

# Stop! In the Name of Flaws: Disentangling Personal Names and Sociodemographic Attributes in NLP

Vagrant Gautam<sup>1</sup> Arjun Subramonian<sup>2</sup> Anne Lauscher<sup>3</sup> Os Keyes<sup>4</sup>

<sup>1</sup>Saarland University, Germany <sup>2</sup>University of California, Los Angeles, USA

<sup>3</sup>Universität Hamburg, Germany <sup>4</sup>University of Washington, USA

## Abstract

Personal names simultaneously differentiate individuals and categorize them in ways that are important in a given society. While the natural language processing community has thus associated personal names with sociodemographic characteristics in a variety of tasks, researchers have engaged to varying degrees with the established methodological problems in doing so. To guide future work that uses names and sociodemographic characteristics, we provide an overview of relevant research: first, we present an interdisciplinary background on names and naming. We then survey the issues inherent to associating names with sociodemographic attributes, covering problems of validity (e.g., systematic error, construct validity), as well as ethical concerns (e.g., harms, differential impact, cultural insensitivity). Finally, we provide guiding questions along with normative recommendations to avoid validity and ethical pitfalls when dealing with names and sociodemographic characteristics in natural language processing.

## 1 Introduction

A person’s identity is a complex and paradoxical thing - it simultaneously identifies someone’s *uniqueness*, and categorizes them, identifying what they have in common with others (Strauss, 2017). A perfect example of this phenomenon is a person’s *name*. Personal names are proper nouns used to refer to individuals. They play an important distinguishing role in our lives, as they let us uniquely represent people mentally, refer to them directly in speech, and underscore their significance as individuals (Jeshion, 2009). For these reasons, personal names are a linguistic universal, i.e., they appear across languages and cultures, although naming customs vary across the world (Hough, 2016).

But alongside differentiating people, names also categorize them in their society. Names assigned to




Figure 1: Overview of the methodological issues (concerning validity and ethics) of the use of personal names and sociodemographic characteristics in NLP.

people often index aspects of identity that are important in the context of their society, including sex, religion, tribe, stage of life, etc. Personal names are thus rich resources to understand the social organization of communities, and have been studied across anthropology (Alford, 1987; Hough, 2016), sociology (Marx, 1999; Pilcher, 2017), linguistics (Allerton, 1987; Anderson, 2003), and onomastics (Alvarez-Altman et al., 1987; Adams, 2009).

In natural language processing (NLP) as well, personal names have a long history of use—NLP researchers have worked on identifying and disambiguating uses of personal names (Mann and Yarowsky, 2003; Minkov et al., 2005; Färber and Ao, 2022) and have examined name translation (Sennrich et al., 2016; Wang et al., 2022; Sandoval et al., 2023) and name transliteration (Li et al., 2007; Benites et al., 2020; Sälevä and Lignos, 2024). Increasingly, NLP researchers also use personal names along with sociodemographic characteristics for passive analysis of media and scholarly content (Vogel and Jurafsky, 2012; Knowles et al., 2016; Mohammad, 2020; Asr et al., 2021), or to examine model biases and harms (Maudslay et al., 2019; Romanov et al., 2019; Webster et al., 2021). However, these papers engage to varying degrees with concerns that have been raised outside of NLP about the methodological validity and

ethics of associating names with sociodemographic characteristics. We argue that neglect of these issues is a significant barrier to valid and respectful research, as well as more inclusive NLP systems.

Hence, we contribute an overview of the issues with associating names with sociodemographic attributes (focused on gender and race, two popular categories used in NLP research), as shown in Figure 1. We begin with background on names in other fields and in NLP (§2), and lay out the problems with validity (§3) and ethical concerns (§4) raised when associating personal names with sociodemographic characteristics. Finally, we present guiding questions along with normative recommendations (§5) to guide future work in the field with names.

**Bias statement.** We consider group-level, individual, and representational harms, terms which we explain where we use them in Section 4.

**Positionality statement.** All authors have a background in data science and ethics, and one has a background in philosophy. Three of the authors are trans and have names that are likely unintelligible to popular name-based sociodemographic inference methods. Two authors are trans people of colour; as such, many examples in this paper reflect concerns about misgendering and racialization.

## 2 Background

We begin with some background: on names and naming, and on the use of names in NLP.

### 2.1 Names and Naming

Names are generally regarded as social phenomena that serve two central functions that are sometimes in conflict: differentiation and categorization of individuals (Alford, 1987). Differentiation is important psychologically and semantically for us to be able to directly refer to and mentally represent individuals, and names also serve to underscore their referent’s significance as an individual (Jeshion, 2009). Categorization, on the other hand, is important for the social organization of communities, and naming conventions tend to reflect factors that are important to a community at a given point in time, e.g., gender, religion, descent, transition to adulthood, and so on (Hough, 2016). For instance, the practice of naming someone after their father or grandfather—patronymic naming—was once common across Europe, and was popular in Sweden until the nineteenth century (e.g., *Samuels-*

*son*) and continues into Iceland today (e.g., *Gunnarsdóttir*) (Hough, 2016). This example shows how names and naming can *only* be understood in a specific (geographic, cultural, temporal) context, and even then includes a lot of variation. As folk assumptions about names tend to overlook the wide variation in names and naming (McKenzie, 2010), we present an overview of naming as it relates to sociodemographic characteristics below.

**Variation in societal conventions.** The markers considered important to index in a name vary widely across cultures. For example, almost all European naming systems and indeed most societies across the world tend to assign sex-typed names (Hough, 2016), while South Indian naming conventions often index caste (Meganathan, 2009). However, convention does not mean that every single individual is assigned a name that neatly follows that convention, as shown by the long history of gender-ambiguous names in the U.S. (Barry and Harper, 1982). Additionally, gendered associations for specific names change over time (Barry and Harper, 1993), as do naming conventions in societies—for example, it is becoming increasingly popular to assign non-gendered names in the U.S. and in Israel (Hough, 2016; Obasi et al., 2019). Apart from conventions, names themselves are not static and unchanging from birth, with many names changing due to partnerships, adoption, transition to a different life stage or gender, and so on (Hough, 2016; Obasi et al., 2019; AIATSI, 2022).

**Assimilation and resistance to convention.** Trends in big-picture naming conventions are complicated by factions of society who want to resist imposed classification. Increasingly heterogeneous societies are a natural setting for such tensions; cross-cultural associations with sociodemographic characteristics can differ and sometimes clash, complicating naming, e.g., names like *Nicola* and *Andrea* tend to be assigned to boys in Italy but to girls in Germany. As Germany is a society with highly regulated naming practices, inclusion of these names necessitated a court judgment (Hough, 2016). Immigrant families thus have to juggle the delicate balance of asserting their identity but avoiding name-based stigma and discrimination in the new culture. Their naming practices have therefore been studied as an indicator of attitudes towards assimilation or its rejection, showing how names are not a transparent indicator of race (Sue and Telles, 2007; Becker, 2009). Even among adults, imperial-

ism and colonialism are forces that affect naming. Indigenous individuals have been forced to adopt Western names in settler colonial and postcolonial societies, e.g., the U.S., Canada, Australia (AIAT-SIS, 2022). Similarly, Chinese individuals around the world adopt Western names in conversational (Li, 1997) and professional settings (Chan, 2016). Among trans and gender-nonconforming adults, many choose a new name to reflect and express their gender, walking the tightrope between normativity and self-assertion (Konnelly, 2021); Obasi et al. (2019) find that 50% of gender-nonconforming respondents who change their name pick a gender-neutral name. Beyond transgender people, new names and pseudonyms are also often self-selected to assert agency in identity creation, e.g., bell hooks, Sojourner Truth, and Malcolm X (Baker and Green, 2021).

**Quantitative aspects of naming.** As naming involves a trade-off between differentiation and categorization, names often recur, a quantitative assumption that a lot of sociological, anthropological and NLP classification relies on (Alford, 1987). However, the distributions of names and people can be very different. Weitman (1981) finds that in 100 years of first names from Israel’s Population Registry, the most frequent names (101+ occurrences) account for the majority of the population of a society (91%), but this corresponds to just a tiny minority of all assigned *names* (2.93%). These numbers could vary widely depending on the society, as, for example, the Chuukese people of Micronesia have a tradition of giving entirely unique names to children (Alford, 1987). Hence, it is important to distinguish when names are the object of study and when people are, to contextualize any results that involve the analysis of names.

## 2.2 Names and Sociodemographic Characteristics in NLP

Here, we present a non-comprehensive list of papers to illustrate some common uses of names and sociodemographic characteristics in NLP.

**NLP tasks and problems.** Numerous NLP works have developed algorithms to infer sociodemographic attributes from names (Chang et al., 2010; Liu and Ruths, 2013; Knowles et al., 2016), e.g., for passive analysis of social media content. Another line of NLP papers have relied on names to quantify gender disparities in academic publishing (Vogel and Jurafsky, 2012; Mohammad, 2020)

or media representation (Asr et al., 2021). Some NLP works have identified preserving dominant gender associations as an important criterion for transliteration and translation (Li et al., 2007; Wang et al., 2022). Names are also used to investigate social biases in NLP systems and language models (Kotek et al., 2023; An et al., 2023; Ibaraki et al., 2024). For example, De-Arteaga et al. (2019) study how first names, which they consider “explicit gender indicators,” affect the gender bias of occupation prediction from biographies. Similarly, Jeoung et al. (2023) assess the causal impact of first names, which they posit “may serve as proxies for (intersectional) socio-demographic representations,” on the commonsense reasoning performance of language models. Smith and Williams (2021) measure racial biases as well, evaluating generative dialogue models by having “one conversational partner [...] state a name commonly associated with a certain gender and/or race/ethnicity.” In this line of research, it is commonplace to use skewed reference populations such as U.S. census data (U.S. Census, 2020) and Social Security Administration baby names (U.S. Social Security Administration, 2023) for gender associations (Lockhart et al., 2023).

**Engagement with pitfalls.** In these works, researchers engage to varying degrees with the established methodological and ethical problems of associating names with sociodemographic characteristics. Some NLP papers make unfounded assumptions about names, e.g., Vogel and Jurafsky (2012) posit that certain names are “unambiguous” with respect to gender across languages, and Wang et al. (2022) claim that there exist “names with obvious gender.” Other papers are more critically reflective, acknowledging the limitations of their work: Knowles et al. (2016) state that their classifier to predict gender from names is biased towards the U.S. and assumes gender is binary, but leaves these issues “to be addressed in future work.” Mohammad (2020) acknowledges that inferring gender from names can yield misgendering because “names do not capture gender fluidity or contextual gender,” but suggest a trade-off with “the benefits of NLP techniques and social category detection.” Encouragingly, some recent papers opt for more inclusive study designs after engaging deeply with the pitfalls of using names and sociodemographic characteristics (Sandoval et al., 2023; Saunders and Olsen, 2023; Lassen et al., 2023).

### 3 Validity Issues

In this section, we present issues of validity when associating names with sociodemographic categories, or using names to infer them. Issues of validity mean that results with these operationalizations may neither be indicative of what we actually want to measure, nor of reality.

#### **Error is not quantifiable without asking humans.**

The accuracy of using names to infer sociodemographic characteristics of individuals cannot be quantified without ground truth data, which for people’s identities, can *only* be obtained by asking them. Multiple studies thus empirically analyze the error rates of name-based gender and race inference systems as compared to gold data in different contexts (Karimi et al., 2016; Kozłowski et al., 2022; Van Buskirk et al., 2023; Lockhart et al., 2023).<sup>1</sup> For example, Lockhart et al. (2023) evaluate gender and race inference systems using self-reported data from nearly 20,000 individuals. Importantly, their self-reported data does not directly transfer to other contexts, as their respondents are authors of English language social science journal articles who are mostly located in the U.S. Using this data as reference data for a system with users located primarily in India, or for U.S. authors in a different century, makes little sense. In new environments, it is simply not possible to reasonably estimate the bounds of error of a name-based analysis, and results without a corresponding analysis of self-reported data should not be taken seriously.

#### **Popular design choices lead to systematic error and selection bias.**

Names that are uninformative of a sociodemographic characteristic present an issue for tools that aim to label everyone. In the context of gender, names like *Alex* have no unique gendered association in the U.S. and Canada; with race, names assigned by Black and white parents overlap in the U.S. (Lockhart et al., 2023), and religious names are used around the world (Curtis, 2005; Olúwáfémi, 2014)); at the intersection of gender and race, many Chinese names are not gender-associated when Romanized, and infrequent names are also not informative. Two common design

choices for handling uninformative names are to assign the majority class label anyway, or, alternatively, to just exclude them. Assigning the majority class (i.e., classifying *all* people named *Miaoran* as female if a gender prediction tool predicts the name to be “60% female”) results in systematic error (Kirkup and Frenkel, 2006). On the other hand, excluding uninformative names from the analysis completely alters the makeup of the data and therefore the results (Mihaljević et al., 2019), resulting in selection bias. Both choices affect internal validity, i.e., gaps in the translation from measurements to overall conclusions (Liao et al., 2021), leading to less robust and trustworthy results.

**Poor construct validity.** Construct validity asks how well an abstract concept can be measured through some indicator (Messick, 1995); in our case, the question is: how valid is it to assign sociodemographic categories via names?<sup>2</sup> The answer to this depends on what aspects of the sociodemographic category we are interested in: identity, socialization, expression, perception—all of which could differ and are frequently conflated (Keyes et al., 2021). As discussed previously, many names are simply not informative of certain sociodemographic *identities* in given contexts and with homogeneous populations; Lockhart et al. (2023) find that overall error rates of name-based gender and race imputation tools range from 4.6% to 86% overall, and up to 100% for particular subgroups, depending on the tool. However, when it comes to the *perception* of names as indexing a sociodemographic category, some names may have stronger construct validity, an assumption used by Sandoval et al. (2023) in their examination of names assigned at birth that are strongly associated with the baby’s sex and the parents’ race/ethnicity. On the other hand, Mohammad (2020) uses names to operationalize both identity (to investigate trends in authorship) and perception (to investigate trends in citation) in a bibliometric analysis of the ACL Anthology, even though these need not match, and many underrepresented names are uninformative of identity as well as perception (Van Buskirk et al., 2023). As names do not neatly line up with sociode-

<sup>1</sup>All these studies look at imputing an individual’s gender, but the gold labels they compare to are, confusingly, not always self-reported gender! Some use gender assigned by annotators as the ground truth, which would be fine if comparing to *perceptions* of an individual based on their name, but these studies do not, raising further questions about their methodological validity.

<sup>2</sup>While we focus on the construct validity of names in this section, we note that poor construct validity also applies to the sociodemographic categories themselves (Benthall and Haynes, 2019; Hanna et al., 2020) and to abstract concepts such as “bias” and “fairness,” which show up frequently in the study of names and sociodemographic categories in NLP (Blodgett et al., 2020; Jacobs and Wallach, 2021).

mographic identities, perceptions, or experiences in a context-independent way, it is critical to investigate construct validity of names in any setting where they are used.

**Systems of classification create results.** Although classification is inherently human, classification systems are produced by culture and politics and end up *creating* a view of the world (Bowker and Star, 2000). In computing, researchers have power and our positionality shapes how we view and operationalize categories of classification such as race and gender (Scheuerman et al., 2020b; Scheuerman and Brubaker, 2024). However, many such categories are unstable and contested (Keyes et al., 2021; Mickel, 2024). For instance, it has been shown that different ways of operationalizing race can result in entirely different conclusions (Steidl and Werum, 2019; Benthall and Haynes, 2019; Hanna et al., 2020). Individuals and groups thus cannot be treated as monoliths that can be characterized one-dimensionally via names.

## 4 Ethical Issues

The issues we have examined so far impact the scientific validity of claims made using personal names and sociodemographic categories. Many of these problems arise from assumptions that can also be criticized on ethical grounds, as we show.

**Errors cause harms.** Harms can be broadly described as a setback in the interests or progress of people due to, e.g., the outcomes of an automatic process (Feinberg, 1984). Group-level harms are experienced collectively by people in a sociodemographic group, while individual harms (which might result from group membership) are experienced at the person-person or person-technology level. Inferring gender from names frequently misgenders trans people and erases non-binary people (Keyes, 2018). This perpetrates group-level erasure, as well as individual harms including damaging autonomy and dignity (Mcnamarah, 2020), inflicting psychological harms (Dev et al., 2021), and a failure to show recognition respect to people (Darwall, 1977). Certain types of name-based classification (e.g., of persecuted ethnic or religious groups) can threaten individual safety, and when NLP infrastructure is used for surveillance and targeting, this also threatens the safety of entire

groups of people (Wadhawan, 2022).<sup>3</sup> NLP systems reinforce group-level structural discrimination in other ways as well; name-based studies of racial disparities in academia have been shown to systematically discount the intellectual contributions of Black researchers (Kozlowski et al., 2022).

**Errors and harms are not distributed equally.** In their work on name-based gender classification, Van Buskirk et al. (2023) note that for names with no available data, assigning the majority class (in their case, male) maximizes accuracy, but results in 0% error for the male class and 100% error for any other classes. For non-binary people, who are generally excluded from gender classification by design, the error rate is also almost always 100%. As for name-based race/ethnicity classifiers, Lockhart et al. (2023) show that people who self-identify as Filipino, Black, or Middle Eastern and North African, are misrecognized 55-75% of the time, as compared to those who identify as white, Chinese, or Korean, who are mislabelled less than 10% of the time. As described above, misrecognition errors cause harms, which are then disproportionately experienced by these individuals. We echo the conclusions of Mihaljević et al. (2019) and Lockhart et al. (2023), i.e., that inclusive analyses are only possible when names are no longer used as a proxy to infer individuals' gender or race/ethnicity.

**Representational harms.** The erasure of identities and the flattening of variation in naming customs leads to representational harms, which include the reinforcement of essentialist categories and power structures (Chien and Danks, 2024). These harms primarily affect sociodemographic groups, e.g., non-binary people, who are often incorrectly and unjustly treated as a novel social phenomenon. Groups of people with a certain name are often subject to a different type of representational harm, i.e., stereotyping. For instance, the name *Kevin* is associated with lower socioeconomic class in Germany (Kaiser, 2010). This stereotype, if encoded in an NLP system, could lead to quality-of-service differentials, as class is a sociodemographic characteristic that correlates with lower NLP performance in other contexts (Curry et al., 2024).

**Cultural insensitivity.** Conceptualizations of names and sociodemographic characteristics in

---

<sup>3</sup>Regulation efforts such as the AI Act (Commission, 2021) in the EU try to mitigate this, but this does not apply to authoritarian regimes' use of such technology (Briglia, 2021).

NLP are often Western-centric, with folk assumptions about what names look like and the application of U.S. racial categories and naming preferences to areas outside the U.S., where they are unintelligible (Field et al., 2021). Non-Western naming practices are only sometimes described in papers where there is a specific language of study that is not English, e.g., name tagging in Arabic (Shaalan and Raza, 2007) and Uyghur (Abudukelimu et al., 2018). Even within English, there is little recognition of, e.g., English common nouns used as names in China (*Billboard, Shooting, Pray*, etc.; Chan, 2016), names containing spelling variations (AIAT-SIS, 2022), and names that overlap in different cultures but have different associations, e.g., *Jan* in the U.S. compared to *Jan* in Germany. Beyond names, even gender, race, and other sociodemographic categories of relevance are different across cultures. Many cultures have definitions of gender that go beyond the binary. Enforcing binary gender can thus be seen as an example of what Lugones (2016) calls the “coloniality” of gender, which also results in epistemic violence, i.e., inhibiting people from producing knowledge, or silencing and discrediting their knowledge (Chilisa, 2019).

**No shifts in power.** Names are a site for enforcing institutional power, as seen in “real name” policies (Haimson and Hoffmann, 2016), the (non-consensual) permanence of names in data infrastructure including Google Scholar (Speer, 2021), governmental name regulation (Te Tari Taiwhenua, 2021), and the “collective delusion” of legal names, at least in the U.S. (Baker and Green, 2021). Names are also regulated socially through norms and expectations, many of which end up baked into our NLP systems. We exercise power as NLP researchers and practitioners via our assumptions, which may reify sociodemographic categories, codify (or dismantle) associations between names and these categories, and create infrastructure that harms people at scale through surveillance or mislabelling. Knowles et al. (2016) open-sourced their name-based gender inference tool, and Vogel and Jurafsky (2012) published (binary) gender labels with names of authors of NLP papers, which continue to be used in research (Mohammad, 2020; Van Buskirk et al., 2023). This data reflects folk assumptions about gender, i.e., that it is binary, immutable and in perfect correspondence with names (Keyes, 2018; Cao and Daumé III, 2021). These datasets also deadname and misgen-

der scientists from the NLP community, some of whom have spoken about its harms (Mielke, 2024). Transgender people can only be counted in such a system if they conform to normative expectations (Johnson, 2016; Konnelly, 2021), and if not, the burden is disproportionately on them to seek redress. Even Asr et al. (2021)—a system relying on name-based gender inference that considers gender beyond the binary and does not publicly misgender individuals—does not shift power, as workarounds are a patch rather than built-in to the method; gender inference still relies on APIs that use binary gender, and mistakes (typically, famous non-binary people) are manually corrected. As all these examples show how power remains centralized, we echo previous calls to reimagine and reconfigure power relations in service of user autonomy (Keyes et al., 2019; Blodgett et al., 2020; Hanna and Park, 2020).

## 5 Guiding Questions and Recommendations

In the previous sections, we have reviewed the myriad of issues surrounding the accuracy, validity and ethical use of names along with sociodemographic characteristics, and noted that all these issues arise from the same assumptions and inform each other. In addition, we have shown that these problems apply overwhelmingly to those who are not cisgender, white, normatively named in a Western context, and well-represented in publicly available data. Thus, work that uses names to operationalize people’s sociodemographic categories most misrepresents and further marginalizes those who are already at the margins. We take the normative position that this is not acceptable collateral damage, even (and especially!) in the name of ostensible fairness. Thus, we come up with guiding questions and recommendations for NLP practitioners who are considering the use of names as they relate to sociodemographic categories. These are summarized in Table 1.

**What are you aiming to study—names? Or people, via their names?** It is acceptable to investigate what concepts NLP models associate with names, e.g., *Madeleine* with *kindness*. It is even acceptable to demonstrate that NLP models associate *Marius* with the pronoun *he* or with being male, and that these associations mirror common human associations (Caliskan et al., 2017; Crabtree et al., 2023). It is marginally acceptable to associate names with sociodemographic characteristics using imaginary people, e.g., drawing insights about gen-

Theme	Guiding questions
Names vs. people	What are you aiming to study—names? Or people, via their names? What aspects of names are you interested in? What aspects of people are you interested in?
Context	What is your context? Is processing names with NLP systems necessary to answer your questions?
Harms and power	What kinds of harms apply? How can you mitigate them? Are you describing or prescribing? How does your work reify/redistribute power?
Refusal	Is it still worth it?

Table 1: Our list of guiding questions for the use of names and sociodemographic categories in NLP, grouped by theme. See paragraphs in Section 5 for detailed recommendations.

der bias more broadly based on how NLP models handle synthetic names of people assumed to be exclusively female; while doing so does not compromise people’s autonomy and dignity, it does further entrench hegemonic folk theories of names and people’s identities, which has cultural harms. Finally, it is unacceptable to present results about real people based solely on the assumption that their names provide a reliable signal about their identities, e.g., NLP papers authored by people named *Madeleine* and *Marius* cannot on their own provide trustworthy insights into gender and racial representation in the field, unless those specific individuals are asked about their gender.

#### What aspects of names are you interested in?

Names are rich objects of study with variation in form, length, training data frequency, tokenization, associations, the strengths of these associations, and more.<sup>4</sup> Once you have decided what aspects to study, they must be operationalized and measured carefully, with attention to the context of the study or eventual system deployment. This includes the scope of what counts as a “name.” For instance, considering the use of English common nouns as names (e.g., *Cloud*) is particularly important when working with data from or systems deployed in China, where this naming practice is common (Chan, 2016). Ensure that pre-processing choices are contextualized and do not distort results, that names are understood within context, and that error can be quantified robustly in the given context. Thus, when measuring training data frequency

<sup>4</sup>Some of these aspects have already been explored in prior work in NLP (Shwartz et al., 2020; Wolfe and Caliskan, 2021; Sandoval et al., 2023).

of names, counting *Cloud* tokens as names must consider when it is used as a name and when it is used simply as a noun. Error could be quantified through manual analysis on a subset of the data.

#### What aspects of people are you interested in?

People’s identity and perceptions of them can differ, and these shape their experiences in various ways. Therefore, it is first necessary to decide which aspects are relevant for a study. *Attempting to infer someone’s identity using names is simply unacceptable due to the range of methodological and ethical concerns we list in this paper.* We echo onomastic advice from nearly 40 years ago (Weitman, 1981), i.e., that “inferences from names must be to the givers of these names, not to their bearers. What is more, inferences must always be to sociological formations (such as social classes, ethnic groups, historical generations, and the like), not to individual name-givers.” In addition to studying formations of name-givers, it can also be acceptable to study perceptions of identity based on names. For instance, numerous sociology papers have investigated racial and ethnic perceptions, as well as occupational stereotypes, based on names (King et al., 2006; Gaddis, 2017a,b). Again, we emphasize that perceptions based on names are also highly contextual and non-universal.

**What is your context?** It is essential to understand the geographical, temporal, and cultural context of data with names, and document this information for datasets, e.g., with datasheets (Geburu et al., 2021). What is the geographic, temporal, cultural and political context of the name data, name-bearers, models and sociodemographic categories you use? Who are the people who will be impacted

by your work, and what is their context? What do you know about the naming practices in these contexts and the heterogeneity in these practices? Are you quantifying error with self-reported data? We posit that it is unacceptable to use names without deeply engaging with context in these senses, and stress that ascribing contemporary Western identity categories to historical peoples without acknowledging the difference in contexts is reductive.

**Is processing names with NLP systems necessary to answer your questions?** For questions about human identity and perception based on names, NLP may not be the only or best method available. We warn against technical solutionism (Green, 2021); researchers should reflect on whether their questions could be approached with interviews, case studies, fourth-world paradigms, and so on (Cameron, 2004). Qualitative methods can provide deeper, richer evidence while respecting people’s autonomy, dignity and context. If your questions are instead about NLP systems, then processing names with them is certainly necessary, but we note that methodological pluralism and interdisciplinarity can enrich our practice as NLP researchers and practitioners regardless (Wahle et al., 2023).

**What kinds of harms apply? How can you mitigate them?** Our paper provides a starting point for harms that are relevant to the use of names and sociodemographic characteristics in NLP, and we encourage transparency about methodological and ethical problems (Bietti, 2019; Hao, 2019). It is unacceptable to sideline these problems in the name of “social good” (Green, 2019; Greene et al., 2019; Bennett and Keyes, 2020), and rather than treating entire segments of the world as limitations of or future work for your research, we encourage changing the methods themselves, as Lauscher et al. (2022) do with neopronouns. We recommend firmly grounding work in the ethical principles of autonomy, justice, and beneficence for people (Floridi and Cowls, 2019), which we note are sadly under-represented in machine learning research (Birhane et al., 2022).

**Are you describing or prescribing?** Descriptions of social phenomena are often conflated with normative behaviour (i.e., assumptions and assertions that create and reinforce norms) in NLP (Vida et al., 2023). This is the subtle but significant difference between showing that sociodemographic name associations in language models mirror the

judgements of some group of humans, versus stating that model associations *should* mirror the judgements of some group of humans. The latter “cannot avoid creating and reinforcing norms” (Talat et al., 2022). Therefore, researchers should clearly distinguish descriptive and normative behaviours in the design, execution, and presentation of their experiments (Vida et al., 2023). System designers do have to make decisions about how systems *should* behave, i.e., they need to choose to perpetuate harmful structures in service of usability or to impose their own values on users and stakeholders when they take an advocacy position. This is an ethical dilemma in design that participatory methods and feminist epistemologies are uniquely positioned to help with (Bardzell, 2010).

**How does your work reify or redistribute power?** Central to NLP and computer science at large are scale thinking (Hanna and Park, 2020), quantitative methodologies (Birhane et al., 2022), and the illusion of objectivity (Waseem et al., 2021). All these values serve to reify existing hierarchies and power structures. We must first recognize our own power as NLP researchers and practitioners, and how our work can reinforce infrastructure for (mis)classifying real people and enable surveillance and harms at scale. We recommend a counter-power stance (Keyes et al., 2019), situated knowledges (Haraway, 1988), and methods informed by a politic, e.g., intersectionality, a critical framework that centers justice, power, and reflexivity, and mandates praxis with teeth (Collins, 2019; Erete et al., 2018; Ovalle et al., 2023). Particularly for those of us who are interested in using NLP for social good, we should constantly be asking: “social good for whom?” The differential impact on people matters, and as researchers and practitioners, we have a responsibility to attend to it and resist the othering perpetuated by classification systems.

**Is it still worth it?** After considering all these guiding questions, we remind the reader that refusal is possible (Honeywell, 2016; Tatman, 2020; Lockhart et al., 2023; Mihaljević et al., 2019), and indeed an important part of the history of science (Williams, 1924; United States National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research, 1978; Weindling, 2001).



## 6 Related Work

Several papers study and critically interrogate the inference and use of sociodemographic information in computing (Larson, 2017; Keyes, 2018; Benthall and Haynes, 2019; Hanna et al., 2020; Keyes et al., 2021; Field et al., 2021; Devinney et al., 2022), many of which touch upon names but do not address them in detail. The work that deals with names in particular are all outside of NLP: Karimi et al. (2016); Keyes (2017); Tzioumis (2018); Mihaljević et al. (2019); Scheuerman et al. (2019); Lockhart et al. (2023); Van Buskirk et al. (2023). These papers have different scopes and take a variety of positions with regards to the ethics of name-based inference, some of which we find insufficiently radical. Finally, our recommendations echo those from prior work (particularly in the fields of human-computer interaction and science and technology studies), but are contextualized for names in NLP. Among others, we take inspiration from Keyes et al. (2019); Hanna and Park (2020); Blodgett et al. (2020); Scheuerman et al. (2020a); and Green (2021).

## 7 Conclusion

We present the field with an overview of names and naming as discussed in other disciplines. We lay out background on naming practices around the world and describe how these practices create issues of validity (e.g., selection bias and construct validity) and ethical concerns (e.g., harms, cultural insensitivity), that affect NLP uses of names and sociodemographic characteristics. Finally, we present a list of guiding questions and normative suggestions towards addressing these concerns in future work involving names in NLP.

## Acknowledgments

We thank Lucy Li for recommending literature about personal names, sociodemographic characteristics, and social perceptions. We are also grateful to our reviewers, and to members of the Critical Media Lab Basel, Switzerland, and the Interdisciplinary Institute for Societal Computing, Germany, for their feedback on the ideas in this paper.

## Limitations

Our background on names and naming is limited, and meant only as a brief introduction to onomastics and related fields that use names and

sociodemographic characteristics; space prevents us from being more comprehensive and we refer the interested reader to our references for deeper discussion of onomastic variation. Additionally, we know that problematic and decontextualized assumptions about names are rife within NLP based on our background as authors within or adjacent to the field, as well as writing in other fields about methods that are also popular in NLP. However, as we do not undertake a comprehensive, critical survey of NLP papers that use names and sociodemographic characteristics, we cannot empirically quantify the extent to which the problems we outline plague NLP research, and we leave a more systematic study of this to future work.

## References

- Halidanmu Abudukelimu, Abudoukelimu Abulizi, Boliang Zhang, Xiaoman Pan, Di Lu, Heng Ji, and Yang Liu. 2018. Error analysis of Uyghur name tagging: Language-specific techniques and remaining challenges. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Michael D. Adams. 2009. Power, politeness, and the pragmatics of nicknames. *Names*, 57:81–91.
- AIATSIS. 2022. Indigenous names. <https://aiatsis.gov.au/family-history/you-start/indigenous-names>.
- Richard Alford. 1987. *Naming and identity: A cross-cultural study of personal naming practices*. Hraf Press.
- David J. Allerton. 1987. The linguistic and sociolinguistic status of proper names what are they, and who do they belong to? *Journal of Pragmatics*, 11:61–92.
- Grace Alvarez-Altman, Frederick M. Burelback, Luis A. Oyarzun, and Walter P. Bowman. 1987. *Names in literature : essays from Literary onomastics studies*. University Press of America.
- Haozhe An, Zongxia Li, Jieyu Zhao, and Rachel Rudinger. 2023. SODAPOP: Open-ended discovery of social biases in social commonsense reasoning models. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 1573–1596, Dubrovnik, Croatia. Association for Computational Linguistics.
- John Anderson. 2003. On the structure of names. *Folia Linguistica*.
- Fatemeh Torabi Asr, Mohammad Bagheri Mazraeh, Alexandre Lopes, Vagrant Gautam, Junette Fatima Gonzales, Prashanth Rao, and Maite Taboada. 2021.

- The gender gap tracker: Using natural language processing to measure gender bias in media. *PLoS ONE*, 16.
- Austin A Baker and J Remy Green. 2021. There is no such thing as a ‘legal name’: A strange, shared delusion. *Columbia Human Rights Law Review*, 53:129.
- Shaowen Bardzell. 2010. **Feminist hci: taking stock and outlining an agenda for design**. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, page 1301–1310, Atlanta Georgia USA. ACM.
- Herbert Barry and Aylene S. Harper. 1982. **Evolution of unisex names**. *Names*, 30(1):15–22.
- Herbert Barry and Aylene S. Harper. 1993. **Feminization of unisex names from 1960 to 1990**. *Names*, 41(4):228–238.
- Birgit Becker. 2009. **Immigrants’ emotional identification with the host society: The example of turkish parents’ naming practices in germany**. *Ethnicities*, 9(2):200–225.
- Fernando Benites, Gilbert François Duivestijn, Pius von Däniken, and Mark Cieliebak. 2020. **TRANSLIT: A large-scale name transliteration resource**. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 3265–3271, Marseille, France. European Language Resources Association.
- Cynthia L. Bennett and Os Keyes. 2020. **What is the point of fairness? disability, ai and the complexity of justice**. *SIGACCESS Access. Comput.*, (125).
- Sebastian Benthall and Bruce D. Haynes. 2019. **Racial categories in machine learning**. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, FAT\* ’19, page 289–298, New York, NY, USA. Association for Computing Machinery.
- Elettra Bietti. 2019. **From ethics washing to ethics bashing: a view on tech ethics from within moral philosophy**. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*.
- Abeba Birhane, Pratyusha Kalluri, Dallas Card, William Agnew, Ravit Dotan, and Michelle Bao. 2022. **The values encoded in machine learning research**. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, FAccT ’22, page 173–184, New York, NY, USA. Association for Computing Machinery.
- Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. 2020. **Language (technology) is power: A critical survey of “bias” in NLP**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5454–5476, Online. Association for Computational Linguistics.
- Geoffrey C. Bowker and Susan Leigh Star. 2000. *Sorting Things Out: Classification and Its Consequences*. The MIT Press.
- Makenzie D. Briglia. 2021. Big brother xi: How china’s surveillance of the uyghur population violates international law note. *George Washington International Law Review*, 53(1):85–118.
- Aylin Caliskan, Joanna J. Bryson, and Arvind Narayanan. 2017. **Semantics derived automatically from language corpora contain human-like biases**. *Science*, 356(6334):183–186.
- Graham Cameron. 2004. Evidence in an indigenous world. In *Australasian Evaluation Society 2004 International Conference*, Adelaide, South Australia.
- Yang Trista Cao and Hal Daumé III. 2021. **Toward gender-inclusive coreference resolution: An analysis of gender and bias throughout the machine learning lifecycle\***. *Computational Linguistics*, 47(3):615–661.
- Cherie Chan. 2016. **Why Chinese speakers use Western names**.
- Jonathan D. Chang, Itamar Rosenn, Lars Backstrom, and Cameron A. Marlow. 2010. **e pluribus: Ethnicity on social networks**. *Proceedings of the International AAAI Conference on Web and Social Media*.
- Jennifer Chien and David Danks. 2024. **Beyond behaviorist representational harms: A plan for measurement and mitigation**. *arXiv preprint arXiv:2402.01705*.
- Bagele Chilisa. 2019. *Indigenous research methodologies*. Sage publications.
- Patricia Hill Collins. 2019. *Intersectionality as Critical Social Theory*. Duke University Press.
- European Commission. 2021. **Proposal for a regulation of the european parliament and of the council laying down harmonised rules on artificial intelligence (artificial intelligence act) and amending certain union legislative acts com(2021) 206 final**.
- Charles Crabtree, Jae Yeon Kim, S. Michael Gaddis, John B. Holbein, Cameron Guage, and William W. Marx. 2023. **Validated names for experimental studies on race and ethnicity**. *Scientific Data*, 10.
- Amanda Cercas Curry, Giuseppe Attanasio, Zeerak Talat, and Dirk Hovy. 2024. **Classist tools: Social class correlates with performance in nlp**. *Preprint*, arXiv:2403.04445.
- Edward E. Curtis. 2005. **African-american islamization reconsidered: Black history narratives and muslim identity**. *Journal of the American Academy of Religion*, 73:659–684.
- Stephen Darwall. 1977. **Two kinds of respect**. *Ethics*, 88:36 – 49.

- Maria De-Arteaga, Alexey Romanov, Hanna Wallach, Jennifer Chayes, Christian Borgs, Alexandra Chouldechova, Sahin Geyik, Krishnaram Kenthapadi, and Adam Tauman Kalai. 2019. [Bias in bios: A case study of semantic representation bias in a high-stakes setting](#). In *Proceedings of the Conference on Fairness, Accountability, and Transparency, FAT\* '19*, page 120–128, New York, NY, USA. Association for Computing Machinery.
- Sunipa Dev, Masoud Monajatipoor, Anaelia Ovalle, Arjun Subramonian, Jeff Phillips, and Kai-Wei Chang. 2021. [Harms of gender exclusivity and challenges in non-binary representation in language technologies](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1968–1994, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Hannah Devinney, Jenny Björklund, and Henrik Björklund. 2022. [Theories of “gender” in nlp bias research](#). In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency, FAccT '22*, page 2083–2102, New York, NY, USA. Association for Computing Machinery.
- Sheena Erete, Aarti Israni, and Tawanna Dillahunt. 2018. [An intersectional approach to designing in the margins](#). *Interactions*, 25(3):66–69.
- Michael Färber and Lin Ao. 2022. [The microsoft academic knowledge graph enhanced: Author name disambiguation, publication classification, and embeddings](#). *Quantitative Science Studies*, 3:51–98.
- J. Feinberg. 1984. *Harmless Wrongdoing*. Moral Limits of the Criminal Law. Oxford University Press.
- Anjalie Field, Su Lin Blodgett, Zeerak Waseem, and Yulia Tsvetkov. 2021. [A survey of race, racism, and anti-racism in NLP](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1905–1925, Online. Association for Computational Linguistics.
- Luciano Floridi and Josh Cowsls. 2019. [A Unified Framework of Five Principles for AI in Society](#). *Harvard Data Science Review*, 1(1). <https://hdr.mitpress.mit.edu/pub/l0jsh9d1>.
- S. Michael Gaddis. 2017a. [How black are lakisha and jamal? racial perceptions from names used in correspondence audit studies](#). *Randomized Social Experiments eJournal*.
- S. Michael Gaddis. 2017b. [Racial/ethnic perceptions from hispanic names: Selecting names to test for discrimination](#). *Socius*, 3.
- Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. 2021. [Datasheets for datasets](#). *Commun. ACM*, 64(12):86–92.
- Ben Green. 2019. [Good” isn’t good enough](#). In *Proceedings of the AI for Social Good workshop at NeurIPS*, volume 17.
- Ben Green. 2021. [The contestation of tech ethics: A sociotechnical approach to technology ethics in practice](#). *Journal of Social Computing*, 2(3):209–225.
- Daniel Greene, Anna Lauren Hoffmann, and Luke Stark. 2019. [Better, nicer, clearer, fairer: A critical assessment of the movement for ethical artificial intelligence and machine learning](#). In *Proceedings of the 52nd Hawaii International Conference on System Sciences*.
- Oliver L Haimson and Anna Lauren Hoffmann. 2016. [Constructing and enforcing" authentic" identity online: Facebook, real names, and non-normative identities](#). *First Monday*.
- Alex Hanna, Emily Denton, Andrew Smart, and Jamila Smith-Loud. 2020. [Towards a critical race methodology in algorithmic fairness](#). In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, FAT\* '20*, page 501–512, New York, NY, USA. Association for Computing Machinery.
- Alex Hanna and Tina M Park. 2020. [Against scale: Provocations and resistances to scale thinking](#). *arXiv preprint arXiv:2010.08850*.
- Karen Hao. 2019. [In 2020, let’s stop AI ethics-washing and actually do something](#). <https://www.technologyreview.com/2019/12/27/57/ai-ethics-washing-time-to-act/>.
- Donna Haraway. 1988. [Situated knowledges: The science question in feminism and the privilege of partial perspective](#). *Feminist Studies*, 14(3):575–599.
- Leigh Honeywell. 2016. [neveragain.tech](#). <https://neveragain.tech/>.
- Carole Hough. 2016. *The Oxford Handbook of Names and Naming*. Oxford University Press.
- Katsumi Ibaraki, Winston Wu, Lu Wang, and Rada Mihalcea. 2024. [Analyzing occupational distribution representation in Japanese language models](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 959–973, Torino, Italia. ELRA and ICCL.
- Abigail Z. Jacobs and Hanna Wallach. 2021. [Measurement and fairness](#). In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, page 375–385, Virtual Event Canada. ACM.
- Sullam Jeoung, Jana Diesner, and Halil Kilicoglu. 2023. [Examining the causal impact of first names on language models: The case of social commonsense reasoning](#). In *Proceedings of the 3rd Workshop on Trustworthy Natural Language Processing (TrustNLP 2023)*, pages 61–72, Toronto, Canada. Association for Computational Linguistics.

- Robin Jeshion. 2009. [The significance of names](#). *Mind & Language*, 24(4):370–403.
- Austin H Johnson. 2016. Transnormativity: A new concept and its validation through documentary film about transgender men. *Sociological inquiry*, 86(4):465–491.
- Astrid Kaiser. 2010. „kevin ist kein name, sondern eine diagnose!“der vorname in der grundschule–klangwort, modewort oder reizwort. *Die Grundschulzeitschrift*, 24:26–29.
- Fariba Karimi, Claudia Wagner, Florian Lemmerich, Mohsen Jadidi, and Markus Strohmaier. 2016. [Inferring gender from names on the web: A comparative evaluation of gender detection methods](#). In *Proceedings of the 25th International Conference Companion on World Wide Web - WWW '16 Companion*, page 53–54, Montréal, Québec, Canada. ACM Press.
- Os Keyes. 2017. [Stop mapping names to gender](#).
- Os Keyes. 2018. [The misgendering machines: Trans/hci implications of automatic gender recognition](#). *Proc. ACM Hum.-Comput. Interact.*, 2(CSCW).
- Os Keyes, Josephine Hoy, and Margaret Drouhard. 2019. [Human-computer insurrection: Notes on an anarchist hci](#). In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, CHI '19*, page 1–13, New York, NY, USA. Association for Computing Machinery.
- Os Keyes, Chandler May, and Annabelle Carrell. 2021. [You keep using that word: Ways of thinking about gender in computing research](#). *Proc. ACM Hum.-Comput. Interact.*, 5(CSCW1).
- Eden B. King, Saaid A. Mendoza, Juan M. Madera, Mikki R. Hebl, and Jennifer L. Knight. 2006. [What’s in a name? a multiracial investigation of the role of occupational stereotypes in selection decisions](#). *Journal of Applied Social Psychology*, 36:1145–1159.
- L. Kirkup and R. B. Frenkel. 2006. *Systematic errors*, page 83–96. Cambridge University Press.
- Rebecca Knowles, Josh Carroll, and Mark Dredze. 2016. [Demographer: Extremely simple name demographers](#). In *Proceedings of the First Workshop on NLP and Computational Social Science*, pages 108–113, Austin, Texas. Association for Computational Linguistics.
- Lex Konnelly. 2021. [Nuance and normativity in trans linguistic research](#). *Journal of Language and Sexuality*, 10:71–82.
- Hadas Kotek, Rikker Dockum, and David Sun. 2023. [Gender bias and stereotypes in large language models](#). In *Proceedings of The ACM Collective Intelligence Conference, CI '23*, page 12–24, New York, NY, USA. Association for Computing Machinery.
- Diego Kozłowski, Dakota S Murray, Alexis Bell, Will Hulsey, Vincent Larivière, Thema Monroe-White, and Cassidy R Sugimoto. 2022. [Avoiding bias when inferring race using name-based approaches](#). *Plos one*, 17(3):e0264270.
- Brian Larson. 2017. [Gender as a variable in natural-language processing: Ethical considerations](#). In *Proceedings of the First ACL Workshop on Ethics in Natural Language Processing*, pages 1–11, Valencia, Spain. Association for Computational Linguistics.
- Ida Marie S. Lassen, Mina Almasi, Kenneth Enevoldsen, and Ross Deans Kristensen-McLachlan. 2023. [Detecting intersectionality in NER models: A data-driven approach](#). In *Proceedings of the 7th Joint SIGHUM Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature*, pages 116–127, Dubrovnik, Croatia. Association for Computational Linguistics.
- Anne Lauscher, Archie Crowley, and Dirk Hovy. 2022. [Welcome to the modern world of pronouns: Identity-inclusive natural language processing beyond gender](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 1221–1232, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- David C.S. Li. 1997. [Borrowed identity: Signaling involvement with a western name](#). *Journal of Pragmatics*, 28(4):489–513. Language and Discourse Issues in Hong Kong’s Change of Sovereignty.
- Haizhou Li, Khe Chai Sim, Jin-Shea Kuo, and Minghui Dong. 2007. [Semantic transliteration of personal names](#). In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*, pages 120–127, Prague, Czech Republic. Association for Computational Linguistics.
- Thomas Liao, Rohan Taori, Inioluwa Deborah Raji, and Ludwig Schmidt. 2021. [Are we learning yet? a meta review of evaluation failures across machine learning](#). In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*.
- Wendy Liu and Derek Ruths. 2013. [What’s in a name? using first names as features for gender inference in twitter](#). In *AAAI Spring Symposium: Analyzing Microtext*.
- Jeffrey W. Lockhart, Molly M. King, and Christin Munsch. 2023. [Name-based demographic inference and the unequal distribution of misrecognition](#). *Nature Human Behaviour*, 7(7):1084–1095.
- María Lugones. 2016. *The Coloniality of Gender*, pages 13–33. Palgrave Macmillan UK, London.
- Gideon S. Mann and David Yarowsky. 2003. [Unsupervised personal name disambiguation](#). In *Conference on Computational Natural Language Learning*.

- Gary T. Marx. 1999. [What's in a name? some reflections on the sociology of anonymity.](#) *Inf. Soc.*, 15:99–112.
- Rowan Hall Maudslay, Hila Gonen, Ryan Cotterell, and Simone Teufel. 2019. [It's all in the name: Mitigating gender bias with name-based counterfactual data substitution.](#) In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5267–5275, Hong Kong, China. Association for Computational Linguistics.
- Patrick McKenzie. 2010. [Falsehoods Programmers Believe About Names | Kalzumeus Software.](https://www.kalzumeus.com/2010/06/17/falsehoods-programmers-believe-about-names/) <https://www.kalzumeus.com/2010/06/17/falsehoods-programmers-believe-about-names/>.
- Chan Tov Mcnamarah. 2020. [Misgendering.](#) *The SAGE Encyclopedia of Trans Studies*.
- Ramanujam Meganathan. 2009. [The politics of naming.](#) *Contributions to Indian Sociology*, 43:317–324.
- Samuel Messick. 1995. [Standards of validity and the validity of standards in performance assessment.](#) *Educational Measurement: Issues and Practice*, 14(4):5–8.
- Jennifer Mickel. 2024. [Racial/ethnic categories in ai and algorithmic fairness: Why they matter and what they represent.](#) *ArXiv*, abs/2404.06717.
- Sabrina Mielke. 2024. Personal communication.
- Helena Mihaljević, Marco Tullney, Lucía Santamaría, and Christian Steinfeldt. 2019. [Reflections on gender analyses of bibliographic corpora.](#) *Frontiers in Big Data*, 2:29.
- Einat Minkov, Richard C. Wang, and William W. Cohen. 2005. [Extracting personal names from email: Applying named entity recognition to informal text.](#) In *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing*, pages 443–450, Vancouver, British Columbia, Canada. Association for Computational Linguistics.
- Saif M. Mohammad. 2020. [Gender gap in natural language processing research: Disparities in authorship and citations.](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7860–7870, Online. Association for Computational Linguistics.
- Sharon N. Obasi, Richard MocarSKI, Natalie Holt, Debra A. Hope, and Nathan Woodruff. 2019. [Renaming me: Assessing the influence of gender identity on name selection.](#) *Names*, 67(4):199–211.
- Ìkòtún Reuben Olúwáfémi. 2014. [New trends in yorùbá personal names among yorùbá christians.](#) *Linguistik Online*, 59(2).
- Anaelia Ovalle, Arjun Subramonian, Vagrant Gautam, Gilbert Gee, and Kai-Wei Chang. 2023. [Factoring the matrix of domination: A critical review and reimagination of intersectionality in ai fairness.](#) In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '23, page 496–511, New York, NY, USA. Association for Computing Machinery.
- Jane Pilcher. 2017. [Names and “doing gender”: How forenames and surnames contribute to gender identities, difference, and inequalities.](#) *Sex Roles*, 77:812–822.
- Alexey Romanov, Maria De-Arteaga, Hanna Wallach, Jennifer Chayes, Christian Borgs, Alexandra Chouldechova, Sahin Geyik, Krishnamurthy Kenthapadi, Anna Rumshisky, and Adam Kalai. 2019. [What's in a name? Reducing bias in bios without access to protected attributes.](#) In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4187–4195, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jonne Sälevä and Constantine Lignos. 2024. [ParaNames 1.0: Creating an entity name corpus for 400+ languages using Wikidata.](#) In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 12599–12610, Torino, Italia. ELRA and ICCL.
- Sandra Sandoval, Jieyu Zhao, Marine Carpuat, and Hal Daumé III. 2023. [A rose by any other name would not smell as sweet: Social bias in names mistranslation.](#) In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 3933–3945, Singapore. Association for Computational Linguistics.
- Danielle Saunders and Katrina Olsen. 2023. [Gender, names and other mysteries: Towards the ambiguous for gender-inclusive translation.](#) In *Proceedings of the First Workshop on Gender-Inclusive Translation Technologies*, pages 85–93, Tampere, Finland. European Association for Machine Translation.
- Morgan Klaus Scheuerman and Jed R. Brubaker. 2024. [Products of positionality: How tech workers shape identity concepts in computer vision.](#) In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, page 1–18, Honolulu HI USA. ACM.
- Morgan Klaus Scheuerman, Jacob M. Paul, and Jed R. Brubaker. 2019. [How computers see gender: An evaluation of gender classification in commercial facial analysis services.](#) *Proc. ACM Hum.-Comput. Interact.*, 3(CSCW).
- Morgan Klaus Scheuerman, Katta Spiel, Oliver L Haimson, Foad Hamidi, and Stacy M Branham. 2020a. [Hci guidelines for gender equity and inclusivity.](#)

- Morgan Klaus Scheuerman, Kandrea Wade, Caitlin Lustig, and Jed R. Brubaker. 2020b. [How we've taught algorithms to see identity: Constructing race and gender in image databases for facial analysis](#). *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW1):1–35.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural machine translation of rare words with subword units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.
- Khaled Shaalan and Hafsa Raza. 2007. [Person name entity recognition for Arabic](#). In *Proceedings of the 2007 Workshop on Computational Approaches to Semitic Languages: Common Issues and Resources*, pages 17–24, Prague, Czech Republic. Association for Computational Linguistics.
- Vered Shwartz, Rachel Rudinger, and Oyvind Tafjord. 2020. [“you are grounded!”: Latent name artifacts in pre-trained language models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6850–6861, Online. Association for Computational Linguistics.
- Eric Michael Smith and Adina Williams. 2021. [Hi, my name is martha: Using names to measure and mitigate bias in generative dialogue models](#). *ArXiv*, abs/2109.03300.
- Robyn Speer. 2021. [Google scholar has failed us](#).
- Christina R. Steidl and Regina Werum. 2019. [If all you have is a hammer, everything looks like a nail: Operationalization matters](#). *Sociology Compass*, 13(8):e12727.
- Anselm L. Strauss. 2017. *Mirrors and masks: The search for identity*. Routledge.
- Christina A. Sue and Edward E. Telles. 2007. [Assimilation and gender in naming](#). *American Journal of Sociology*, 112(5):1383–1415.
- Zeerak Talat, Hagen Blix, Josef Valvoda, Maya Indira Ganesh, Ryan Cotterell, and Adina Williams. 2022. [On the machine learning of ethical judgments from natural language](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 769–779, Seattle, United States. Association for Computational Linguistics.
- Rachael Tatman. 2020. What i won't build. <https://www.rctatman.com/talks/what-i-wont-build>.
- Department of Internal Affairs | Te Tari Taiwhenua. 2021. Press Releases - dia.govt.nz — dia.govt.nz. <https://www.dia.govt.nz/press.nsf/d77da9b523f12931cc256ac5000d19b6/d1288ac08d7758c2cc25838200107411!OpenDocument>.
- Konstantinos Tzioumis. 2018. [Demographic aspects of first names](#). *Scientific Data*, 5.
- United States National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research. 1978. *The Belmont report: ethical principles and guidelines for the protection of human subjects of research*, volume 2. United States National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research.
- U.S. Census. 2020. [First name frequency by gender](#).
- U.S. Social Security Administration. 2023. [Top 10 baby names of 2023](#).
- Ian Van Buskirk, Aaron Clauset, and Daniel B. Larremore. 2023. [An open-source cultural consensus approach to name-based gender classification](#). *Proceedings of the International AAAI Conference on Web and Social Media*, 17:866–877.
- Karina Vida, Judith Simon, and Anne Lauscher. 2023. [Values, ethics, morals? on the use of moral concepts in NLP research](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5534–5554, Singapore. Association for Computational Linguistics.
- Adam Vogel and Dan Jurafsky. 2012. [He said, she said: Gender in the ACL Anthology](#). In *Proceedings of the ACL-2012 Special Workshop on Rediscovering 50 Years of Discoveries*, pages 33–41, Jeju Island, Korea. Association for Computational Linguistics.
- Subhah Wadhawan. 2022. Let the machines do the dirty work: Social media, machine learning technology and the iteration of racialized surveillance. *Canadian Journal of Law and Technology*, 20(1):1.
- Jan Philip Wahle, Terry Ruas, Mohamed Abdalla, Bela Gipp, and Saif Mohammad. 2023. [We are who we cite: Bridges of influence between natural language processing and other academic fields](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12896–12913, Singapore. Association for Computational Linguistics.
- Jun Wang, Benjamin Rubinstein, and Trevor Cohn. 2022. [Measuring and mitigating name biases in neural machine translation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2576–2590, Dublin, Ireland. Association for Computational Linguistics.
- Zeerak Waseem, Smarika Lulz, Joachim Bingel, and Isabelle Augenstein. 2021. [Disembodied machine learning: On the illusion of objectivity in nlp](#). *Preprint*, arXiv:2101.11974.
- Kellie Webster, Xuezhi Wang, Ian Tenney, Alex Beutel, Emily Pitler, Ellie Pavlick, Jilin Chen, Ed Chi, and Slav Petrov. 2021. [Measuring and reducing gendered correlations in pre-trained models](#). *Preprint*, arXiv:2010.06032.

- Paul Weindling. 2001. The origins of informed consent: The international scientific commission on medical war crimes, and the nuremberg code. *Bulletin of the History of Medicine*, 75(1):37–71.
- Sasha Weitman. 1981. Some methodological issues in quantitative onomastics. *Names*, 29(3):181–196.
- John F. Williams. 1924. The geneva protocol of 1924 for the pacific settlement of international disputes1. *Journal of the British Institute of International Affairs*, 3(6):288–304.
- Robert Wolfe and Aylin Caliskan. 2021. Low frequency names exhibit bias and overfitting in contextualizing language models. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 518–532, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.