests on our ability to seed our model with terms based on our own scholarly understanding of this discourse. This allows us to statistically situate our quantitative approach within the critical work on racial discourse in America. Our initial step, therefore, was to create a list of fourteen racial and ethnic categories and then populate each category with lists of relevant words, or "target terms." In both cases, we proceeded subjectively, albeit with the help of our training in American history. The categories include broad racial/ethnic groups (black, white, Native American), ethno-religious distinctions (Catholic, Jewish, Middle Eastern and Muslim), geographical origins (East Asian, Eastern European, German/Dutch, Irish, Italian, Latin American, Scandinavian), and one catch-all field (Immigrants). Within each category we hand-selected target terms associated with that group in historic American discourse. These terms range from common terms like "Indian" or "migrant" to outdated technical terminology like "octoroon" or "Mohammedan" to specific nationalities or ancestries like "Sioux" or "Brazilian" to slurs. In all, we wound up with 208 terms spread unevenly across the fourteen categories. This approach has its disadvantages.

Most notably, it is not comprehensive, but limited by our own specialized knowledge. To take just two examples, in subsequent conversations with other researchers in race and ethnicity studies, we have come to believe that we should have included terms for Pacific Islanders and Filipinos; eliding those groups is an important and unjustified oversight. Nevertheless, our initial goal was not comprehensiveness—which is probably impossible—but a large-scale study of race and ethnicity. Even given its important limitations, our approach successfully established a terminology of race and ethnicity from which we could begin to analyze a broader discourse. This goal also informs another apparent weakness of our list: its ahistoricism. Our process did not differentiate *a priori* between terms in use today and terms that seem to have dropped out of the language in the early nineteenth century, and included terms that were unlikely to appear in any great numbers in our corpus’s historical range—but this ahistorical approach was necessary if we were to confirm or disconfirm hypotheses about historical change. If a term that we thought of as taking hold in the 1950s turned out to have an unexpected spike in the 1830s, we did not want to miss it.

For our literary texts, we turned as mentioned to the Gale American Fiction Corpus, a collection of 18,101 novels covering the period from 1789 to 1920. Because it is based on two bibliographies (Lyle Wright’s *American Fiction* and Geoffrey Smith’s *American Fiction: 1901-1925*), it is exceptionally reliable, and covers nearly every work of long prose fiction published in the United States during the periods it represents. Like any corpus, it also leaves some things out. As the “long prose fiction” caveat suggests, the Gale corpus is composed of novels or short story collections that were published as freestanding volumes; as such, it doesn’t directly register the large quantity of fiction being published in magazines during the long nineteenth century. This omission is particularly significant when it comes to nineteenth-century authors of color, who were more likely to publish in magazines than to be picked up by the overwhelmingly white editors at major publishing houses.⁵ Martin Delany’s radical novel *Blake; or, The Huts of America*, for instance, does not appear in the Gale corpus at all, despite its literary and historical significance: *Blake* was serialized in *The Anglo-African Magazine* and *The Weekly Anglo-African* between 1859 and 1862, but was not published in book form until more than a century later, in 1970 (Delany 2017).⁶ This means that our corpus is significantly whiter than the body of *all* fiction published in the long nineteenth century, since it shares its bias toward white writers with the publishing industry itself. We hope that future research will add

depth and precision to our results by comparing them with corpora composed mainly of authors of color—but we also believe there is value in tracking the folk concept of race in the works of predominantly white writers, since racist ideology is primarily a white construction (at least in the United States). Precisely because white writers have had claim to cultural hegemony, their representations of race are, we wager, more likely to unselfconsciously reflect and indeed constitute the constant background noise of racism.




Figure 1 Frequency (per 100,000 words) of terms describing racial, ethnic and ancestral groups in 25-year slices of the Gale American Fiction corpus

The first layer of results shows the frequency of each of our terms over time.⁷ In Figure 1, which reflects term frequency (scaled per 100,000 words) within the Gale corpus, these numbers are shown at the level of the broad categories. A few historical trajectories are immediately clear: terms associated with Native Americans dominate racial and ethnic discourse in the first few periods, before being matched and then surpassed by the “Black” category; “East Asian” terms grow more and more frequent over the whole period; “Irish” stays consistent throughout. Individual terms (Figure 2) provide a high-level window into what is going on beneath the surface of the broader categories: certain words, for instance (“native”, “Indian”, “chief”), dominate the discourse of Native

Americanness.⁸ These figures show one kind of racial and ethnic history: which people were being discussed, specifically in literature. It does not show which people were living in the United States; the convergences and divergences of these two kinds of history show how racial and ethnic discourse responds to, resists, and reshapes race and ethnicity.



Figure 2 Frequency (per 100,000 words) of individual terms in the Gale American Fiction corpus

At times the relationship is fairly straightforward, a window into literature that responded to changing facts about the world. For instance, as Figure 3 shows, the persistent presence of “Native American” as a broad category masks substantial changes within the category. In the first period, the three most-mentioned nations are the Cherokee, Mohegan, and Choctaw—all of whom lived (prior to Andrew Jackson’s Indian Removal Acts) east of the boundaries of the Louisiana Purchase. The second period reflects a similar Eastern bias, with nations like the Iroquois and Lenape (both largely Northeastern) in the top five; at the same time, the Sioux reflect a slight move to the Northwest. By the latest periods of the chart, nations like the Apache and Navajo have risen in the rankings, corresponding with United States encroachment on their territory in the

Southwest. Some of these terms are probably present for reasons having to do with genre, especially the development of the Western toward the later periods. Yet even that generic evolution appears to reflect the history of encounter; there is no Western before the U.S. invades the West.




Figure 3 Frequency (per 100,000 words) of mentions of Native American nations

Even more striking is the substantial uptick in the 1815-1839 period across all but two of the nations (the more western Kiowa and Navajo). These changes are driven by relatively few mentions—just 159 for the word “cherokee”, which tops this list for the period—because there were relatively few books being published in the United States at all in this period; in the Gale corpus there are only 820 *total* works published before 1840. One side effect of this graph, then, is to highlight the importance of one writer who was quite prolific at the time: James Fenimore Cooper, who alone is responsible for a third of the mentions of all the nation words in this graph in the period, with substantial percentages of five of them (“iroquois”, “lenape”, “mohegan”, “mohican”, and “sioux”) including, unsurprisingly, more than 90% of the mentions of “mohican”. Here too we see an interaction of the literary and the political that accords with Jill Lepore’s argument that literature had considerable agency in the contemporaneous process of Indian Removal: “While ... Americans everywhere read Cooper’s *Leatherstocking Tales* by the fireside, the federal government

sought support for removing eastern Indians west of the Mississippi partly by invoking images like those popularized in Indian plays and Indian fiction.”⁹ When it comes to Native Americans, there is good reason to believe that America’s literary imagination was deeply entangled with its imperial march west.

Analysis of our other categories shows a different pattern. The census has detailed data about the national origins of the American population from 1850-1930 (see Figure 4).¹⁰ In some cases, this data does appear to correspond




Figure 4: U.S. Population Claiming National Origin in Countries Belonging to Our Categories (source: U.S. Census)

with literary presence; Scandinavians, for instance, appear in American towns at a fairly similar rate to their appearance in American novels. But this is often not the case. German and Dutch respondents are on the wane heading into 1920, especially as a percentage of all foreign-born Americans; yet their category grows over the same period in our data. This could simply reflect a lag—perhaps novels only reflect new populations after they have been around for a few decades—but it fits in with a general disconnect that is particularly pronounced when comparing categories on a relative basis. This is most evident with the East Asian category, which triples in term frequency from the first period to the last. Immigration from China and Japan really did grow tremendously over this time span, even after the severe limitations imposed by the Chinese Exclusion Act.¹¹ Yet it remained a fairly small fraction of all immigration; it is dwarfed, for instance, by Eastern European and Italian immigration in the last two periods. In the literature, the Eastern European category grows a little, and the Italian category stays about the same, but the

East Asian category easily outpaces both. The literary discourse simply does not reflect racial and ethnic reality in any straightforward demographic sense. This will seem intuitively correct to anyone familiar with the toxic “Yellow Peril” rhetoric that proved so pervasive in American politics through the latter half of the 19th century and first half of the twentieth, even as immigration from the regions at the center of the imaginary peril was severely curtailed. In moments like these, when the discourse diverges from the demographics, we have a window into the operations of race and ethnicity as culturally determined categories; literature does not simply respond to the world, because it is too busy helping to create it.

Results, Part 2: The Racial Unconscious

While the simple appearance of our set of terms in this corpus is illustrative of the *presence* of racial or ethnic discourse in American Fiction, it does not reflect the interrelationships, the points of contact and divergence, that represent the evolution of this discourse over time. After all, we also seek to uncover the historically contingent semantics that attach to various descriptors of socio-cultural identity: how configurations of words attach to specific identities, and how both these configurations and the identities they describe alter with time and in response to socio-cultural shifts. Beyond simply registering how present a word is at any historical point (which the frequency graphs above describe in detail), the relationships between words (and between words and ideas) undergird an understanding of the discourse of race and ethnicity that is both diachronic (registering change over time) and inclusive of each word’s full range of meaning. Depending on the relationships that we seek, two options are available for such a quantitative analysis of language as we propose here.

Methods based on word embeddings¹² use a modeling process (for example, neural networks or least squares) to represent each word in a corpus as a vector of arbitrary length that can be related by distance metrics (such as cosine similarity) to other words represented by equivalent vectors. While there are many advantages associated with this method (including the ability to add and subtract vectors to achieve a more complex representation of word relationships), there are two significant drawbacks for our project. First, as these models relate words based on shared context, similarity is weighted more heavily on substitutability rather than proximity. That is, in a *gloVe* model of a corpus such

as ours, the closest terms to “jew” (a descriptor of both an ethnic identity, a practitioner of a religion and, at times, a racial epithet) are: “gentile”, “turk”, “jews”, “priest”, “peddler”, “jewish”, “trader”, “merchant”, and “dealer”. The top term, “gentile” is the antonym of “jew”, highly placed because it shares many of the same descriptive contexts. The list also includes plurals and adjectives for the word, as well as other near Eastern ethnic groups, all of which share similar contexts. While vector math may be able to assist us in disambiguating antonyms, grammatical forms, and geographically proximate groups, the specific operations required would always privilege one set of connections over another (for example, subtracting “religion” from “jew” gives a list of secular professions, and adding “race” gives a list of other minority group names used in similar contexts). Secondly, and more importantly, we are unable to reconstruct the specific logic behind any of these groupings, beyond the general fact of shared context. If we seek to reconstruct not just which words were associated but why they were associated, we need to be able to examine their contexts in detail. A word embedding model can suggest similarities, but only as the result of aggregate contexts, making it a blunt instrument with respect to historical differentiation and the nuance with which these terms were deployed by authors.

For this project, then, we elected to examine the relationships between our terms through collocate analysis.¹³ The relationship of a word to a collocate is proximal, and by comparing the frequency of each word as a collocate of our target terms (here, words within a 10-word horizon before or after each target) to the frequency of the word overall in the corpus, the significance of each collocate can be calculated. Through this method, two words are related if they appear within the horizon significantly more than can be explained by chance occurrence based on either word’s frequency throughout the corpus.¹⁴ Following the above example, the 10 closest words to “jew” are “renegade”, “gentile”, “jewing”, “gaberline”, “spindler”, “maimon”, “scythian”, “confucian”, “herodian”, “monish”. Again, “gentile” is high, but now surrounded by words speaking to religious apostasy and the origins of the Jewish people in the ancient Middle East. Rather than the plural, it reveals the verb “jewing” that uses the identity as a slur. Not only does this list better capture the complexity of associations linked with the identity itself, but the principle of connection—they are all words that appear in the immediate vicinity of “jew”—better enables us to trace the meaning of each of these terms back into the texts that established it, capturing both associative, as well as syntagmatic, relationships. Put simply, it allows us to return to reading—to interpreting these associations as they appear in natural language—

in a way that word embedding models do not. For these reasons, we chose collocate analysis to trace the associations we sought.¹⁵

As a final stage in our analysis, we filtered the resulting collocates through the Oxford English Dictionary wordlist, retaining only words that appeared in that list (while excluding stop words). Our choice meant that we would lose many character names from our analysis; however, given the relatively poor OCR quality of the Gale corpus, we felt that this step was a critical means to readability in the results. In the Gale corpus, this process yielded 26,976 unique distinctive collocates.¹⁶ At the length of a novella (Hemingway's *The Old Man and the Sea* is only 26,600 words, of which about 2,500 are unique), this list proved unmanageable to parse on a term by term basis. Even when filtered for the *most* significant collocates (those with an observed/expected of greater than 100, or those that appeared more than 10 times overall, which is the list that the below tables are drawn from), there were still more terms than were comprehensible on the human scale of reading. After studying some examples of the most striking collocates, then, we elected to use a set of summary statistics to assist us in our analysis, on the one hand using the collocates as a means of assessing the connections (or similarities) of our target terms and, on the other hand, using these collocates as a basis for ascertaining the overall clustering of each target race term across our periods, as we will describe in our results.

It is worth taking a moment to recognize a few conceptual limitations of this data. Early on, we were forced to acknowledge a surprisingly difficult problem: what to do about the words “black” and “white”. Both are essential to the discourse of their racial groups; both are also ordinary color terms used in a wide variety of non-racial circumstances. Indeed, our analysis of target term collocation (see below) indicated that the term “white” mostly did not reflect discourse about white people. This corresponds with a disproportionate absence of whiteness markers in general: a huge majority of the characters in our corpora are white, but their whiteness is simply assumed; they are racially invisible to their authors and each other. (See the end of this section for a more lengthy discussion of the circumstances in which whiteness *does* become visible in our corpus.) There was no clear right answer to this dilemma, not least because color words can signify race even when they aren't modifying human characters, as Toni Morrison points out in *Playing in the Dark*—in one of the most powerful moments of racial discourse in 19th-century fiction, “white” is a collocate mostly of “fog”—and as segregationists in Alabama proved when they raised an uproar against the 1958 children's book *The Rabbits' Wedding* because of the fear that

its black and white bunnies were tacitly promoting interracial marriage.¹⁷ In spite of these important caveats, in the end we elected to remove the words “black” and “white,” which seemed based on our collocate list to be more indicative of literal color than of race.¹⁸ Moreover, because we did not grammatically parse our corpus, we are unable to definitively determine the syntactic role typically played by each collocate—a key shortcoming, given that many of our ethnic terms can be used as nouns or as adjectives (for instance, “indian,” “mexican,” or “catholic”). Computationally, then, we did not distinguish between a Chinese child and a Chinese vase and “the Chinese” in the abstract: at most, a look our collocates and target terms in context can give us hints as to whether particular target terms are modifying characters, modifying objects, serving as individual or collective nouns, and so on. Insofar as this minimalism drew our attention to persistent slippages between persons, objects, and cultural abstractions, however, the ambiguity it produced proved productive—and inspired further investigation into the way that different ethnic terms connote animacy or inanimacy.

Despite these weaknesses, significant collocates gave us a useful window onto the changing contexts in which the target terms were used, as well as the different valences of particular terms within each group. For instance, within the apparently monolithic racial category “Black”, closer analysis of collocates across our historical range reveals very distinct discursive registers and sociopolitical agendas from one term to the next. Perhaps most immediately striking are the different groups of collocates associated with the terms “negro” and “n[-----]”.

| Term | NObs | Obs_Exp | p-value |
|------------------|------|------------|------------|
| rapist | 11 | 869.94785 | 0.00048826 |
| disfranchisement | 19 | 546.413526 | 1.907E-06 |
| stealer | 21 | 244.750878 | 4.7674E-07 |
| lobo | 22 | 206.428303 | 2.3837E-07 |
| roustabouts | 22 | 152.240874 | 2.3837E-07 |
| insurrections | 15 | 151.672103 | 3.0515E-05 |
| mulattoes | 22 | 144.561067 | 2.3837E-07 |
| enfranchised | 19 | 134.852056 | 1.907E-06 |
| m[----] | 24 | 133.532425 | 5.9589E-08 |
| minstrels | 77 | 128.10652 | 2.607E-22 |
| dialect | 274 | 127.979136 | 3.0423E-79 |
| enfranchisement | 15 | 117.78791 | 3.0515E-05 |
| cabins | 394 | 114.483481 | 3.01E-113 |
| mammies | 12 | 113.076394 | 0.00024413 |
| carpetbagger | 23 | 111.447467 | 1.1918E-07 |

| | | | |
|--------------|----|------------|------------|
| stevedores | 16 | 108.350469 | 1.5257E-05 |
| scalawag | 19 | 103.886028 | 1.907E-06 |
| amalgamation | 21 | 101.534207 | 4.7674E-07 |

Table 1: Filtered collocates of “negro”

| Term | NObs | Obs_Exp | p-value |
|-----------|------|------------|------------|
| bluegum | 18 | 1833.78847 | 3.81E-06 |
| buckra | 128 | 1489.37135 | 2.89E-39 |
| heaper | 23 | 727.19198 | 1.19E-07 |
| cust | 19 | 558.365079 | 1.91E-06 |
| n[-----] | 2314 | 459.28073 | 0 |
| gentman | 20 | 325.1398 | 9.53E-07 |
| trooly | 14 | 258.800792 | 6.10E-05 |
| wut | 39 | 250.763501 | 1.82E-12 |
| onery | 14 | 239.4873 | 6.10E-05 |
| wenches | 26 | 211.34087 | 1.49E-08 |
| woodpile | 88 | 210.670216 | 3.21E-27 |
| paisley | 26 | 180.327157 | 1.49E-08 |
| wich | 290 | 180.051008 | 4.96E-84 |
| stink | 14 | 178.28499 | 6.10E-05 |
| lowdown | 42 | 169.496294 | 2.27E-13 |
| cotch | 49 | 169.283412 | 1.77E-15 |
| brack | 41 | 167.974368 | 4.54E-13 |
| udder | 48 | 163.487828 | 3.55E-15 |
| skeeter | 58 | 162.828738 | 3.46E-18 |
| runway | 16 | 159.459867 | 1.53E-05 |
| gim | 13 | 158.928334 | 0.00012205 |
| sassy | 56 | 156.734057 | 1.38E-17 |
| dat | 3349 | 150.957102 | 0 |
| wench | 96 | 147.98562 | 6.13E-28 |
| wite | 62 | 140.781185 | 2.16E-19 |
| n[--] | 69 | 140.527993 | 1.69E-21 |
| n[-----]s | 438 | 140.517731 | 2.10E-126 |
| lub | 56 | 137.06908 | 1.38E-17 |
| ole | 1053 | 136.068779 | 5.52E-301 |
| fotch | 43 | 128.007961 | 1.13E-13 |
| gib | 126 | 124.090949 | 7.41E-37 |
| masser | 19 | 123.728626 | 1.91E-06 |
| massa | 469 | 123.612567 | 6.92E-134 |

Table 2: Top filtered collocates of “n[-----]”

The collocates linked to “negro” show a mixture of terms associated with both racist and anti-racist political rhetoric: *insurrections*, *carpetbaggers*, and *amalgamation*, yes, but also *enfranchisement* and *disenfranchisement*.¹⁹ Whether the rhetorical goal is to expand or to curtail African Americans’ rights, this is clearly the language of “the negro problem,” of black-white race relations as a

prompt for deliberative argument and civic action. The collocates associated with “n[-----],” by contrast, evoke not the public sphere but the private (if those two can be differentiated in the context of slavery and Jim Crow): a world of dialect and colloquialisms (*wich, cust, dat, lub*), of regional slurs and idioms (*woodpile, bluegum*), and, most distinctively, of interpersonal affects and behavior patterns. The conjunction of *onery* and *sassy* as collocates, together with *masser* and *massa*, establish a kind of narrative script associated with this term, one in which a playful or recalcitrant slave interacts comedically with a white authority figure. Although *onery* and *sassy* might be rendered in an attempt at southern African-American dialect, they describe predictable attitudes of black subordinates *as seen from the perspective of a white person*. In addition to their distinct discursive registers, then, we see that “negro” and “n[-----]” differ in their potential use in group self-determination: the latter anchors a discourse applied to African-Americans from without rather than generated from within, while the former seems more closely split between the two.²⁰

In this latter respect, it is instructive to compare “negro” with another target term, “slave.” As with “negro,” the collocates of “slave” seem ethically and emotionally ambivalent: terms like *manumission* and *emancipate* (and, again,

| Term | NObs | Obs_Exp | p-value |
|--------------|------|------------|------------|
| panthea | 15 | 566.513277 | 3.05E-05 |
| manumitted | 25 | 444.726606 | 2.98E-08 |
| catchers | 32 | 311.732173 | 2.33E-10 |
| slaveowner | 18 | 264.599015 | 3.81E-06 |
| moabite | 14 | 242.031494 | 6.10E-05 |
| nubian | 65 | 228.280266 | 2.70E-20 |
| chattel | 88 | 205.547475 | 3.22E-27 |
| emancipating | 11 | 200.028145 | 0.00048825 |
| freedman | 51 | 190.562305 | 4.44E-16 |
| bondman | 26 | 170.205767 | 1.49E-08 |
| circassian | 57 | 165.010353 | 6.93E-18 |
| galley | 282 | 163.814298 | 1.25E-81 |
| serf | 79 | 162.289856 | 1.65E-24 |
| holders | 160 | 160.515954 | 5.47E-47 |
| fugitive | 698 | 156.499573 | 3.74E-202 |
| slaveholder | 60 | 154.233981 | 8.66E-19 |
| manumission | 14 | 150.738913 | 6.10E-05 |
| emancipate | 33 | 142.375041 | 1.16E-10 |
| slave | 2732 | 130.827907 | 0 |
| oligarchy | 23 | 128.323841 | 1.19E-07 |
| drudging | 19 | 125.384211 | 1.91E-06 |
| runaway | 333 | 123.225604 | 4.27E-95 |
| drudge | 91 | 121.542475 | 1.87E-26 |
| holder | 130 | 114.797055 | 4.81E-38 |

| | | | |
|--------------|-----|------------|------------|
| quadroon | 73 | 113.422173 | 3.96E-21 |
| octoroon | 22 | 111.816975 | 2.38E-07 |
| enfranchised | 14 | 110.155359 | 6.10E-05 |
| baseborn | 11 | 100.760446 | 0.00048825 |
| abject | 269 | 100.542881 | 4.26E-76 |

Table 3: Filtered collocates of “slave”

enfranchised) sit next to *fugitive*, *chattel*, and *baseborn*. Unlike “negro,” however, “slave” encompasses a set of collocates that serve to demarcate the line between black and white, enslaved and free: *quadroon* and *octoroon* are significant collocates, as are *freedman* and *bondman*. Reflecting a somewhat narrower historical period than “negro,” these terms sketch out a semantic field that aims to police the boundaries of racial identity and political citizenship by focusing attention on liminal cases: individuals of mixed race and ambiguous political status, or those who have recently experienced a change in status (*freedman*). It is a discourse based on the close physical proximity and socio-economic entanglement of dominant and oppressed racial groups, standing in stark contrast to the collocates that we found for the target term “african,” which instead evoke a script of exoticism and peripheral contact: *explorer*, *colonization*, *jungle*. Knowing this, one might expect “ethiopian,” another target term, to center

| Term | NObs | Obs_Exp | p-value |
|------------|------|------------|------------|
| serenaders | 36 | 12247.0524 | 1.43E-11 |
| leopard | 104 | 3817.78916 | 4.13E-32 |
| christy | 81 | 2979.01275 | 3.71E-25 |
| ethiopian | 44 | 2455.95456 | 5.51E-14 |
| melodies | 99 | 1986.68178 | 1.34E-30 |
| mongolian | 11 | 1437.384 | 0.00048739 |
| songster | 16 | 1300.52233 | 1.52E-05 |
| caucasian | 14 | 1194.54734 | 6.09E-05 |
| melodeon | 12 | 1186.84886 | 0.00024361 |
| minstrels | 20 | 1185.95716 | 9.48E-07 |
| eunuch | 19 | 814.54836 | 1.90E-06 |
| spots | 98 | 509.29703 | 2.69E-30 |
| whites | 74 | 425.319198 | 4.84E-23 |
| skin | 156 | 270.934998 | 5.81E-46 |
| plantation | 41 | 144.694304 | 4.43E-13 |
| songs | 45 | 125.234316 | 2.75E-14 |
| song | 128 | 122.280614 | 1.46E-37 |
| egyptian | 11 | 110.892826 | 0.00048739 |

Table 4: Filtered collocates of “ethiopian”

a similar semantic field, with some modifications based on Ethiopia's relative insulation from colonial takeover—but here, instead, we found an entirely different thematic cluster: *serenaders*, *songster*, *melodeon*, *minstrels* ... It was not until we learned of the existence of a prominent and much-imitated blackface minstrel troupe called the “Ethiopian Serenaders” that we began to make sense of the cluster. Given the lack of overlap between genetic ancestry and cultural models of race, it perhaps should not be surprising that seemingly nested racial and ethnic categories can in fact occupy almost entirely disconnected semantic regions—but it was nonetheless startling to find this combination of local thematic coherence and global heterogeneity.

Since the fourteen categories into which we grouped our hundreds of target terms were more divisions of convenience than considered hypotheses, it is to be expected that not all of the categories would be either as internally rich or as thematically interconnected as “Black.” We created a single category for “German or Dutch,” for instance, but found little semantic overlap between those two nationalities: “german” collocates clustered around intellectual culture (*rationalism*, *universities*, *kultur*, *goethe*) and war (*bombed*, *militarism*, *submarines*), while the “dutch” were mostly represented as *navigators* and *merchantmen*. Our “Eastern European” category, on the other hand, does show internal semantic coherence, although not necessarily in the form we expected: from “russian” to “polish” to “hungarian,” these target terms conjured up the language of political radicalism and of classical music in equal numbers, with *pianist*, *waltz*, and *ballet* counterpoised against *nihilist*, *revolutionist*, and *refugee*. The absurd aptness of these collocates hints at the potential for interpretive work on this data to become a kind of spot-the-stereotype parlor game: collocates of “dutch,” after all, also included *phlegm*, *cheeses*, and *windmills*. The semantic content of these European ethnic collocates is uninteresting, almost by design: even an individual who has worked hard to counteract her own tacit racial biases has probably not extended that work to these far less poisonous (and, often, quite locally positive: *kultur*!) associations with white ethnic groups. In the absence of the horror that stems from systematic racial oppression, the concentrated stereotypes represented in our European collocate data elicit a kind of amused recognition—of *course* Russians are ballerinas and revolutionaries—that gives a contemporary reader some uneasy sense of how a white American reader in the 1840s might have reacted to seeing the collocates of, for instance, “d[----]” (*banjo*, *yah*, *massa*). While the mechanisms behind this recognition might be similar, however, their affective divergence for a contemporary reader reflects a

genuine difference between the categories, as we will see below: our historical findings suggest that the semantic associations of the “Black” and “Native American” categories are indeed far more durable (“stickier”) than those of European ethnicities, implying that these stereotypes retain a greater potential for insidious action.

In the aggregate, our collocates reveal something about the linguistic and semantic *structure* of stereotypes, irrespective of the particular group in question. In keeping with social-psychological research on stereotyping, which suggests that stereotypes function by attributing an individual’s behavior to *traits* rather than *situations*, our findings point toward the hypothesis that adjective-noun pairs may constitute a particularly salient locus for ethnic stereotypes.²¹ While collocate analysis alone does not allow us to determine which words are associated with, say, “oriental” as noun vs. “oriental” as adjective, collocates themselves tend to be nouns and adjectives rather than, say, verb forms: we have a *perplexing* “oriental” or an “oriental” *face*, a *handsome* “mulatto” or a “mulatto” *woman*. Indeed, using ethnic terms to characterize objects or abstractions seems to be as important a force in the development of stereotypes as the use of descriptive adjectives to characterize ethnicities; while psychological research might have led us to expect adjectives as collocates of ethnic terms (say, *passionate* for “italian”), many of our most significant collocates were nouns likely being modified by the adjectival term. So, for instance, we find *future* as a collocate of “anglo-saxon” and *powers* of “european”—high-minded abstractions sketching out the fate of the white race—whereas “chinese,” “japanese,” and “oriental” are more likely to be associated with concrete nouns, especially for commodities and aesthetic objects: *butterflies*, *calligraphy*, *silk*, *tea*, *illustration*, *verses*. The latter results, which echo Anne Anlin Cheng’s thesis on the role of “ornamentalism” in constructing East Asian and Asian-American women’s identities, suggest that it is not *only* as descriptors of characters or of groups of people that our ethnic target terms come into play; racial and ethnic background discourse can tinge a rug or a teapot as readily as an individual.²²

When ethnic terms do attach to persons, whether named or unnamed, our collocate data shows that ethnicity does not combine as seamlessly with other identity markers as one might assume. Although any racial background can theoretically be assigned to either male or female characters, for instance, our collocates showed a significant preference for some ethnicity-gender combinations over others, revealing intersectionality in action.²³ The collocate *man* was most likely to be associated with racialized terms in the “Black” and

“Native American” categories—“colored,” “indian,” “d[-----],” and “i[----],” for instance. Shifting to feminine collocates, however, one finds an increased drift toward exoticism and a faint tinge of sexualization: both *woman* and *girl* are significant collocates of “slave,” for instance—emphasizing female subordination and lack of power—while such diverse ethnicities as “irish,” “german,” “chinese,” “mexican,” “arab,” and “hebrew” also have female collocates. Given the prevalence of the “tragic mulatta” storyline in American fiction, it is particularly appropriate to find *woman* as a collocate of “mulatto,” hinting at the sexualization of light-skinned African American women in many nineteenth- and twentieth-century American fictions.²⁴ And *girl* in particular frequently appears as a noun modified by an ethnicity: the “Mexican girl,” “Indian girl,” “Chinese girl,” and “Arab girl” are in an important sense interchangeable—all similarly eroticized, all defined primarily by their distance from an implied white male protagonist. One telling line from Louise C. Ellsworth’s 1892 romance *Furono Amati* reveals the misogyny that underlies and feeds upon these racialized caricatures: “he hated girls in general,” the narrator tells us of a focalizing character—“the genus girl of which ‘Irish Lizzie’ was a specimen.” Tellingly, “Irish Lizzie” is invoked in invidious comparison to a beautiful girl that our hero *does* admire; though this “Isabel” is not explicitly racially marked, her “liquid blue eyes” and “golden halo of curls” strongly point toward a white Anglo-Saxon identity, indicating how readily sexist ideals can be invoked to enforce racial hierarchies (and vice versa).²⁵

The slightly different target terms associated with *woman* and *girl*, moreover, suggest that age as well as gender has a relationship of mutual influence with racial and ethnic identity. *Young*, for instance, is a collocate of “italian,” while *old* attaches to “dutchman,” echoing the relative historical moments at which these two ethnic groups immigrated to the United States (at least during the 1789-1920 period of the Gale corpus).²⁶ Interestingly, many slurs—including especially those from the Black category—have a preferential association with *old*, which, an examination of collocates in context reveals, directly modifies for the target terms in question: attaching “old” to a slur not only emphasizes the degrading familiarity that the term is intended to convey, but, in a postbellum context, reflects the plantation nostalgia that fictions by and for white Southerners produced and capitalized upon. One significant exception to the association of *old* with “Black” terms, however, is “mulatto,” which has *young* as a collocate—again, perhaps, hinting at the vulnerability and sexual desirability associated with that racial term, together with the implicit

generational dynamic (one black parent, one white parent, one mixed-race child) that it evokes. The link between *young* and “mulatto” also reflects a broader tendency, revealed in contextual analysis, for *young* ethnically marked characters to be represented as heroic and/or sexualized, from the “gallant young Cuban Fernando Perez” to “a tall, graceful, and exceedingly handsome young Aztec woman named Margarita Ayla” to “a handsome young Italian laborer who had on his person no clothes whatsoever”.²⁷ Even when characters whose ethnicities are coded as *young* remain peripheral or are represented as sinister, they bear semantic associations with physical health and beauty, quick wit, and the potential for upward mobility: several “mulatto” servants are described as “young and active,” “intelligent [and] gayly turbaned,” or “as attractive ... as a Moorish statue would be” even as they are narratively subordinated to minor roles, while many a “young Jew” appears “pushing his way from the Ghetto to the places of power”.²⁸

In our corpus, then, we find a large cast of racialized stock characters occupying not only particular ethnic niches but also predictable social positions determined by gender, age, and foreignness. Where, in all this, is that dominating figure of American history and literature, that maximally unmarked character—the white man? A clue came to us when, investigating appearances of the collocate *man* in context, we found that its association with a target term like “indian” was driven by contrast rather than equivalence: instead of modifying *man* as an adjective, “indian” tends to appear in these passages as a noun positioned in opposition to the phrase “white man”. The same, it should be noted, is true for “n[----]”—and the fact that it is *these* two categories, “Black” and “Native American,” that nineteenth-century American authors explicitly contrast to whiteness suggests a basic assumption of racial difference far more deep-seated and dehumanizing than in the case of other ethnic groups.²⁹ Indeed, while some passages that insist upon this racial difference depict “the white man” and his racial others as two unequal but human groups—as in the melancholic-genocidal certainty of one novel’s declaration that “the white man and the Indian cannot live together[, t]he latter dies while the first lives and prospers”—others use the contrast to assign the racialized group to a not-quite-human status, sometimes literally between white people and animals: another text describes “a sloop called the *Sea Fox* manned by a white man, an Indian, and a dog.”³⁰

Yet while these references, largely delivered in the narratorial voice, overtly reinforce a white supremacist and misogynist system that positions white men as the only fully human beings, the bigram “white man” often appears in

quite another context: the voice of racially marked characters themselves. These references usually occur in represented speech, as when an enslaved character remarks, “what a careless creetur’ dat white man is”—albeit in the context of a deferential conversation with his “mars’.”³¹ At times, though, they even make their way into a kind of free indirect discourse, as when a narrator partially enters the mind of the character Chee Ming only to reinforce orientalist ideas of mysteriousness and opacity: “It was impossible to guess what Chee Ming thought” or to know “[i]f he had been holding any suspicion against the white man who had ridden with him from the Yanggun gate.”³² In books written by and (presumably) for white people, such passages suggest, whiteness becomes visible only through the adopted perspective of a racialized character to whom these writers ascribe wariness, animosity, or disdain. Yet far from actually opening up space for a critique of whiteness, these moments re-inscribe racial difference as an inevitable fact legible even to those whom it most disadvantages, representing a world in which racialized characters cannot escape whiteness even in their own consciousness. By placing the acknowledgment that a *man* can (and indeed, in the racist framework of these texts, must) be “white”—the least personalized, recall, of all the racial and ethnic categories we tracked—in the mouths and minds of racially marked characters, the white writers in our corpus foist the recognition and maintenance of racial difference onto those who lose most from it, abdicating responsibility even for their own so carefully constructed and defended whiteness. To compel their own belief in their performance of racial logic, these writers must also act the part of the audience—sometimes appreciative, sometimes hostile, always (imagined to be) captive.

Results, Part 3: On “Stickiness”

Although the word frequency results in the first section are able to provide a macroscopic overview of how much the language of identity was used over the nineteenth century, and the collocate results offer a microscopic view of the semantic richness of individual terms, we still lacked a way of understanding the ways in which the constellations of associated words revealed through the collocate analysis work in concert to define the various identities found in American Literature. Although the collocate analysis reduces the problem from 18,000 texts to 144,000 words, the complexities in these relationships are still

difficult to systematize. In particular, we wanted a way to capture two key concepts that we combined under the rubric of “stickiness”: the tendency of groups of collocates to “stick” to particular ethnic groups (which we came to call Sticky 1) and the tendency of collocates to “stick” to each other as they travel between identities (Sticky 2). In other words, Sticky 1 labels the consistency of an association between a target term and its collocates over time: perhaps “irish” is consistently associated with *catholic*, *priest*, and *famine* from the 1840s on. Sticky 2, on the other hand, labels the consistency of collocate groups themselves: perhaps *catholic* and *priest* stick together even when they attach to the target term “italian,” while *famine* no longer travels with them.

To quantify how our collocates were related to each other based on their textual co-occurrence — that is, their tendency to appear in the same frequencies in the same texts — we used a Term-Document matrix to calculate the scaled frequency (per 100000 words) of each collocate in each text, thereby placing each collocate in relation to every other through the 18,202 dimensional vector of their mutual frequencies in our corpus. The resulting model is a sparse vector representation of collocate co-occurrence, which differs from the dense vector models of word embedding analyses such as word2vec or GloVe by virtue of the calculable association between terms, rather than a probabilistic model.³³ Given the size of our corpus, visualizations of the sparse vectors produce similar results to a word embedding model with the benefit that our measurements of word similarities are intuitively interpretable, and different embedding models can be directly compared between time periods.

Given the immense complexity of the resulting data, we turned to t-Distributed Stochastic Neighbor Embedding (tSNE) to create graphs that could act as a window onto both forms of stickiness, showing how collocates associated with each target group change their relationships to each other over time. The advantage of tSNE over other methods of dimensionality reduction lies in its distributed method of embedding observations within local structures while retaining meaningful global patterns.³⁴ The overall placement of clusters within the graph, therefore, gave us the ability to interpret the broad trends in how different clusters of words related to each other, while individual clusters revealed their internal associations based on the individual associations of words they contained.³⁵ The relative positions of collocates in the clusters revealed by the tSNE graphs for each of our 25-year periods offer an intuitive way to make qualitative assessments about the semantic “regions” each group inhabits across history. For instance, in the 1815-1839 period (see Figure 5.1), many of the terms




Figure 5.2: Detail of t-SNE for the period 1865-1889. To make it more legible, this image has been filtered to retain only those collocates that were primarily attached to terms in the Native American group.

A similar genre transformation informs a major cluster of collocates in our Black group. The word “chile”, here most often a dialect spelling of “child”, tells the story well. In the 1815-1839 period, it does not appear at all, and in general terms from the Black category are widely dispersed, in our view failing to cluster in any readily legible fashion. In short, the discourse of African Americans is not yet especially internally coherent, nor is it obviously tied to any genres or clusters of stereotypes a current reader would be apt to recognize. In the 1840-1864 period, there are several large clusters on the periphery of the main mass of collocates, including one area that clearly centers around ocean tales and another that is highly indicative of religious terminology. The word “chile” appears in the vicinity of a similarly well-defined new cluster, near words like “mammy”, “masser”, and “missus”. A nearby island is predominately a cluster of dialect words, and the word “banjo” appears as well. In these groups we glimpse the beginnings of the discourse of the “Happy Slave” narrative, in which slaves are figured as comical, carefree figures under the benevolent care of white masters.³⁶




Figure 6.2: Detail of t-SNE for the period 1890-1914

easily visible clusters simply do not exist on the same scale. Qualitatively, this might suggest a distinction between the discourses of race (centering on legally and phenotypically “othered” Native Americans and Black Americans) and ethnicity (centering on “everyone else,” meaning, in this period, mostly people of European ancestry). If this explanation is correct, our t-SNEs would provide indirect evidence for Eric Foner’s claim that that “immigrant groups suffered severe discrimination, but being discriminated against did not make them non-white.”³⁷ In any case, however, the result certainly points to the historical contingency of these groups in the first place. We can only truly assess the stickiness of our collocates by answering a bigger question: whether, in any given period, there was even a group for the collocates to stick *to*.

To answer this question we require a different way of measuring stickiness: rather than measuring whether individual collocates cohere to discrete clusters over time, we need to instead measure how stable our groups themselves were compared to other identity categories at a given time. That is, to what degree do the groups in each of our 25-year periods a) contain a group of collocates that is unique to that group (not shared among other groups of collocates), and b) have a significant overlap between the collocates of their constituent targets such that they represent a holistic group (rather than a diverse set of potentially unrelated identities)? For the first metric, we measured what we call here the *external distinctness* of our groups: the percentage of collocates of each group that belong

only to that group. The more collocates a group contains that are unique to that group, the more distinct it is: its terms denote the specific identity it names rather than a range of possible identity-based subject positions. In the second metric, we measure what we call the *internal coherence*: the percentage of collocates shared among the target terms that make up that group. As in our analysis described above, we found collocates based on their co-occurrence with target terms (for example “african”), which we then group under broad identity headings (in this case “Black”). This metric allows us to test how meaningful that overall group heading is within a given period. If a group heading shares a significant percentage of collocates among its constituent target terms, then by the logic of our project, those terms are all working together to identify a single cohesive group. Conversely, the less coherence among the target members of a group, the less likely that group describes a single, stable, identity in that period.




Figure 7: External distinctness (a) and internal coherence (b) of groups based on shared collocates across target terms. The size of the dots represent the number of collocates in each group.

Figure 7 shows the graph of both external distinctness (a) and internal coherence (b) for our groups at each of our individual periods. In a), the higher the group’s identifier, relative to the other groups, the more unique the words in that group are, and the more distinct the group. In b), the higher the group, the more coherent it is: more terms are shared among the collocates of each of its target members. Within each period, this graph shows substantial differences in

the coherence of each group. Most notably, the Black and Native American categories are consistently more coherent than the other groups across both metrics, with the Black category containing collocates that overlap the least with every other group, and the Native American category containing collocates that form a tight, interlinked discourse among our target terms. And as we near the later periods, there is a rise in the prominence and coherence of the Chinese/East Asian category, one that accords with our sense in other areas of this project that literature was increasingly focused on and discursively consistent about this group in the latter half of the 19th century. In other cases the categories themselves seem to fail; the Middle Eastern/Muslim category, though extremely high in internal coherence in the middle periods, has quite poor external coherence overall. This likely owes something to its substantial overlap with the Jewish category, which in turn probably reflects a pronounced Biblical discourse that recurred throughout our time periods and entangled discussion of Jewish people (and the Middle East) with a particular set of valences that were not always racialized/ethnicized in the same way as, say, a Natty Bumppo or Uncle Remus story. Most apparent of all, the graph shows fluctuation. Differences in corpus size make it difficult to track any one group over time; later periods just have many more *words* in total, and this drives many of the changes in collocate behavior. But we can say that the relative rankings are very different from one period to the next; sometimes the Native American and Black categories are farther away from the crowd than at other times; sometimes Irish is higher than Scandinavian, and sometimes it is lower; and so on. The specific mix of racial and ethnic discourses changed dramatically over the course of the long 19th century.

These final results point to a few linked conclusions. First, they support our sense that the discourse surrounding Black Americans and Native Americans was uniquely coherent and distinct in comparison with the language applied to other racial and ethnic groups. We attribute this finding to differential racialization: while most of our target groups retained at least some association with ethnic identity (national heritage, a shared language, and so on) over the course of the nineteenth century, white discourse around people of African and indigenous American ancestry deemphasized these ethnic markers in favor of biologized racial characteristics. Indeed, these categories functioned in part to erase distinctions among, say, Wolof and Igbo speakers, or members of the Sioux and Apache nations. Although the distinctions between race and ethnicity articulated in footnote one were important to explaining the logic of our target terms, we did

not expect our collocates data to differentiate itself consistently on the basis of ethnic or racialized groups; on the contrary, we chose our groups under the assumption that each of them had been racialized during at least *some* part of the long nineteenth century. These group results are thus particularly striking, and while it is important not to identify the internal coherence and external distinctness of a discourse with racialization itself, our findings certainly support the idea that non-Black or Native American ethnicities—whether European, East Asian, Latin American, or Middle Eastern—were discursively permeable in a way that Black and Native American identities were not. Since much discourse around demographic change in America still assumes that what Robert Blauner calls the “framework of immigration and assimilation that is applied to European ethnic groups” works equally well for all other racial and ethnic identities, the distinctness of discourse around our Black and Native American categories is significant—as is the relative *indistinctness* of most ethnicities.³⁸ Rather than a number of discrete groups that move closer to or further from an unmarked neutrality as the ethnicities they represent are more or less othered by white Americans, we find a relatively fluid “ethnic” discourse—similar to “ethnic” cuisine in its status as simultaneously marked and generic—largely separated from the much larger and more internally coherent discourses of blackness and “indian-ness”.

Many mainstream and authoritative representations of race, ethnicity, and ancestry—the U.S. Census, for instance, or the results of a DNA test—represent these human population categories as relatively symmetrical. That is, “Black,” “White,” “East Asian,” “Native American,” and so on label different clusters of variables in a coherently organized space; even though the *content* of each category is different, they are all *structurally* similar (in that they label particular genetic patterns, particular phenotypes, or whatever the case may be). Our results suggest that, when it comes to literary language, racial categories are in fact radically asymmetrical—not simply in valuation, but in form. Discursively, it is not the case that positive or negative language attaches to preexisting kinds of persons as their social fortunes rise and fall; rather, that rising or falling is itself complexly indicated by the availability of coherent, discrete language to describe those kinds of persons. Our results suggest that racial categories, despite the veneer of common sense that gives stereotypes their apparent obviousness, have often failed to achieve lasting consensus, particularly as components of an overarching taxonomy. Instead, racial common sense, to slightly repurpose Omi and Winant’s term, seems to be constituted by a number of freestanding character

types that vary widely in discursive detail, narrative flexibility, and perceived distance from normative whiteness.

This fact reminds us of the ultimate contingency and instability of racial and ethnic categories, not only by comparison with the concept of a timeless and unified human species, but by comparison *with themselves* over historical time. In our contemporary moment, when some geneticists and philosophers of science are seeking to recuperate the folk concept of race by identifying it with “biogeographical ancestry”—asserting, for instance, that “what ordinary folk in the U.S. mean by ‘race’” corresponds to a 5-category taxonomy that is echoed in population genetics³⁹—it is especially important to underline that “what ordinary folk in the U.S. mean by ‘race’” is itself neither fixed nor internally consistent. This awareness can prepare us to ask different questions, and obtain more precise answers, about the discourse of race and ethnicity in the contemporary United States. By framing our current cultural moment, for instance, as a particularly intense and aggressive episode in the longer story of the creation of “Latiné/Hispanic” as a “race”, we can recognize the ways that language is being used to construct an internally cohesive category of persons—just as it was for the “Black” and “Native American” categories in the nineteenth century. That solidification may take place *both* through the shaping of explicit policies excluding members of this group from the category of “citizen” or “American,” *and* through linguistic associations shared by both malicious actors and the people resisting them. As we fight back against the former process, it seems important to keep our scholarly eyes trained on the latter—to keep racialization in the foreground, precisely so that it cannot slip into the background of common sense.

References

¹ Brigitte Fielder, *Relative Races: Genealogies of Interracial Kinship in Nineteenth-Century America* (Durham: Duke University Press, 2020), 247.

² Our usage of the terms “race” and “ethnicity” and “ancestry” will both loosely follow that of Paula Moya and Hazel Markus in their anthology *Doing Race: 21 Essays for the 21st Century* (New York: Norton, 2010), as well as current usage of the terms in anthropological literature that explicitly takes race as its subject (see, for example, Charmaine D. Royal, John Novembre, Stephanie M. Fullerton, David B. Goldstein, Jeffrey C. Long, Michael J. Bamshad, Andrew G. Clark, “Inferring Genetic Ancestry: Opportunities, Challenges, and Implications,” *American Journal of Human Genetics* 86.5 (2010): 661-673). That is, we will use the term “ethnicity” to describe a social group based on shared nationality, language, religion, and other cultural markers. We will use “race” to describe social groups based on phenotypical markers like skin color and the pseudo-biological justifications that have developed to explain these markers. And we will use “ancestry” to describe groups based on bio-evolutionary geographies that inform both global migration and immigration. In what follows, then, “Polish” or “Congolese” will be ethnicities, “white” and “black” will be races, and “African” or “European” will be ancestries. Occasionally, for the sake of brevity, we may use the term “race” to describe the

overarching system of human classification that results in our more specialized uses of “ethnicity” and “race”; this usage will be clear from context. These distinctions should be understood as discursive artifacts rather than as descriptors of inevitable reality.

³ This research is exemplified by the work of Jennifer Eberhardt, who has contributed to popularizing the concept of implicit bias; see, for instance, Jennifer L. Eberhardt, Phillip A. Goff, Valerie J. Purdie, and Paul G. Davies, “Seeing Black: Race, Crime, and Visual Processing,” *Journal of Personality and Social Psychology*, 87.6 (2004): 876-893. In recent years, both the methodology of implicit bias studies and the rhetorical use of the concept in public conversations about racism have been critiqued by some scholars: Jonathan Kahn, for instance, argues that implicit bias offers a convenient way to avoid addressing structural racism. While these concerns about the rhetorical and political uses of psychological research seem quite warranted, the existence of implicit bias itself has been substantiated across many types of psychological measures and tests and constitutes a subset of the much broader, well-understood phenomenon of unconscious processing of perceptual information. Moreover, critiques like Kahn’s suggest that implicit bias is flawed as a framework precisely because it *doesn’t* sufficiently shift the focus away from the individual; we share that goal with analyses of institutional racism. Jonathan Kahn, *Race on the Brain: What Implicit Bias Gets Wrong About the Struggle for Racial Justice* (New York: Columbia University Press, 2017). For an overview of psychological methods of evaluating implicit attitudes — including the most famous such tool, the Implicit Association Test (developed by Anthony Greenwald) — see Bernd Wittenbrink and Norbert Schwarz, eds., *Implicit Measures of Attitudes* (New York: Guilford Press, 2007); for a summary of critiques of the Implicit Association Test, see Olivia Goldhill, “The World is Relying on a Flawed Psychological Test to Fight Racism,” *Quartz*, <https://qz.com/1144504/the-world-is-relying-on-a-flawed-psychological-test-to-fight-racism/>.

⁴ Michael Omi and Howard Winant, *Racial Formation in the United States*, third edition (New York: Routledge, 2015).

⁵ We thank Kinohi Nishikawa for directing our attention to this point.

⁶ Martin Delany, *Blake; or, The Huts of America: A Corrected Edition*, ed. Jerome McGann (Cambridge: Harvard University Press, 2017), xv, xxxiii-xviii.

⁷ Throughout our analysis, we approached questions of historical change by dividing our corpus into 25-year segments. In the earlier nineteenth century, when relatively few novels were published per year in the United States, this allowed us to accumulate enough texts to draw meaningful conclusions; even so, we ended up discarding the first segment (from 1789-1814) because it simply did not contain enough text. In the later segment of our corpus, 25-year chunks were small enough that the quantity of texts was not overwhelming. We still took care, however, to normalize our word counts so that the sizes of the 25-year subcorpora did not drive our results.

⁸ “Chief” was one of the target terms that, we soon realized, was perhaps too polysemic to tell us much about racial discourse; its prevalence is partly accounted for by non-racialized uses (chief of police, chief executive, “chief” as a term of address, and so on). As we discuss below (with respect to color terms), it was difficult to decide what to do with such polysemic terms; ultimately, we left in “chief” because it did seem to capture a substantial amount of discourse about Native Americans.

⁹ Jill Lepore, *The Name of War: King Philip’s War and the Origins of American Identity* (New York: Knopf, 1998), 193.

¹⁰ Note that this data reflects where respondents were born; it is not a measure of the number of immigrants in each decade, and many people should appear in the data more than once—e.g., a 30-year-old Irish Bostonian in 1850 will still be from Ireland in 1860; she’ll just be 40.

¹¹ This act was signed in 1882, and was not repealed until 1943, although legal Chinese immigration did not resume in large numbers until the Immigration Act and Nationality Act of 1965. Our category “East Asian” includes both Chinese and Japanese immigrants; increased immigration by the latter in the later periods of our study counterbalances some of the decline in immigration by the former. For more on this, and on the “Yellow Peril” rhetoric mentioned below, see Gordon Chang, “Eternally Foreign: Asian Americans, History, and Race” in Markus and Moya, eds., *Doing Race*, 216-233.

¹² See, for instance, Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean, “Efficient Estimation of Word Representations in Vector Space”, *Arxiv.Org* (2013), and Jeffrey Pennington, Richard Socher and Christopher Manning, “Glove: Global Vectors for Word Representation,” *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing* (Association for Computational Linguistics: Doha, 2014), 1532-1543.

¹³ Paul Baker, *Using Corpora in Discourse Analysis* (London: Bloomsbury, 2006), Chapter Five.

¹⁴ To assess significance, we used a Fisher’s Exact Test using the frequency of a word as a collocate of a target as the observed, and the probable number of occurrences of that word as a collocate of the target given its overall frequency in the corpus. Significance was assessed using $\alpha=0.05$.

¹⁵ Given the scale of the corpus, both methods, word embeddings and collocate analyses, begin to converge. The sheer number of collocates captures both synonyms and antonyms, albeit at a lower significance than the vector model; while the vector model reveals that many proximal words, over the length of the corpus, share

contexts, although again at a lower frequency. Given the interpretability of the collocates, however, this method remained the most useful.

¹⁶ The actual number of collocates, including duplicates, in the table was 142,177, a staggering result.

¹⁷ Werner Sollors, *Neither Black Nor White Yet Both: Thematic Explorations of Interracial Literature* (Cambridge: Harvard University Press, 1999), 18-22.

¹⁸ We also removed “yellow,” which we thought might collocate with a racist discourse about people of East Asian descent, but which turned out to be overwhelmingly literal. As with “black” and “white,” however, we retain some sense that a yellow silk gown, for instance, is not *merely* a yellow silk gown but an indirect evocation of Chinese culture and people, even when no explicit term referencing them is present.

¹⁹ The strongest connection to “negro,” measured by statistical significance, is of course the word at the top of table 1. In order to avoid reinforcing the connection between this target and collocate, we name the term here in a second sentence and well outside of a ten-word window: the word “rapist”. While this term clearly has a strong negative valence and is important for understanding how Black figures were depicted in nineteenth-century American fiction, it is worth noting that these two words appear together only 11 times. Further investigation reveals that 8 of these uses come from Marion Curtis’s 1907 novel *The Note of Discord*. That said, the word “rapist” appears in the corpus only 26 times and it is significant that this paring accounts for 46% of its total mentions in our corpus. Whatever its cultural force, Curtis’s forgotten text still has an undeniable statistical effect on the use of both words. These two important words really are linked, and it really is (at least within this corpus) because of this one novel among many thousands of others. For researchers studying linguistic artifacts at scale, the potential statistical importance of small numbers and single texts should occasion serious ethical reflection.

²⁰ In part, this may reflect the historical limitations of the Gale corpus, which stops in 1920; if we were to follow the term “[n-----]” into the present, we may find that its partial and contested reclamation leads its semantic field to bifurcate in a way that resembles that of “negro.”

²¹ On this heuristic, known in the psychological literature as the “Ultimate Attribution Error,” see Thomas Pettigrew, “The Ultimate Attribution Error: Extending Allport’s Cognitive Analysis of Prejudice,” *Personality and Social Psychology Bulletin* 5.4 (1979): 461-476; Susan T. Fiske, “Social Cognition and the Normality of Prejudgment,” in John F. Dovidio, Peter Glick, and Laurie A. Rudman, *On the Nature of Prejudice: Fifty Years After Allport* (Oxford: Blackwell, 2005), 36-53; and Tracie L. Stewart, Ioana M. Latu, Kerry Kawakami, and Ashley C. Myers, “Consider the Situation: Reducing Automatic Stereotyping through Situational Attribution Training,” *Journal of Experimental Social Psychology* 46.1 (2010): 221-225.

²² Anne Anlin Cheng, “Ornamentalism: A Feminist Theory for the Yellow Woman,” *Critical Inquiry* 44.3 (2018), 415-446. No doubt these references to Chinese and Japanese objects offer a partial explanation for the disproportionate increase in East Asian terms even as East Asian immigration increased more slowly (mentioned above); it would be interesting to investigate whether these periods of frequent East Asian references correspond to *chinoiserie/japonaiserie* trends in fashion or interior decoration.

²³ Such effects, as Sharon Block points out, were built into language itself at least as of the eighteenth century, when advertisements for fugitive slaves and servants “integrated sex and race divisions” with terms like “wench” (for African-American women) or “[s-----]” (for Native American women) “that reframed their gender into a silent component of their racial designation” (100). Sharon Block, *Colonial Complexions: Race and Bodies in Eighteenth-Century America* (Philadelphia: University of Pennsylvania Press, 2018).

²⁴ According to Eve Allegra Raimon, the “tragic mulatta/o” concept dates at least back to Sterling Brown’s 1937 *The Negro in American Fiction*; Raimon herself insists on the specific importance of mixed-race women in American literature, noting that “the sexual vulnerability of a female light-skinned slave” is key to “the very tragedy of the figure’s fate” (5). Eve Allegra Raimon, *The “Tragic Mulatta” Revisited: Race and Nationalism in Nineteenth-Century Antislavery Fiction* (New Brunswick: Rutgers University Press, 2004).

²⁵ Louise C. Ellsworth, *Furono Amati: A Romance* (New York: United States Book Company, [c1892]), *American Fiction* (accessed June 12, 2019), <http://tinyurl.galegroup.com/tinyurl/AErs48>.

²⁶ “Jew” also has *young* as a collocate, but “hebrew” is more strongly associated with *old*, hinting that the latter connotes a religious and cultural tradition more than a particular demographic group — a hypothesis supported by some other collocates of “hebrew” with a very high observed/expected rating: *chaldaic*, *aramaic*, *syriac*, and *septuagint*. This suggests a bimodal set of associations with Judaism: the Biblical past, and an ethnic immigrant group of the late nineteenth/early twentieth century present.

²⁷ Antoinette Sheppard, *The Heroine of Santiago De Cuba: A Sequel, or, What Followed the Sinking of the Merrimac* (New York: Abbey Press, [c1900]), *American Fiction* (accessed June 12, 2019), <http://tinyurl.galegroup.com/tinyurl/AErh28>; Francis Fontaine, *The Modern Pariah: A Story of the South* (Atlanta, Ga.: Published by Francis Fontaine, [c1892]), *American Fiction* (accessed June 12, 2019), <http://tinyurl.galegroup.com/tinyurl/AErjT9>; and Gouverneur Morris *The Penalty*, illustrated by Howard Chandler Christy (New York: Charles Scribner’s Sons, 1913), *American Fiction* (accessed June 12, 2019), <http://tinyurl.galegroup.com/tinyurl/AErkG2>.

²⁸ Robert Montgomery Bird, *The Hawks of Hawk-Hollow: A Tradition of Pennsylvania*. Vol. 2. (Philadelphia: Carey, Lea & Blanchard, 1835), *American Fiction* (accessed June 12, 2019), <http://tinyurl.galegroup.com/tinyurl/AErig3>; J.R. Beckwith, *The Winthrops: A Novel* (New York: Carleton, Publisher, 1864), *American Fiction* (accessed June 12, 2019), <http://tinyurl.galegroup.com/tinyurl/AErkt2>; Teresa Hammond Strickland, *Under the Ban: A South Carolina Romance* (Chicago: Rand McNally, [c1898]), *American Fiction* (accessed June 12, 2019), <http://tinyurl.galegroup.com/tinyurl/AEmp6>; and James Oppenheim, *Pay Envelopes: Tales of the Mill, the Mine and the City Street*, illustrated by Harry E. Townsend (New York: B. W. Huebsch, 1911), *American Fiction* (accessed June 12, 2019), <http://tinyurl.galegroup.com/tinyurl/AErom5>.

²⁹ It is worth noting that at least one nonfictional genre seems to show the same pattern: Sharon Block has found that eighteenth-century advertisements for runaway slaves or servants rarely referred to European-descended people as “white,” but when they did, “it was a comparison to someone described as ‘Negro’ or ‘Indian’” (Block, *Colonial Complexions*, 66).

³⁰ The first citation comes from John Hovey Robinson, *Silver-Knife, or the Hunters of the Rocky Mountains: A Romance of the Wild West* (Boston: Hotchkiss & Company, 1850), *American Fiction* (accessed June 12, 2019), <http://tinyurl.galegroup.com/tinyurl/AEryR9>; the second from Charles Clark Munn, *Pocket Island: A Story Of Country Life In New England* (New York: The Abbey Press Publishers, [1900]), *American Fiction* (accessed June 12, 2019), <http://tinyurl.galegroup.com/tinyurl/AErpu3>. Interestingly, the “white man” in the latter scene is in fact Jewish — “a Polish Jew peddler named Wolf” with “sinister black eyes.” In the context of the coastal Maine towns (populated, as the novel soon shows us, by white New England Protestants) where Captain Wolf runs his bootlegging operation, the more salient identity is indeed “Jew”; when compared to his “roving Micmac Indian” partner, however, his whiteness comes to the fore.

³¹ E.D.E.N. Southworth, *The Rejected Bride: “Only a Girl’s Heart”*, illustrated by Hugh M. Eaton (New York: R. Bonner, 1894), *American Fiction* (accessed June 12, 2019), <http://tinyurl.galegroup.com/tinyurl/AEruq8>.

³² Jeremy Lane, *Yellow Men Sleep*, illustrated by George W. Gage (New York: The Century Co., 1919), *American Fiction* (accessed June 12, 2019), <http://tinyurl.galegroup.com/tinyurl/AErrQ3>.

³³ Mikolov et al., “Efficient Estimation of Word Representations in Vector Space”; Pennington et al., “Glove: Global Vectors for Word Representation.”

³⁴ Laurens Van der Maaten and Geoffrey Hinton, “Visualizing High-Dimensional Data Using t-SNE,” *Journal of Machine Learning Research* 9 (2008): 2579-2605.

³⁵ As our interpretation relies on a visual assessment of the graph, we took into account the local distortions introduced by t-SNE plots by increasing the number of iterations and varying the perplexity of our analysis. Moreover, our analysis relies more on cluster membership than the relative shapes of the clusters. See Martin Wattenberg, Fernanda Viegas and Ian Johnson, “How to Use t-SNE Effectively,” *Distill* 1.10 (2016): e2.

³⁶ Though this discourse reached its popular apotheosis in the post-Reconstruction works of Joel Chandler Harris, it was already evident in Southern writing by around 1835, especially in magazine fiction, where it served a clear pro-slavery rhetorical function. See, for example James Hall, *The Harpe’s Head: A Legend of Kentucky* (Philadelphia: Key & Biddle, 1833); William Alexander Caruthers, “The Knights of the Golden Horse-Shoe: A Traditionary Tale of the Cocked Hat Gentry in the Old Dominion. Chapter VI: A Kitchen Fire-Side in the Old Dominion,” *Magnolia; or Southern Monthly* (1841-1842): 75-76; and especially Henry Ruffner, “Judith Bensaddi: A Tale,” *Southern Literary Messenger* 5.7 (1839): 477.

³⁷ Eric Foner, “Review: *The History of White People* By Nell Irvin Painter,” *Harper’s* September 2010: 74-78.

³⁸ Robert Blauner, *Racial Oppression in America* (New York: Harper & Row, 1972), 52.

³⁹ Quayshawn Spencer, “The Unnatural Racial Naturalism,” *Studies in History and Philosophy of Biological and Biomedical Sciences* 46 (2014): 38-43. Spencer draws on the work of Noah Rosenberg to support his claim that genetics can help us sort human populations into five clusters, but does not cite any social science research to support his contention that the U.S. folk conception of race is similarly organized; this asymmetry, far from being unique to Spencer, is characteristic of contemporary philosophical investigations of race.