

RESEARCH ARTICLE

Digital accents, homogeneity-by-design, and the evolving social science of written language

AJ Alvero¹, Quentin Sedlacek², Maricela León² and Courtney Peña³

¹Center for Data Science for Enterprise and Society, Cornell University, Ithaca, NY, USA; ²Department of Teaching & Learning, Annette Caldwell Simmons School of Education & Human Development, Southern Methodist University, Dallas, TX, USA and ³Stanford University School of Medicine, Stanford, CA, USA
Corresponding author: AJ Alvero; Email: ajalvero@cornell.edu

Abstract

Human language is increasingly written rather than just spoken, primarily due to the proliferation of digital technology in modern life. This trend has enabled the creation of generative artificial intelligence (AI) trained on corpora containing trillions of words extracted from text on the internet. However, current language theory inadequately addresses digital text communication's unique characteristics and constraints. This paper systematically analyzes and synthesizes existing literature to map the theoretical landscape of digitized language. The evidence demonstrates that, parallel to spoken language, features of written communication are frequently correlated with the socially constructed demographic identities of writers, a phenomenon we refer to as “digital accents.” This conceptualization raises complex ontological questions about the nature of digital text and its relationship to social identity. The same line of questioning, in conjunction with recent research, shows how generative AI systematically fails to capture the breadth of expression observed in human writing, an outcome we call “homogeneity-by-design.” By approaching text-based language from this theoretical framework while acknowledging its inherent limitations, social scientists studying language can strengthen their critical analysis of AI systems and contribute meaningful insights to their development and improvement.

Keywords: large language models; computational text analysis; AI homogenization; sociolinguistics; sociology

Introduction

Much has been written about the dominance of screens in daily life, especially in digital technology like computers and cell phones. Producing and consuming text on those screens has become integral in the era of digital technology – not just in the United States, but around the world, even in highly remote communities (Porter, 2012). There is a robust literature from the cognitive and non-cognitive social sciences that has examined the relationship between spoken (and signed) language and society;

however, given the increasing ubiquity of (almost always digitized) text, written language should receive more scholarly attention as a fundamental medium of language, communication, and social interaction. This is especially true in light of developments in generative artificial intelligence (AI) and large language models (LLMs) which enable anyone with access to the internet to have an LLM produce text based on any given prompt.

We draw on recent studies of AI, text analysis, language, and sociology to illuminate the origins and implications of two theoretical constructs we introduce: digital accents and homogeneity-by-design. Sociolinguists have long studied variation in spoken language and its relationship to demography, identity, and social context. Our paper brings together disparate lines of recent research that suggest similar patterns in written language: a strong relationship between demography, writing, text, and various social outcomes and conditions. In doing so, we highlight how our theories point to key tensions in established modes of conceptualizing and studying language and AI. These tensions matter because failing to address them could leave language scholars on the outside looking in when it comes to the ethical development of generative AI and to grounded, thoughtful critiques of this technology. Though this paper is primarily focused on conceptual and theoretical development, we supplement our arguments (particularly those concerning the concept of homogeneity-by-design) with insights from a survey of dual-language teachers in the United States.

While this paper takes more of a critical perspective on relationships between written language, AI, and society, we also want to be clear that our goal is not to argue that generative AI and LLMs are inherently bad or good. Rather, we argue that these technologies have important implications for the ways that humans use language and communicate, and that social scientists must reexamine extant theories and perspectives in order to effectively grapple with written language and society, especially given the advances made over a short period of time with large language models. We conclude by highlighting the ways that social scientists studying language could contribute to a broader understanding of generative AI and LLMs in two ways. One could be by developing an informed critique of LLMs, their uses and development, and their implications for human language practices. Another way could be through the provision of linguistic expertise to computer scientists about language and society in order to promote ethical and equitable generative AI development.

Digital accents

The well-known saying “everyone has an accent” has been used by language scholars to argue that there is no “good” or “bad” language. Instead, the ways we speak, our vocabulary, grammatical construction, pronunciation, and other language features are all substantially associated with demography (e.g., race, gender, social class, geography). Of course, demography is not destiny; there is no inherent or essentialist reason why people who share a particular racial identity, gender identity, social class, or geographic location would necessarily use similar language patterns. Indeed, we are wary of associating these identities or experiences too closely with language practices, lest we fall into the trap of endorsing commonplace ideologies that associate language practices with social identities such as race regardless of the actual features of

individuals' language use (Flores & Rosa, 2015; Kutlu, 2023; Rubin, 1992). Put differently, we are not suggesting a causal or deterministic relationship between identity and language but a correlational one. Recognizing such correlations can open up important lines of inquiry. Specifically, we consider correlations between named categories of race, gender, social class, education, and other social markers used by social scientists, governments, corporations, and everyday people with features of primarily digitized written language. We are intentional with our wording here: the categories of demography are socially constructed yet remain useful in understanding trends among groups of people, particularly marginalized communities at the most risk for technological harm.

At the same time, it is essential to acknowledge that speech communities are shaped by the interconnected social systems that structure our society. For example, the likelihood that two individuals will regularly communicate and identify with each other – and thus may engage in similar speech practices to strengthen this communication and identification – may well be influenced by the geographic proximity, occupations, and social circles of these individuals, which are in turn shaped by systems such as city zoning ordinances, labor and housing markets, and the long history of redlining and other forms of systemic racism, many of which persist to the present day (Labov, 2010; Rey & Knaap, 2024; Rothstein, 2017). Insofar as these systems shape social interaction, they can also contribute to correlations between social identities and individuals' language practices (Fishman, 1997). Understanding language in this way allows us to analyze how our speech is related to grounded social conditions, and the abundance of data reflecting these social conditions has reinvigorated the study of language in society across disciplines in addition to its centrality to sociolinguistics.

Notably, a growing body of literature from many different social science disciplines has found similar relationships between written text with demography (e.g., spatial and geographic distributions of people) and the socially constructed identities used to study and discuss demography (e.g., racial and gender identities), suggesting that writers also have “accents” akin to spoken accents. Such ideas have long circulated in the arts and humanities in the form of authorship controversies (Ostrowski, 2020), and have been the subject of quantitative research in the social and computational sciences (Mosteller & Wallace, 1963; Stuhler, 2024). For example, recent work has found that the socioeconomic backgrounds and social identities of high school students applying to college are associated with the topics and word choices in their admissions essays (Alvero, 2023; Alvero *et al.*, 2020; Alvero *et al.*, 2022; Alvero & Pattichis, 2024). In fact, this association is even stronger than the association between socioeconomic background and standardized test scores (Alvero *et al.*, 2021), further highlighting the social dynamics at play in written language. Similarly, analyses of scientific abstracts and patents reveal small associations between authors' socially constructed gender identities and the rate at which certain writing features (e.g., questions, pronouns, past-tense verbs) are used (Kedrick *et al.*, 2022) or the relative innovation (measured through a bibliometric analysis of novelty in published research) of scientific ideas described in dissertation abstracts and how they are unequally taken up by the global scientific community for scholars from marginalized backgrounds (specifically women and racial minorities in academia) (Hofstra *et al.*, 2020). These discoveries should attract the attention of social scientists in general and linguists in particular.

There is no reason to believe such associations originate from inherent or essential attributes of individuals. Instead, a more parsimonious explanation for such associations is that they arise from well-documented socialization processes, such as the gendered socialization of language (Wallentin & Trecca, 2023). Furthermore, such differences tend to be small; measured variation in writing features within a given demographic category often appears to be vastly greater than measured variation between categories. For example, a recent study that found a statistically significant difference between male and female scientists' use of past tense in their writing also found that the standard deviation within each gender category was nearly 50 times larger than the mean difference between categories (Kedrick et al., 2022). Beyond the issue of within-category and between-category variation, these differences also illustrate how less obvious writing features, such as tense, align along demographic lines. We expand on this important distinction about digital accents later on in this section.

To the extent between-category variation does exist, this variation is important to investigate, because it provides an avenue for real-world social structures and power relations to influence text-based social processes. For example, human college admissions officers (or LLMs assigned to provisionally evaluate college admissions essays) may exhibit conscious or unconscious biases for and against particular essay topics or word choices, and such biases may serve as vehicles for socioeconomic class discrimination. Furthermore, some colleges have recently begun the practice of using LLMs to provisionally evaluate admissions essays (generally marketed as "AI"; for example, see this news article about the University of North Carolina: <https://www.dailytarheel.com/article/2025/01/university-admissions-essays>). Given the possibility that LLM judgments may turn out to be even more consistently biased than those of human reviewers (a trend suggested by work like Hofmann et al. [2024] about AI and dialect bias), this practice carries immense potential for harm. For example, marginalized students who share personal stories that refer to their marginalized identity or identities might face consistent and opaque forms of algorithmic bias (a not unlikely scenario given the 2023 Supreme Court ruling on affirmative action and the court's approval for students to write about "racial hardship"). While we caution against the use of LLMs in these evaluative settings, we also argue that it is imperative to critically investigate the ways in which we know LLMs are already being used, in order to (at the very least) advocate for informed regulation of such uses and in order to mitigate the biases that may be created or amplified by such uses.

We refer to associations between social identities and writing features as digital accents. We use this term, rather than alternatives such as "written accents," for two reasons. First, these patterns were largely unearthed through the analysis of large, digitized corpora using computational methods. This is why research in this area can identify subtle patterns (such as the Kedrick et al. [2022] study previously mentioned) as well as more plain, straightforward patterns (e.g., people writing about their identity and background). Second, the term "digital" can help draw attention to the possibility that some of these patterns may be, at least in part, artifacts of our increasingly digital society. For example, studies have found that geotagged social media text varies significantly and consistently by region (Eichstaedt et al., 2015; H. Huang et al., 2024; Y. Huang et al., 2016; van Loon et al., 2022); it is at least possible that the rise in both the creation of

digital text and audiences' access to such texts can contribute to intentional or inadvertent processes of language change and social or geographic differentiation (Bailey & Durham, 2021; Eisenstein *et al.*, 2014). Regardless, findings from these and other related streams of research have shown that writing is not a purely idiosyncratic process reflective of highly individualized modes of one's thoughts but instead is reflective of and intertwined with demography, geography, and social structures.

In spite of the fact that publishers and educational institutions have historically enforced the standardization of written language, it remains common for writers to deliberately encode accents, dialectal forms, regionalisms, and other purportedly "non-standardized" linguistic forms into writing. For example, stage plays like *Pygmalion* by George Bernard Shaw and novels like *Oliver Twist* by Charles Dickens use such "eye dialect" to give characters Cockney "accents" intended to denote their social location. More recently, many internet users of varied identities and backgrounds tend to use (or appropriate) features of African American Language and other marginalized language varieties online (Ilbury, 2020; Masis *et al.*, 2023). Recent research has identified approaches to quantifying and computationally analyzing such practices (see for example the book by Watson & Jensen, 2020). However, the phenomenon we refer to as digital accents is far more expansive because recent computational text analyses point to the existence of diffuse, little recognized, and less obvious "accents" in extant corpora.

Sociolinguists in the pre-digital era may have sought to use demographic variables such as ethnicity to shape their analyses – for example, by investigating the frequency with which Latina/o or non-Latina/o college applicants include Spanish words and phrases in their admissions essays. Modern computational text analyses have added important nuances to such questions by showing considerable variation across social classes and specific Latina/o identities (e.g., Mexican) in the types of Spanish words used in such essays and the ways these tend to be incorporated into otherwise English language texts (Alvero & Pattichis, 2024). Furthermore, modern linguistic analyses have found relationships between text, demography, and socially constructed identities that had not previously been hypothesized (see e.g., Eckert & Rickford, 2001), and some of these relationships likely could not have been discovered prior to the advent of computational tools. Recognizing the obvious and much less obvious ways that these patterns hold will be an important step for social scientists studying language and digital technology.

Regardless of whether digital accents are overt or subtle, surprising or stereotypical, they matter – in part because modern language technologies like LLMs can be used to identify and act upon such patterns. This was dramatically demonstrated in a 2024 study combining linguists and computer scientists which found that many of the world's most popular LLM models, when tasked with evaluating testimony from a hypothetical defendant in a first-degree murder trial, were much more likely to recommend the death penalty if the defendant's testimony exhibited (written, textual) features of African American English than if it did not (Hofmann *et al.*, 2024). Although such high-stakes uses of LLMs are rare at present, we are concerned that they may not remain so: generative AI models are reportedly already used to make life-or-death decisions in armed conflict (Adam, 2024) and are increasingly being tested for use in other decision-making processes such as college admissions (Knox, 2023). Apart

from the threat that AI may be used to actively automate social harms, even uses that appear benevolent or morally neutral carry inherent risks. Ceding human decision-making to LLMs (or, more accurately, to the small number of companies creating and maintaining them) could have frightening consequences, simultaneously consolidating social power while reducing accountability by obscuring the decision-making processes associated with that power. Understanding the ways that AI technologies can reproduce injustice and inequity will better inform efforts at both resistance and regulation, and social scientists could take the lead in analysis and description to stand out from the many voices making many claims about generative AI. As more social processes become reliant on both text and digital technology, it becomes increasingly urgent for social scientists and LLM developers to understand the relationships between written language and demographics and their implications for fairness and inequality.

Extending the linguistic maxim that “everyone has a [spoken] accent” to the idea that “everyone has a written accent” (what we refer to as “digital accents” when examined using digital means and media) has important social scientific implications. For example, when we study processes that involve text and evaluation, such as college admissions and job applications, it is important to consider the extent to which evaluators’ preferences for specific writing styles are not abstract, neutral preferences unrelated to societal inequality, but rather preferences for – or against – writing styles associated with particular demographic groups. The ideology that certain forms of writing are “standard,” unmarked, and superior to other writing styles is precisely that – an ideology shaped by societal racism, sexism, class inequality, and other societal structures (Flores & Rosa, 2015). These same structures and ideologies that reproduce societal stratification also influence the creation and interpretation of digital text, which then becomes incorporated into LLMs and other generative AI models (e.g., text-to-image, text-to-audio, and text-to-video models that use similar prompting structures like popular LLM interfaces). Connecting and comparing these trends in human text back to LLM-generated text will become an important way to understand the social ramifications of LLMs now and in the future.

Social filters of written language

Historically, literacy has been restricted to those with social and political power. While gains have been made in global literacy, published work nevertheless falls far short of reflecting the global diversity of human language. In most societies, the language of those in power is the language that is most valued and thus most likely to be reflected in written and published forms (Bourdieu, 1991). This has important implications for the development of LLMs, which are trained on corpora of what computer scientists sometimes refer to as “natural language.” This purportedly “natural” language is in fact a narrow and extensively curated subset of human language practices. These issues have been subject to scrutiny in the computational research community, such as studies of hate speech detection datasets finding that text written by Black social media users are more likely to be labeled as hate speech (Davidson et al., 2019; Sap et al., 2019). More recently, computer scientists studying peer review data found a sharp uptick in specific words and phrases, such as “delve,” pointing to “unnatural” trends likely explained by

the use of LLMs to assist in writing peer reviews (Liang *et al.*, 2024). In this section, we use prior research and empirical data to describe the social filters that limit the range of language practices eventually fed into LLMs through training data.

Consider published authors in the United States. Historically privileged demographic groups such as White Americans are over-represented among copyright holders, while other populations such as Latina/o Americans are dramatically under-represented (Brauneis & Oliar, 2018). Societal stakeholders with power over publishing and writing (e.g., publishers and educators) tend to value the voices of those with comparatively greater power, while the voices of those with comparatively lesser power are often excluded or “filtered out” from the published and written record. This is often reflected in educational practices, cultural trends, and daily life through things like the seemingly standardized way that news anchors speak. Thus, written works disproportionately reflect the language practices of those who have held some access to social power, which reinforces and replicates linguistic hierarchies. Common ideologies of language justify this filtering process by suggesting that “standard” language is fixed and homogeneous, despite a broad consensus among linguists that all living languages are actually heterogeneous and ever-changing (Horner & Weber, 2017). While it is true that LLMs are trained on as much textual data as possible, it is still the case that remnants of these social filters persist even in the text that finds its way to training data. For example, in Hofmann *et al.*’s (2024) study about dialect prejudice in LLMs, it is likely that the models had plenty of training data from social media that reflects African-American English. Despite this, the models still had clear tendencies to derogate dialectal markers in hypothetical social scenarios and attempted to filter it out in its own way. This has led to calls in the computational research community to apply critically oriented filters (rather than the purely social filters we discuss here) about the type of data being used as a means to mitigate these issues (Bender *et al.*, 2021). This ideological filtration process results in a written linguistic record that does not reflect the diversity of human language but instead reflects a narrowed ideological construct depicting what language is prescriptively imagined to be.

This filtering process is not absolute, but it is nevertheless real and empirically measurable. Here, we illustrate the filtering process using survey data from a recent mixed-methods study designed to capture US dual language immersion teachers’ beliefs about translanguaging – deviations from an imagined monolingual “standard” – in the context of classroom writing instruction in science. The survey was conducted in early 2024 among 259 teachers across the United States working in schools with Spanish–English dual language instruction programs. The survey included Likert-type items asking teachers how acceptable it was for student writers to deviate from an imagined monolingual “standard” at different stages of the writing process by using features of African American Language (Figure. 1) or by mixing features of Spanish and English in their writing (Figure. 2). We use this survey not as the central focus of our argument, but merely as a tool for illustrating the gradual filtering process which constrains linguistic diversity in written corpora.

Each figure shows how, as students move through successive stages of the writing process, teachers – even teachers accustomed to working with multilingual students and who actively support multilingualism in their work – appear to become increasingly uncomfortable with deviations from a purported monolingual “standard”

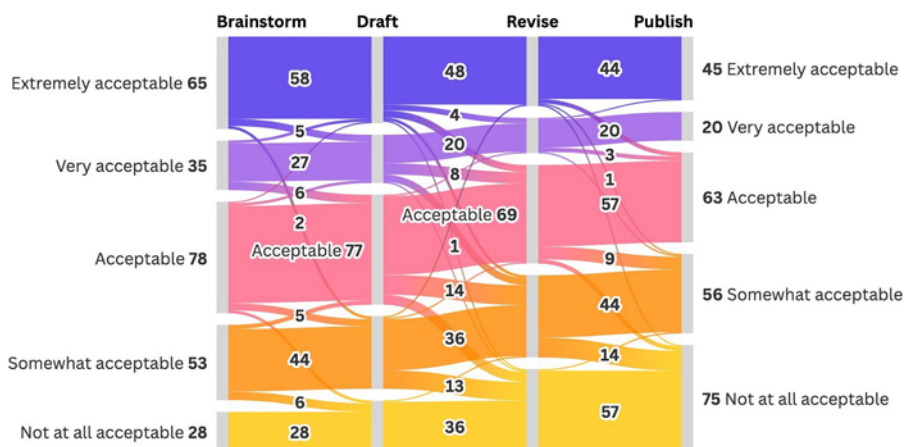


Figure 1. In your classroom, how acceptable is it for students to use African American English or other “non-standardized” dialects or varieties to [do each of the following tasks] during the writing process?




Figure 2. In your classroom, how acceptable is it for students to mix languages to [do each of the following tasks] during the writing process (e.g., to use Spanish when English is the language of instruction or vice versa)?

norm. For example, in Figure 1, 65 of the 259 participating teachers felt it was “extremely acceptable” for students to use African American English during the initial (Brainstorming) phase of the writing process; however, when advancing to the second (Drafting) phase, only 58 teachers still felt this was extremely acceptable, while 5 of the original 65 now felt it was only “very acceptable” and 2 now felt it was only “somewhat acceptable.” At each stage of the writing process, as students move closer to the final (Publishing) phase, fewer and fewer teachers tended to see African American English as “extremely acceptable” in students’ writing, while more and more teachers tended to

see it as “not at all acceptable” in students’ writing. Patterns in teachers’ acceptance of mixing Spanish and English are similar (see [Figure. 2](#)).

This filtering process drives student writing toward a norm that is more homogeneous than the breadth of students’ everyday language practices. Such filtering effects occur throughout society as schools, journalistic organizations, book publishers, and other gatekeepers continually enforce and reinforce homogenization toward an imagined prescriptive linguistic norm. The texts created and disseminated through schools, journalism, and publishing are then fed into LLMs as training data which has itself already undergone this filtering process. LLMs are then likely to mirror and perpetuate this narrow range of linguistic practices, a range of practices disproportionately shaped by the voices of those with privileged access to technology and publication opportunities. Just as written language in educational contexts is the outcome of “filtering” that marginalizes diverse linguistic expressions, LLMs build upon and arguably reinforce this narrowing, amplifying language practices associated with social elites while excluding and further marginalizing linguistic diversity. Their broad use, from personal creative projects to daily processes inside large corporations, also likely contributes to patterns of linguistic hegemony.

Ontological tensions with the social science of language

A careful reader well-versed in ontological debates about the nature of language might take issue with our claims about digital accents and whether they are naturally occurring rather than social constructs. Language scholars have long debated the ontology of language and how we, as social scientists, impose categories onto practices that are famously resistant to boundaries (Horner & Weber, 2017). The reification of social and linguistic categories, even when done for overtly anti-oppressive purposes, can still cause unintended harm (Bucholtz, 2003; King, 2020; Morgan, 1994). The rise of generative AI complicates these debates even further, posing new challenges to the ways we define and differentiate among named languages and language varieties. Generative AI tools will readily act upon users’ requests to differentiate between language varieties: for example, when a user prompts an LLM to generate text that reflects a particular dialect or variety, such as Dominican Spanish (e.g., https://www.reddit.com/r/Dominican/comments/14z411x/asking_chatgpt_to_write_in_dominican_spanish/; see Hinojos, 2023), it will produce text that reflects lexical patterns reminiscent of that particular style (the same is true for audio and speech). While reifying social and linguistic categories can lead to harmful stereotyping or over-simplification, we nevertheless argue that associations between demographic categories and written text, even small and relatively insignificant ones, inherently contribute to situations in which generative AI tools implicitly or explicitly reify social and linguistic categories. It is therefore essential for linguists and computer scientists to grapple with the tensions between social patterning in writing and language and critical issues about the nature of categorization.

Other social sciences have grappled with similar tensions. For example, the sociology of race and ethnicity has helped thoroughly debunk the idea that racial identity is intrinsic, stable, essential, or grounded in phenotypic features such as skin tone. Rather, race is widely understood to be a social construct (Roth *et al.*, 2023). At the

same time, sociologists studying race have provided strong evidence that skin tone is nevertheless related to social outcomes (e.g., discrimination, income stratification) because of the mediating role that perceptions of skin tone play in mechanisms of systemic racism (Adames, 2023; Keith & Herring, 1991; Monk, 2015). Importantly, these studies are not conducted because their authors believe or seek to prove that race is an essential or meaningful construct grounded in skin tone, but instead because widely held ideologies that associate racial categories with skin tone can help explain the way certain forms of discrimination operate in our society, and because there is meaning and value in understanding such mechanisms in order to more effectively disrupt or address them.

As generative AI becomes more pervasive in modern life, language scholars must contend with similar tensions in their work. The development and use of LLMs heavily relies on discrete language categories, something that can be easily observed in its marketing (not least as a machine translation tool). In supervised machine learning, including the deep learning and transformer architectures that generative AI is built with, the statistical objective for each model is to label data along a set of discrete categories correctly: this is often binary (e.g., this text is English or not English) but also multinomial (e.g., predict which language this text comes from a set of 100 possibilities). The same is directly true for generative AI: predicting the next token is no different than the objectives of prior machine learning algorithms for text even though the scenario presents many more possibilities. It would be challenging for language scholars to carry out empirical work without acknowledging how language categorization shapes generative AI (even if languages are socially constructed); the same is true for studying the relationships between LLM output vis-à-vis written language features that are statistically associated (even to a small degree) with demographic categories.

Studying language and how it operates in society, even when it conflicts with our intellectual commitments and theories, must be central to our efforts to make thoughtful, critical contributions to public knowledge and discourse (see this statement from the Linguistic Society of America for one example: https://www.lsadc.org/statement_on_race; Hudley et al., 2024; Ramjattan, 2022). Sociologists and psychologists studying race and ethnicity have addressed similar tensions, using theoretical lenses such as color-evasive theory (also known in the literature as “colorblind racism” (Annamma et al., 2017; Bonilla-Silva, 2021), to show how systematically ignoring (or professing to ignore) differences between social categories in order to avoid reifying such categories can have its own unintended and potentially harmful consequences (Bonilla-Silva, 2021; Thomas et al., 2023). For linguists, adopting a similar stance might require more attention toward situated perspectives that incorporate structural and systemic considerations typically favored by sociolinguistics research. Primarily behavioral and cognitive theorizations of language are also helpful and important, but complementing them with sociolinguistic perspectives could generate new inroads. Doing so may also help clarify correlational and causal claims about the relationships with language, technology, and demography.

Homogeneity-by-design

While it is true that there is broad interest in generative AI among academics, it is also true that this technology was not necessarily designed with social scientific applications in mind. A significant aspect of the popularity of generative AI is its potential use in business, such as enterprise software, and many of the ideas about what AI can and should do come from those perspectives. In theory, when used by skilled workers, generative AI could make tasks more straightforward; when used by executives and high-level decision-makers, it could be used as a way to cut down on employment costs (i.e., job displacement and automation). These intended applications also extend to explicitly shaping or reshaping language use in business settings; for example, generative AI is actively being used by call centers to alter the spoken accents of workers to sound more appealing to American consumers with linguistic biases against accents associated with marginalized identities (Payne *et al.*, 2024; Ramjattan, 2019). Essentially, generative AI is not designed to represent the diversity of human language; it is designed to be (or at least appear to be) relatively homogeneous, plausibly as part of a business strategy. Reminding ourselves that generative AI and LLMs are primarily business tools designed for profit rather than scientific, cultural, or social tools can be useful framing for understanding why homogeneity would be intentionally designed into these systems.

Recent studies have shown that, compared to humans, LLMs have relatively little variation in their output, especially compared to patterns in human writing (Agarwal *et al.*, 2024; Alvero *et al.*, 2024; Anderson *et al.*, 2024; Moon *et al.*, 2024; Zhang *et al.*, 2025). While there may be technical explanations for this, related to modeling decisions or dimensions of the training data, we argue that AI homogenization is also an organizational decision on the part of emergent AI companies. Past chatbots that were also trained on massive amounts of data scraped from the internet demonstrated much more variation and malleability in their output, such as the infamous example of Microsoft's chatbot Tay that was easily manipulated by online trolls to generate sexist, racist responses to user interactions (Neff & Nagy, 2016). Modern AI companies have been able to learn from these mistakes in order to develop technology they can credibly market as consistent, low variance, and somewhat predictable, despite working in a highly inconsistent, high variance, and unpredictable medium (human language). Others have described analogous trends elsewhere in society as "McDonaldsization," when social processes and contexts become more consistent and predictable (Ritzer, 2021). Generative AI is similar: no matter who is creating a prompt, and no matter what small idiosyncrasies their prompt might include, they will likely get similar responses (i.e., homogeneity-by-design).

Computer scientists are trained to consider how any user, regardless of their background, can use a given piece of hardware or technology (e.g., see the ACM/IEEE Computer Society Joint Task Force on Computing Curricula, 2013). The trouble is that this ideal is often pushed aside. In theory, designing with homogeneity in mind might help ensure a broad range of users can effectively engage with any computational technology; however, there are many instances where this ideal does not pass muster (Benjamin, 2019; Haraway, 2013). The potential issues with homogeneity-by-design become further apparent once we consider the many social scenarios where

linguistic variation is highly stigmatized and/or suppressed. Schools, especially those with colonial histories or embedded in colonial or post-colonial contexts, are well-known sites where linguistic variation manifests and is stigmatized (Mallinson, 2024). Introducing technology that is deliberately designed to homogenize language has the potential to produce new, digitally mediated versions of linguistic stratification and inequality.

Critically, the existence of digital accents means that homogeneity-by-design can reproduce societal inequality in innumerable ways, some of which may be subtle and hard for researchers to detect. If there are writing styles, pragmatics, and ways of communicating that are correlated with human demographic and geographic categories, then AI-generated text will likely exhibit an affinity toward approximating the styles and voices of particular groups, with past studies showing affinity toward high-income, males (Alvero et al., 2024) or those from Protestant Western countries (Tao et al., 2024). Taken one step further, the burgeoning AI homogenization literature has also found that even when compared to the language use of these privileged demographic groups, AI-generated text exhibits less variation than human writers (Zhang et al., *in press*). Our current social systems, schools, businesses, and other organizations engage in similar homogenization processes – including the filtering described above – in ways that reproduce social inequality and stratification. Since generative AI also operates in this way (at least partly by design), it is reasonable to expect that AI, too, could reproduce social inequality and stratification.

The evolving social science of written language in a world of generative AI

Just as race is a social construct with concrete, real-world material consequences, language categories are social constructs with concrete, real-world material consequences. While humans would use language regardless of whether or not those languages were named (unlike race), the ways that languages are named, sorted, stratified, and classified make them amenable to processes where people (i.e., speakers of a given language) are likewise sorted and stratified based on language (like race and racial categories). For example, accent- or dialect-focused biases in spoken language serve as mechanisms for enacting racist and/or nativist discrimination in hiring (Schulte et al., 2024), housing (Wright, 2023), service industry interactions (Wang et al., 2013), and courtroom testimony (Romero-Rivas et al., 2022). These same ideologies and biases shape written language through the filters described above, and such filters shape LLMs as well. There is an urgent need for further research by social scientists of language to explore how specific language ideologies and varied linguistic biases may have already been incorporated into LLMs in unanticipated ways.

It also suggests an urgent need to examine the language ideologies of the populations responsible for developing, training, and aligning LLMs: computer scientists in general, and generative AI engineers and developers in particular. In the short term, understanding the language ideologies of computer scientists could help researchers identify and disrupt the types of bias most likely to have been incorporated into the design of LLMs. In the longer term, understanding the language ideologies of computer scientists could help inform the design of reforms in computer science education to address such ideologies proactively.

To date, some research has begun to explore the language ideologies encoded into LLMs and the language ideologies of their programmers. This work has found trends such as the relationship between what is considered “high quality” text for training a model tending to come from socioeconomically privileged communities or the tendency for LLMs to misclassify non-standard dialects (Gururangan *et al.*, 2022; Höhn, *et al.*, 2023; Smith *et al.*, 2024). Further quantitative and qualitative work is needed to understand how the language ideologies of the broader society shape those of computer scientists; whether unique language ideologies (or unique mechanisms that reproduce these ideologies) operate in computer science education and in software companies; and how different types of educational, regulatory, or grassroots interventions might disrupt oppressive language ideologies and their effects on LLM development and alignment.

Additional research is needed to explore how LLMs reproduce (or disrupt) the language ideologies of their users. LLMs generate text (and, increasingly, audio) for millions of people worldwide daily. Thus, they represent a unique and powerful platform where language ideologies can be reproduced; if the default voice of ChatGPT, for example, is a voice that invariably writes and speaks Standardized American English (or rather, a variety perceived and described by users as Standardized American English), this may reinforce the perception of Standardized American English as an unmarked “voice from nowhere,” when in reality this voice has been shaped by innumerable ideological filters driven by racism, colonialism, and other oppressive systems.

The increasing ubiquity of LLMs may also make them a uniquely powerful platform for disrupting oppressive language ideologies. Such disruption will require a critical analysis of language ideologies and the underlying ontological assumptions of LLM development. For example, LLM developers have recently begun training models to produce multiple languages or language varieties. ChatGPT’s new audio interface can create a variety of different accents and dialects (e.g., English with a Dominican accent) upon request and in real time (see <https://www.theverge.com/2024/8/1/24211087/openai-chatgpt-advanced-voice-mode-demo-accents-language>). Taken at face value, some commentators might argue this is an important step toward solving the problem of homogenization and disrupting standard language ideology because it enables ChatGPT to replace a purported Standardized American English with an alternative (and often marginalized) language variety.

However, there are two problems with this argument. First, incrementally adding new accents or dialects to an LLM’s repertoire still conceptualizes the world’s languages as a finite set of clearly differentiated, internally homogeneous types – an ontological perspective long ago rejected by linguists, who recognize that variation is a typical feature of all living languages and that drawing distinctions between languages is a subjective and political process rather than an empirical one (Horner & Weber, 2017). ChatGPT’s “Dominican accent” may represent a probabilistic average of linguistic features associated with the label “Dominican” in its training corpora. Still, this average no more conveys the heterogeneity of speech among the 11 million residents of the Dominican Republic than ChatGPT’s default American voice conveys the heterogeneity of speech among 340 million US residents (to say nothing of the more than 2 million US residents who identify as Dominican). Furthermore, such LLM-generated accents and dialects will always represent a snapshot of how language varieties existed in the

past, heavily “filtered” into corpora of training data, rather than representing how language varieties exist, vary, and dynamically change in the present. Thus, rather than solving the homogenization problem, using LLMs to mimic particular accents, dialects, or varieties simply reproduces the problem of homogenization on a smaller scale.

Furthermore, using LLMs to “mimic” particular accents and dialects at the user’s request is fundamentally unlikely to disrupt standard language ideology. To reach this prediction, we conceptualize standard language ideology as a collection of inter-related stereotypes (associating a particular language variety with qualities such as “appropriateness” and identities such as “White”) and draw upon the extensive social psychological research on stereotype change. We recognize that this differs from many other scholars’ approaches to conceptualizing language ideologies, which often transcend individual cognition to include societal discourses, practices, policies, or artifacts that instantiate language ideologies in the physical world (Flores & Rosa, 2015; Horner & Weber, 2017). Here, however, we apply insights from social psychology to narrowly examine those components of language ideologies which do reside in the cognition of individuals.

Research on stereotype change demonstrates that exposure to one (or a few) strongly stereotype-disconfirming examples does not result in stereotype change but merely creates new mental subcategories to accommodate new examples (Gershman & Cikara, 2023). In the case of LLMs, we predict users are likely to perceive features such as ChatGPT’s “Dominican accent” as fundamentally separate from the “real” voice of ChatGPT, which will likely remain Standardized American English in the public imagination. Thus, we predict LLM users’ standard language ideologies will be unaffected by the availability or even the use of these novel accent- or dialect-mimicking features, and we encourage future researchers to test this prediction.

On the other hand, research on stereotype change also demonstrates that exposure to many, weakly stereotype-disconfirming examples is more likely to create lasting change in stereotypes (Gershman & Cikara, 2023). Thus, there may be conditions under which LLMs can promote change in linguistic stereotypes or broader language ideologies. For example, imagine if an LLM were redesigned to spontaneously exhibit many slightly different “default” voices that incorporate different accents, dialects, or varieties of English. Every voice is from somewhere, and if an LLM randomly alternated among such voices over time, this might help to disrupt users’ belief in a single standardized “voice from nowhere” – and by extension, may disrupt standard language ideology among users. This intervention would require the combined efforts of applied linguists and computer scientists to test; if successful, such practices could subtly affect the day-to-day activities of millions of users and could play an important role in disrupting longstanding language biases rooted in racism and colonialism.

Of course, one important critique of such an intervention is its potential for appropriation and for reinforcing stereotypes. In order to avoid these pitfalls, LLM developers would likely need to draw upon consenting individuals’ own language practices (not from non-consenting individuals’ data contained in existing corpora, nor writers’ stereotypical representations of others’ language practices). LLM developers would further need to ensure just compensation for the people whose language practices are incorporated to train models. We could imagine an LLM structured as a public

(rather than private) good: an opt-in LLM that periodically retrains itself on samples of user-generated input from specific, randomly chosen geographic locations – basing its own linguistic variation not on preexisting human social beliefs that name and differentiate language varieties based on political ideologies or racial stereotypes, but instead on the empirically measured language practices of specific communities. On Monday, the English-language version of this LLM might write or speak like a person from London; on Tuesday, Miami; on Wednesday, Kentucky; on Thursday, El Paso; on Friday, New Delhi. While this would still not completely capture the dynamic nature and variability of human language, it might help mitigate the homogenization problem that seems likely to result from current approaches to LLM design grounded in standard language ideology. Even if this idea were never to be implemented, imagining alternatives for the ways that LLMs communicate is a helpful exercise that social scientists could use for theoretical development and scientific communication.

This approach to addressing the homogenization problem could have other added benefits from an ethical and moral standpoint. The current practice of training LLMs to imitate particular accents and dialects arguably represents a problematic form of appropriation; these new applications enable corporations to commodify marginalized language practices while doing nothing to help populations marginalized by linguistic oppression directly. Of course, this form of linguistic appropriation long predates the rise of LLMs; see, for example, Roth-Gordon *et al.*'s (2020) discussion of the commodification of African American Language by advertisers. Linguists have long argued that it is essential for communities that generate language data to benefit directly from research or development that uses these data (Hudley *et al.*, 2022; Kouritzin & Nakagawa, 2018). Periodically re-training models based on the input of users from randomly selected geographic areas, as we have suggested here, could form a basis for creating reciprocal, mutually beneficial relationships between LLMs and the communities that use them. Users whose data are randomly selected for model re-training could receive financial compensation from the corporation(s) managing such LLMs, operationalizing principles of economic justice while also incentivizing greater consumer use of LLMs. Generative AI could function as a public good, not merely a private one.

We do not wish to promote an idyllic or overly credulous view of generative AI, which can and does readily function as a powerful tool for harm through government and corporate surveillance, cultural and linguistic appropriation, and the reproduction of harmful biases and ideologies. At the same time, we worry about the prospect of social scientists (and applied linguists in particular) eschewing or abandoning work in generative AI because of its potential for harm when their voices and imagination will be fundamental to ensuring harm prevention, reduction, and, ideally, elimination (Bender & Grissom II, 2024). The rise of generative AI brings significant challenges and important questions with very real, concrete implications for the day-to-day lives of millions of individuals worldwide. These are challenges and questions that social scientists in general, and applied linguists in particular, are uniquely positioned to help address; our work can inform both efforts to resist or regulate harmful uses of AI and efforts to promote ethical and just uses. If there is ever to be a version of generative AI that avoids harm and advances equity, its development will require the voices and insights of many – including those who have dedicated their careers to understanding language through critical and applied perspectives.

Acknowledgments. The authors would like to thank Haley Lepp for helping this paper come together. We also thank the anonymous reviewers for their helpful comments and suggestions. The survey work presented in this paper was funded by a performance-based funding grant from the Annette Caldwell Simmons School of Education and Human Development at Southern Methodist University.

References

- Adam, D. (2024). Lethal AI weapons are here: How can we control them? *Nature*, 629(8012), 521–523. <https://doi.org/10.1038/d41586-024-01029-0>
- Adames, A. (2023). The cumulative effects of colorism: Race, wealth, and skin tone. *Social Forces*, 102(2), 539–560. <https://doi.org/10.1093/sf/soad038>
- Agarwal, D., Naaman, M., & Vashistha, A. (2024). AI suggestions homogenize writing toward western styles and diminish cultural nuances. *arXiv preprint arXiv:2409.11360*. <https://doi.org/10.48550/arXiv.2409.11360>
- Alvero, A. (2023). Sociolinguistic perspectives on machine learning with text data. In C. Borch & J. P. Pardo-Guerra (Eds.), *The Oxford handbook of the sociology of machine learning* (pp. 79–97). Oxford University Press. <https://doi.org/10.1093/oxfordhdb/9780197653609.013.15>
- Alvero, A., Arthurs, N., Antonio, A. L., Domingue, B. W., Gebre-Medhin, B., Giebel, S., & Stevens, M. L. (2020). AI and holistic review: Informing human reading in college admissions. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society* (pp. 200–206). Association for Computing Machinery, New York, NY, United States. <https://doi.org/10.1145/3375627.3375871>
- Alvero, A., Giebel, S., Gebre-Medhin, B., Antonio, A. L., Stevens, M. L., & Domingue, B. W. (2021). Essay content and style are strongly related to household income and sat scores: Evidence from 60,000 undergraduate applications. *Science Advances*, 7(42), 1–10. <https://doi.org/10.1126/sciadv.abi9031>
- Alvero, A., Lee, J., Regla-Vargas, A., Kizilcec, R. F., Joachims, T., & Antonio, A. L. (2024). Large language models, social demography, and hegemony: Comparing authorship in human and synthetic text. *Journal of Big Data*, 11(1), 138. <https://doi.org/10.1186/s40537-024-00986-7>
- Alvero, A., Pal, J., & Moussavian, K. M. (2022). Linguistic, cultural, and narrative capital: Computational and human readings of transfer admissions essays. *Journal of Computational Social Science*, 5(2), 1709–1734. <https://doi.org/10.1007/s42001-022-00185-5>
- Alvero, A., & Pattichis, R. (2024). Multilingualism and mismatching: Spanish language usage in college admissions essays. *Poetics*, 105, 101903. <https://doi.org/10.1016/j.poetic.2024.101903>
- Anderson, B. R., Shah, J. H., & Kreminski, M. (2024). Homogenization effects of large language models on human creative ideation. In *Proceedings of the 16th Conference on Creativity & Cognition* (pp. 413–425). Association for Computing Machinery. <https://doi.org/10.1145/3635636.3656204>
- Annamma, S. A., Jackson, D. D., & Morrison, D. (2017). Conceptualizing color-evasiveness: Using dis/ability critical race theory to expand a color-blind racial ideology in education and society. *Race Ethnicity and Education*, 20(2), 147–162. <https://doi.org/10.1080/13613324.2016.1248837>
- Bailey, L. R., & Durham, M. (2021). A cheeky investigation: Tracking the semantic change of cheeky from monkeys to wines: Can social media spread linguistic change? *English Today*, 37(4), 214–223. <https://doi.org/10.1017/S0266078420000073>
- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM conference on Fairness, Accountability, and Transparency (FACt'21)* (pp. 610–623). Association for Computing Machinery. <https://doi.org/10.1145/3442188.3445922>
- Bender, E. M., & Grissom, A., II. (2024). Power shift: Toward inclusive natural language processing. In A. H. C. Hudley, C. Mallinson & M. Bucholtz (Eds.), *Inclusion in linguistics* (pp. 199–224). Oxford University Press. <https://doi.org/10.1093/oso/9780197755303.003.0010>
- Benjamin, R. (2019). *Race after technology: Abolitionist tools for the new Jim code*. Polity Press.
- Bonilla-Silva, E. (2021). *Racism without racists: Color-blind racism and the persistence of racial inequality in America*. Rowman & Littlefield.
- Bourdieu, P. (1991). *Language and symbolic power*. Harvard University Press.
- Brauneis, R., & Oliar, D. (2018). An empirical study of the race, ethnicity, gender, and age of copyright registrants. *George Washington Law Review*, 86(1), 46–98.

- Bucholtz, M. (2003). Sociolinguistic nostalgia and the authentication of identity. *Journal of Sociolinguistics*, 7(3), 398–416. <https://doi.org/10.1111/1467-9481.00232>
- Davidson, T., Bhattacharya, D., & Weber, I. (2019). Racial bias in hate speech and abusive language detection datasets. In *Proceedings of the Third Workshop on Abusive Language Online* (pp. 25–35). Association for Computational Linguistics. <https://doi.org/10.18653/v1/W19-3504>
- Eckert, P., & Rickford, J. R., Eds. (2001). *Style and sociolinguistic variation*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511613258>
- Eichstaedt, J. C., Schwartz, H. A., Kern, M. L., Park, G., Labarthe, D. R., Merchant, R. M., ... Seligman, M. E. (2015). Psychological language on Twitter predicts county-level heart disease mortality. *Psychological Science*, 26(2), 159–169. <https://doi.org/10.1177/0956797614557867>
- Eisenstein, J., O'Connor, B., Smith, N. A., & Xing, E. P. (2014). Diffusion of lexical change in social media. *PLoS One*, 9(11), e113114. <https://doi.org/10.1371/journal.pone.0113114>
- Fishman, J. A. (1997). The sociology of language. In N. Coupland & A. Jaworski (Eds.), *Sociolinguistics: A reader and coursebook* (pp. 25–30). Macmillan Press. https://doi.org/10.1007/978-1-349-25582-5_4
- Flores, N., & Rosa, J. (2015). Undoing appropriateness: Raciolinguistic ideologies and language diversity in education. *Harvard Educational Review*, 85(2), 149–171. <https://doi.org/10.17763/0017-8055.85.2.149>
- Gershman, S. J., & Cikara, M. (2023). Structure learning principles of stereotype change. *Psychonomic Bulletin and Review*, 30(4), 1273–1293. <https://doi.org/10.3758/s13423-023-02252-y>
- Gururangan, S., Card, D., Dreier, S. K., Gade, E. K., Wang, L. Z., Wang, Z., ... Smith, N. A. (2022). Whose language counts as high quality? Measuring language ideologies in text data selection. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing* (pp. 2562–2580). Association for Computational Linguistics. <https://doi.org/10.18653/v1/2022.emnlp-main.165>
- Haraway, D. (2013). A cyborg manifesto: Science, technology, and socialist-feminism in the late twentieth century. In Stryker, S & Whittle, S. (Eds.), *The transgender studies reader* (pp. 103–118). Routledge. https://doi.org/10.1007/978-1-4020-3803-7_4
- Hinojos, S. (2023). 4 “Accented” Latinx Textese: Bilingual scriptural economies and digital literacies. In P. Rangan (Ed.), *Thinking with an accent: Toward a new object, method, and practice* (pp. 73–94). University of California Press. <https://doi.org/10.1525/9780520389748-009>
- Hofmann, V., Kalluri, P. R., Jurafsky, D., & King, S. (2024). AI generates covertly racist decisions about people based on their dialect. *Nature*, 633(8028), 147–154. <https://doi.org/10.1038/s41586-024-07856-5>
- Hofstra, B., Kulkarni, V. V., Munoz-Najar Galvez, S., He, B., Jurafsky, D., & McFarland, D. A. (2020). The diversity–innovation paradox in science. *Proceedings of the National Academy of Sciences*, 117(17), 9284–9291. <https://doi.org/10.1073/pnas.1915378117>
- Höhn, S., Migge, B., Dippold, D., Schneider, B., & Mauw, S. (2023). Language ideology bias in conversational technology. In *International Workshop on Chatbot Research and Design* (pp. 133–148). Springer. https://doi.org/10.1007/978-3-031-54975-5_8
- Horner, K., & Weber, J. J. (2017). *Introducing multilingualism: A social approach*. Routledge. <https://doi.org/10.4324/9781315276892>
- Huang, H., Grieve, J., Jiao, L., & Cai, Z. (2024). Geographic structure of Chinese dialects: A computational dialectometric approach. *Linguistics*, 62(4), 937–976. <https://doi.org/10.1515/ling-2021-0138>
- Huang, Y., Guo, D., Kasakoff, A., & Grieve, J. (2016). Understanding US regional linguistic variation with Twitter data analysis. *Computers, Environment and Urban Systems*, 59, 244–255. <https://doi.org/10.1016/j.compenvurbysys.2015.12.003>
- Hudley, A. H. C., Mallinson, C., & Bucholtz, M. (2022). *Talking college: Making space for black language practices in higher education*. Teachers College Press.
- Hudley, A. H. C., Mallinson, C., & Bucholtz, M. (2024). *Inclusion in linguistics*. Oxford University Press. <https://doi.org/10.1093/oso/9780197755393.001.0001>
- Ilbury, C. (2020). “Sassy queens”: Stylistic orthographic variation in Twitter and the enregisterment of AAVE. *Journal of Sociolinguistics*, 24(2), 245–264. <https://doi.org/10.1111/josl.12366>
- Joint Task Force on Computing Curricula, Association for Computing Machinery (ACM), & IEEE Computer Society. (2013). *Computer science curricula 2013: Curriculum guidelines for undergraduate degree programs in computer science*. Association for Computing Machinery. <https://doi.org/10.1145/2534860>
- Kedrick, K., Levitskaya, E., & Funk, R. J. (2022). Investigating writing style as a contributor to gender gaps in science and technology. *arXiv preprint arXiv:2204.13805*.

- Keith, V. M., & Herring, C. (1991). Skin tone and stratification in the black community. *American Journal of Sociology*, 97(3), 760–778. <https://doi.org/10.1086/229819>
- King, S. (2020). From African American vernacular English to African American language: Rethinking the study of race and language in African Americans' speech. *Annual Review of Linguistics*, 6(1), 285–300. <https://doi.org/10.1146/annurev-linguistics-011619-030556>
- Knox, L. (2023). Admissions offices deploy AI. *Inside Higher Education*. Retrieved January 30, 2025 from <https://www.insidehighered.com/news/admissions/traditional-age/2023/10/09/admissions-offices-turn-ai-application-reviews>
- Kouritzin, S., & Nakagawa, S. (2018). Toward a non-extractive research ethics for transcultural, translingual research: Perspectives from the coloniser and the colonised. *Journal of Multilingual and Multicultural Development*, 39(8), 675–687. <https://doi.org/10.1080/01434632.2018.1427755>
- Kutlu, E. (2023). Now you see me, now you mishear me: Raciolinguistic accounts of speech perception in different English varieties. *Journal of Multilingual and Multicultural Development*, 44(6), 511–525. <https://doi.org/10.1080/01434632.2020.1835929>
- Labov, W. (2010). Unendangered dialect, endangered people: The case of African American vernacular English. *Transforming Anthropology*, 18(1), 15–27. <https://doi.org/10.1111/j.1548-7466.2010.01066.x>
- Liang, W., Izzo, Z., Zhang, Y., Lepp, H., Cao, H., Zhao, X., ... Zou, J. Y. (2024). Monitoring AI-modified content at scale. A case study on the impact of ChatGPT on AI conference peer reviews. *Proceedings of the 41st International Conference on Machine Learning*, 235, 29575–29620. JMLR.org.
- Mallinson, C. (2024). Linguistic variation and linguistic inclusion in the us educational context. *Annual Review of Linguistics*, 10(1), 37–57. <https://doi.org/10.1146/annurev-linguistics-031120-121546>
- Masis, T., Eggleston, C., Green, L. J., Jones, T., Armstrong, M., & O'Connor, B. (2023). Investigating morphosyntactic variation in African American English on Twitter. *Society for Computation in Linguistics*, 6(1), 392–393. <https://doi.org/10.7275/zdgo-0914>
- Monk, E. P. Jr. (2015). The cost of color: Skin color, discrimination, and health among African-Americans. *American Journal of Sociology*, 121(2), 396–444. <https://doi.org/10.1086/682162>
- Moon, K., Green, A., & Kushlev, K. (2024). Homogenizing effect of a large language model (LLM) on creative diversity: An empirical comparison of human and ChatGPT writing. *PsyArXiv*. <https://doi.org/10.31234/osf.io/8p9wu>
- Morgan, M. (1994). Theories and politics in African American English. *Annual Review of Anthropology*, 23(1), 325–345. <https://doi.org/10.1146/annurev.an.23.100194.001545>
- Mosteller, F., & Wallace, D. L. (1963). Inference in an authorship problem: A comparative study of discrimination methods applied to the authorship of the disputed federalist papers. *Journal of the American Statistical Association*, 58(302), 275–309. <https://doi.org/10.1080/01621459.1963.10500849>
- Neff, G., & Nagy, P. (2016). Automation, algorithms, and politics| Talking to bots: Symbiotic agency and the case of Tay. *International Journal of Communication*, 10, 4915–4931.
- Ostrowski, D. (2020). *Who wrote that? Authorship controversies from Moses to Sholokhov*. Cornell University Press. <https://doi.org/10.1515/9781501749728>
- Payne, A. L., Austin, T., & Clemons, A. M. (2024). Beyond the front yard: The dehumanizing message of accent-altering technology. *Applied Linguistics*, 45(3), 553–560. <https://doi.org/10.1093/applin/amae002>
- Porter, G. (2012). Mobile phones, livelihoods and the poor in sub-Saharan Africa: Review and prospect. *Geography Compass*, 6(5), 241–259. <https://doi.org/10.1111/j.1749-8198.2012.00484.x>
- Ramjattan, V. A. (2019). Raciolinguistics and the aesthetic labourer. *Journal of Industrial Relations*, 61(5), 726–738. <https://doi.org/10.1177/0022185618792990>
- Ramjattan, V. A. (2022). Accenting racism in labour migration. *Annual Review of Applied Linguistics*, 42, 87–92. <https://doi.org/10.1017/S0267190521000143>
- Rey, S. J., & Knaap, E. (2024). The legacy of redlining: A spatial dynamics perspective. *International Regional Science Review*, 47(1), 3–44. <https://doi.org/10.1177/01600176221116566>
- Ritzer, G. (2021). The McDonaldization of society. In Mofield, E & Stambaugh, T (Eds.), *In the mind's eye* (pp. 143–152). Routledge.
- Romero-Rivas, C., Morgan, C., & Collier, T. (2022). Accentism on trial: Categorization/stereotyping and implicit biases predict harsher sentences for foreign-accented defendants. *Journal of Language and Social Psychology*, 41(2), 191–208. <https://doi.org/10.1177/0261927X211022785>

- Roth, W. D., van Stee, E. G., & Regla-Vargas, A. (2023). Conceptualizations of race: Essentialism and constructivism. *Annual Review of Sociology*, 49(1), 39–58. <https://doi.org/10.1146/annurev-soc-031021-034017>
- Roth-Gordon, J., Harris, J., & Zamora, S. (2020). Producing white comfort through “corporate cool”: Linguistic appropriation, social media, and @BrandsayingBae. *International Journal of the Sociology of Language*, 2020(265), 107–128. <https://doi.org/10.1515/ijsl-2020-2105>
- Rothstein, R. (2017). *The color of law: A forgotten history of how our government segregated America*. Liveright Publishing.
- Rubin, D. L. (1992). Nonlanguage factors affecting undergraduates’ judgments of nonnative English-speaking teaching assistants. *Research in Higher Education*, 33(4), 511–531. <https://doi.org/10.1007/BF00973770>
- Sap, M., Card, D., Gabriel, S., Choi, Y., & Smith, N. A. (2019). The risk of racial bias in hate speech detection. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics* (pp. 1668–1678). Association for Computational Linguistics. <https://doi.org/10.18653/v1/P19-1163>
- Schulte, N., Basch, J. M., Hay, H.-S., & Melchers, K. G. (2024). Do ethnic, migration-based, and regional language varieties put applicants at a disadvantage? A meta-analysis of biases in personnel selection. *Applied Psychology*, 73(4), 1866–1892. <https://doi.org/10.1111/apps.12528>
- Smith, G., Fleisig, E., Bossi, M., Rustagi, I., & Yin, X. (2024). Standard language ideology in AI-generated language. *arXiv preprint arXiv:2406.08726*. <https://doi.org/10.48550/arXiv.2406.08726>
- Stuhler, O. (2024). The gender agency gap in fiction writing (1850 to 2010). *Proceedings of the National Academy of Sciences*, 121(29), e2319514121. <https://doi.org/10.1073/pnas.2319514121>
- Tao, Y., Viberg, O., Baker, R. S., & Kizilcec, R. F. (2024). Cultural bias and cultural alignment of large language models. *PNAS Nexus*, 3(9), 346. <https://doi.org/10.1093/pnasnexus/pgae346>
- Thomas, A. K., McKinney de Royston, M., & Powell, S. (2023). Color-evasive cognition: The unavoidable impact of scientific racism in the founding of a field. *Current Directions in Psychological Science*, 32(2), 137–144. <https://doi.org/10.1177/09637214221141713>
- van Loon, A., Giorgi, S., Willer, R., & Eichstaedt, J. (2022). Negative associations in word embeddings predict anti-black bias across regions – But only via name frequency. *Proceedings of the International AAAI Conference on Web and Social Media*, 16, 1419–1424. <https://doi.org/10.1609/icwsm.v16i1.19399>
- Wallentin, M., & Trecca, F. (2023). Cross-cultural sex/gender differences in produced word content before the age of 3 years. *Psychological Science*, 34(4), 411–423. <https://doi.org/10.1177/09567976221146537>
- Wang, Z., Arndt, A. D., Singh, S. N., Biernat, M., & Liu, F. (2013). “You lost me at hello”: How and when accent-based biases are expressed and suppressed. *International Journal of Research in Marketing*, 30(2), 185–196. <https://doi.org/10.1016/j.ijresmar.2012.09.004>
- Watson, K., & Jensen, M. M. (2020). Automatic analysis of dialect literature: Advantages and challenges. In P. Honeybone & W. Maguire (Eds.), *Dialect writing and the north of England*. Edinburgh University Press.
- Wright, K. E. (2023). Housing policy and linguistic profiling: An audit study of three American dialects. *Language*, 99(2), e58–e85. <https://doi.org/10.1353/lan.2023.a900094>
- Zhang, S., Xu, J., & Alvero, A. J. (2025). Generative AI Meets Open-Ended Survey Responses: Research Participant Use of AI and Homogenization. *Sociological Methods & Research*. <https://doi.org/10.1177/00491241251327130>

Cite this article: Alvero, A., Sedlacek, Q., León, M., & Peña, C. (2025). Digital accents, homogeneity-by-design, and the evolving social science of written language. *Annual Review of Applied Linguistics*, 45, 50–68. <https://doi.org/10.1017/S0267190525000042>