

Language technology practitioners as language managers: arbitrating data bias and predictive bias in ASR

Nina Markl, Stephen Joseph McNulty

University of Edinburgh

nina.markl@ed.ac.uk, stephen.mculty@ed.ac.uk

Abstract

Despite the fact that variation is a fundamental characteristic of natural language, automatic speech recognition systems perform systematically worse on non-standardised and marginalised language varieties. In this paper we use the lens of language policy to analyse how current practices in training and testing ASR systems in industry lead to the data bias giving rise to these systematic error differences. We believe that this is a useful perspective for speech and language technology practitioners to understand the origins and harms of algorithmic bias, and how they can mitigate it. We also propose a re-framing of language resources as (public) infrastructure which should not solely be designed for markets, but for, and with meaningful cooperation of, speech communities.

Keywords: Automatic Speech Recognition, Algorithmic Bias, Language Policy

1. Introduction

All language communities, even monolingual ones, show linguistic variation. The co-existence of multiple different ways of communicating the same meaning is a fundamental characteristic of natural language. Speakers can employ different words to refer to the same object (e.g. “film” and “movie”), pronounce the same word differently (e.g. “data” and “data”), address interlocutors differentially depending on the context (e.g. “you”, “yous”, “y’all” in many varieties of English), and even utilise different sentence structures (e.g. “data is” and “data are”). Despite that fact that this variation is inevitable, people still form judgements about them. Very often, these judgements reflect biases about (groups of) people, not language per se.

These social-linguistic judgements contribute to differential access to, and performance of, language technologies for speakers of the over 7000 language varieties¹ spoken in the world. Most language communities globally do not have access to them at all, and within those that do, performance for speakers of non-standard(ised) and marginal(ised) varieties is worse. For automatic speech recognition (ASR) systems, this “predictive bias”, defined by Shah et al. (2020) as a systematic error disparity between different user groups, arises in part from data bias in the speech datasets used to train and test them. In this paper, we use the lens of “language policy” to understand the origins and consequences of this data bias, and to facilitate its mitigation. We contend that, perhaps unknowingly, organisations – and particularly individuals – involved in the design and creation of these datasets, whether crowdsourced or curated, perform the function of “language policy arbiters” (Johnson, 2013). In

their selection of widely spoken, prestigious (and often commercially-viable) varieties, these individuals effectively marginalise speakers of minority or lesser-used languages or forms of language. This marginalisation may take the form of limiting access to these technologies and exacerbating stigma towards some varieties in their wider application, thereby amplifying systemic discrimination against particular groups and their language(s). Yet, by recognising the need for proactive, diversity-oriented language management – and their role in engendering it – speech and language technologists can work to mitigate such harms, and work towards more equitable and inclusive technologies.

This paper illustrates the role of language policy in exacerbating predictive bias. While we use the example of English language ASR because of the outsize attention English has received in research on speech technology and algorithmic bias, the framework remains applicable in many contexts in which one variety is considered dominant vis-a-vis another.²

2. Predictive bias in ASR

Recent work shows that state-of-the-art commercial English language ASR systems display significant predictive bias for African American English (AAE) and some regional varieties of English. Koenecke et al. (2020) document dramatic racial error disparities for ASR systems sold by Google, Amazon, Microsoft, IBM and Apple, with much larger error rates for Black speakers of AAE than White speakers of Californian English. Overall, recent research suggests that this predictive bias is driven by under-representation of AAE in training data for both acoustic models (Koenecke et al., 2020) and language models (Martin and Tang, 2020) used by commercial ASR systems. Koenecke et

¹“Language variety” refers to languages (e.g. English), “dialects” (e.g. Scottish English) and accents (e.g. Standard Scottish English). The linguistic features characterising a variety are called “variants”.

²This focus, in and of itself is, of course, also one of the central ways in which speech and language technologies are “biased” and deeply unequal (Joshi et al., 2020).

al. (2020) find error disparities based on pronunciation differences, while Martin and Tang (2020) show that Google Cloud Speech-to-Text handles AAE syntactic features such as “habitual be”³ poorly. A slightly older set of studies has shown similar error disparities for different regional varieties of English, including Scottish English and Southern U.S. English in products sold by Google and Bing (Tatman, 2017; Tatman and Kasten, 2017). This apparent data bias is not limited to commercial ASR systems, as Mozilla’s open-source DeepSpeech system trained on their crowdsourced CommonVoice corpus also performs significantly worse for AAE than Mainstream US English (Martin and Tang, 2020). Other work has focused on the use of ASR as an assistive technology and found that most major systems perform poorly for Deaf and hard of hearing (Glasser, 2019), and dysarthric users (De Russis and Corno, 2019; Young and Mihailidis, 2010).

2.1. Harms of predictive bias

While technical research on bias mitigation is important, especially because particular model structures can amplify data bias (Hooker, 2021), it is crucial to consider the socio-historical origins of predictive bias and its consequences. Most obviously, speech recognition is a component of voice user interfaces which can be used as assistive technologies to access mobile devices and computers. As (mobile) computing becomes increasingly ubiquitous, predictive bias could severely disadvantage AAE speakers and other speakers of stigmatised and under-represented language varieties in completing everyday tasks such as making phone calls, searching for information on the web and sending emails, and engaging agents in private and public sector contexts. Recently speech recognition has joined other AI technologies in moving into very high-stakes contexts such as hiring and healthcare (Lee, 2021)⁴. Companies like HireVue⁵, for example, claim to use “voice data” including information about voice quality and lexical choice to pre-screen and rank job applicants (Raghavan et al., 2020). While HireVue has recently passed an independent audit of their algorithmic systems, according to which their training data is balanced by race, gender, region and job title, accent-based bias among first and second language speakers of English has not been studied⁶. HireVue and its competitors also offer customisation of training data for client companies, which makes identification and mitigation of data bias in practice particularly

³A common feature of AAE not found in Mainstream US English e.g. “I be in my office at 7.30” which is equivalent to MUSE “I am usually in my office at 7.30” (Green, 2002).

⁴Amazon even sells an ASR system specifically for medical transcription: <https://aws.amazon.com/transcribe/medical/>

⁵<https://www.hirevue.com/>

⁶<https://www.hirevue.com/blog/hiring/industry-leadership-new-audit-results-and-decision-on-visual-analysis>

difficult (Raghavan et al., 2020). By disadvantaging marginalised speech communities in accessing technology and resources (up to and including employment), predictive bias can further reify and entrench existing linguistic, and by extension, social hierarchies.

3. Language Policy

As Blodgett et al. (2020) show, discussions of “bias” in the language technology literature often lack grounding in the broader socio-historical context of users or the system, failing to spell out what exactly is meant by “bias,” who is harmed by it and how it relates to larger power structures. In this paper, we use the sociolinguistic concept of “language policy” to understand both where data bias comes from and whom it harms.

Language policy relates to the rules, conventions, choices, values, ideas, or discourses which govern the way that we use or think about languages and their speakers (Spolsky, 2004; Johnson, 2013). These policies can be either explicit or overt, as is the case of language legislation or institutional language policy documents, or can be concealed, covert or de facto – often couched in decisions or actions not specifically related to languages, or in implicit judgements about them or those who speak them (Shohamy, 2006). For Spolsky (2004) it is composed of three distinct but interrelated phenomena, which we will explain in turn. *Language practices* refer to conventionalised or patterned language behaviours; *language ideologies* are value-based judgements of specific language varieties and variants, and by extension their speakers and communities; and *language management* refers to attempts to modify language practice and language ideologies.

3.1. Language practices

As we have noted, language use is characterised by variation. In an influential formulation, Weinreich et al. (1968) refer to this variation as “orderly heterogeneity”. That is, variation in language is patterned and rule-governed. Individuals and speech communities use this variation to construct social identities in interaction (Eckert, 2012). These linguistic choices, which are constrained by the social norms transmitted within a community, make up the community’s language practices (Spolsky, 2004).

Understanding patterns of language variation is crucial to identifying the sources of predictive bias in ASR (and other speech and language technologies) and developing mitigation strategies. For instance, some earlier work on predictive bias in ASR noted apparent differences according to speaker gender (Adda-Decker and Lamel, 2005; Benzeghiba et al., 2007; Tatman, 2017). Adda-Decker and Lamel (2005) locate the source of better performance for women’s speech as compared to men’s speech in data bias in training and test datasets. In the English language broadcast news training and test corpora Adda-Decker and Lamel (2005) use, women are more likely to be newscasters and interviewers who adopt a formal speech style,

while men are more likely interviewees whose speech is more often unplanned and conversational and thus characterised by repetitions, phonetic reduction, backchannels and filled pauses. In addition to the conversational roles of women and men in these datasets, they also attribute differences to broader gendered patterns of language use, whereby women tend to avoid stigmatised linguistic features more than men (Labov, 1990). Overall, the more formal speech styles associated here with women are easier to process for the ASR system⁷. This gendered pattern in language use is also reflected in Koenecke et al. (2020), who find that commercial ASR systems are more error-prone for men. An analysis of their test set shows that men are, generally speaking, more likely to use higher rates of non-standard forms (Koenecke et al., 2020). This “gender gap” in ASR performance and speech patterns is furthermore substantially larger for Black speakers, highlighting that race and gender as interacting axes of oppression cannot be considered separately, as has long been noted by Black feminist scholars like Crenshaw (1991) and Hill Collins (2000). In addition to gender and race, other relevant social factors conditioning language variation are socio-economic class, educational background, linguistic background, disability and ethnicity (Herk, 2018). Which of these factors are particularly important depends on the specific context. Generally, varieties spoken by powerful groups within a society or societal context (e.g. higher social class groups, White groups, particular geographical areas) become associated with prestige (due to their association with power). Often these prestigious varieties are also “standard varieties”, codified in prescriptive (rather than descriptive) grammars and taught in the education system (Herk, 2018). Poor ASR performance on non-standard varieties, then, is more likely to affect already marginalised speech communities.

3.2. Language ideology

From a linguistic perspective, no language variety is inherently “better” or “worse” than any other. However, because language (variation) is always situated within larger social contexts, specific ways of speaking can become indices of particular social identities. Language users create beliefs about language to explain and justify these (arbitrary) associations between speaker and form. As Irvine and Gal (2000, 37) put it, these beliefs “locate linguistic phenomena as part of, and evidence for, what [language users] believe to be systematic behavioral, aesthetic, affective and moral contrasts among the social groups indexed”.

Language users (all of us) lean on these ideologies when we make judgements about other people (al-

⁷Garnerin et al. (2021) show that when women’s speech is under-represented in training sets of read speech, performance is significantly better for men. Notably, adding more women’s voices improves performance for women without degrading performance for men.

beit often unconsciously). Like other ideologies, they often seem to reflect “common sense” and resisting them requires conscious effort. They also have real implications for, in particular, marginal(ised) groups. For example, many studies show that second language speakers of English are less likely to be hired (Hosoda and Stone-Romero, 2010; Timming, 2016) and are frequently rated “less credible” (Lev-Ari and Keysar, 2010) than first language speakers. There are further very strong language ideologies around “professional”, “educated” and “articulate” speech (Lippi-Green, 2012; Baratta, 2017). These ideologies are underpinned by broader structural biases within a society such as racism, classism and sexism. It is because of those broader structures that some social groups (e.g. White, upper and middle class, men) have more power relative to others (e.g. Black and non-white, lower and working class, women and non-binary people), and as a result, their speech becomes associated with power and prestige. Notably, language ideologies are not applied in the same way in every context. This makes sense if we recall that the supposed attitudes about particular linguistic features aren’t about language per se, but about the social identities they are associated with. For example, some voice qualities like creaky voice (“vocal fry”) are more stigmatised in young English speaking women than men (Anderson et al., 2014)⁸.

Language ideologies feed into speech and language technologies in many different ways. As we explore in this paper, they influence which kind of language we use to train and test language technologies, and as a result, who is most impacted by predictive bias. But speech and language technologies also reinforce language ideologies. Better performance on some varieties emphasises their status. And even in cases where there’s no predictive bias, they can reinforce existing ideologies. For example, HireVue (and similar models) may not show predictive bias as such, but since they are trained on interviews with successful job applicants, language ideologies around “professionalism” as expected during a job interview are encoded. At first glance it may seem fairer if these harsh prejudices are part of an algorithmic system since they are at least applied equally (e.g. “vocal fry” = “bad,” “long sentences” = “articulate” = “good”). But more privileged people are much more likely to have access to “the right way of speaking” in an interview, for example because they have been taught how to speak during interviews and what kind of language (features) to avoid. According to the HireVue audit report, applicants from “minority” backgrounds are more likely to give very short answers which potentially puts them at a disadvantage as the system does not pose follow up questions. Beliefs and ideologies about language varieties are in-

⁸Anderson et al. (2014) conclude that women should avoid creaky voice. We reject this conclusion, and point instead to Chao and Bursten (2021) for a detailed feminist critique of the response to women’s creaky voice(s).

evitable and omnipresent. They simply represent “what people think should be done” with regards language use within specific societal contexts (Spolsky, 2004, 14). It is when these sets of preconceived judgments begin to affect language-related decisions that we enter the realm of language management.

3.3. Language management

Language management occurs across all spheres and sectors of society, and involves a wide and diverse range of actors (Spolsky, 2004; Blommaert et al., 2009; Hornberger and Johnson, 2007). What is crucial to remember, is that, even when pertaining to the most mundane, mechanical and technical actions or decisions, language policies are never neutral. By its very nature, language management involves taking a stance on language varieties and variation, by deciding which forms of speech are appealing, acceptable or correct, and which are unattractive, inferior or simply “wrong”. Moreover, as Tollefson (1991) and Shohamy (2006) note, language management often serves to create, reify and reproduce unequal power divisions within society: privileging speakers of dominant, prestigious varieties (e.g. native speakers of a standard form of English) and further marginalising people who use stigmatised forms of language (e.g. non-native speakers of [a non-standard] English, or speakers of minority languages). Furthermore, as Wiley (2012) highlights, the absence of an official policy or non-consideration of issues related to equality and diversity in language, often serves only to reinforce the power and hegemony of prestige varieties, and marginalise others: “The lack of recognition of ‘nonstandard’ varieties of language [...] positions their speakers as merely ‘substandard’ articulators of English.” Inaction, therefore, is action. As noted previously, language managers, planners or policy actors can take many forms. According to Johnson and Johnson (2014), however, certain individuals are endowed with a disproportionate amount of power within specific language policy processes. As a result of their position of influence within a given organisational, institutional or social hierarchy, these “language policy arbiters”, through their interpretations and ideological reflexivity (or lack thereof), can influence how language policies are created or implemented (Hornberger and Johnson, 2007). Given the possible impacts of their actions, if social inequalities are truly to be redressed, it is essential that these individuals recognise how much power they wield. The design and creation of speech technologies, we believe, constitutes a form of language management with consequences across societal scales, and its designers and operators perform the role of language policy arbiters for their end users, as well as for society more generally.

4. State-of-the-art: training & testing

An example of this form of language management would be the curation of speech datasets used in the

training and testing of ASR systems. It is through this process that decisions about what kind of language to include or exclude in training and test datasets are made. These decisions then shape for which kinds of language, and therefore for which kinds of speakers, these technologies are useful rather than harmful.

4.1. Training ASR in industry

ASR systems by corporations like Amazon and Google, or large foundations such as Mozilla, are trained on very large datasets. In the case of commercial ASR these datasets consist (at least in part) of voice commands and dictation snippets which are collected from customers during their interactions with voice user interfaces and transcribed by employees⁹. Mozilla’s corpora are made up of voice recordings which are submitted, transcribed and validated by volunteers via an online platform¹⁰. As explored in 2, both types of systems exhibit predictive bias towards less prestigious varieties, in particular African American English. In the following section, we explore how corporate language policies influence the apparent data bias giving rise to these error disparities.

4.1.1. Corporate: Proprietary user data

Corporations like Amazon, Google and Microsoft do not provide detailed model documentation for the ASR systems they sell to third parties (e.g. Amazon Transcribe, Google Cloud Speech-to-Text) or the ones embedded in their own products such as voice user interfaces (e.g. Siri, Alexa, Cortana) and video platforms (e.g. YouTube captions) but their privacy notices and academic publications suggest that large proprietary datasets which include data collected from users are involved. For example, Chiu et al. (2018) (Google) present a system which is trained on “representative voice search data” from their user base. Similarly, Facebook AI trained a multilingual ASR system on “publicly shared user videos” in 51 languages (Pratap et al., 2020). A fundamental problem with training on user data is that even if this data is “representative” of the user base, the user base is not necessarily representative of the population at large. According to a 2021 Pew Research Center survey, 85% of residents of the United States own a smartphone¹¹. However, there are still quite big gaps between different age and social class groups. There are further even larger gaps in home broadband access depending on income in particular. As has been raised in the context of large language models, while digital spaces are in theory “open to everyone”, participation in online communities is not equally accessible or attractive everyone (Bender et al., 2021). Any dataset based on online communication, then, risks mis- or under-representing

⁹With consent of the users, as indicated in the privacy notices of e.g. Apple, Microsoft, Amazon and Google

¹⁰<https://commonvoice.mozilla.org/>

¹¹<https://www.pewresearch.org/internet/fact-sheet/mobile/>

marginalised (speech) communities who may not be able or willing to in participate (Bender et al., 2021). Indeed, the findings by Koenecke et al. (2020), Martin and Tang (2020) and Tatman and Kasten (2017) suggest that, in the context of US English, Black talkers in particular remain under-represented. To avoid predictive bias, data from different groups would have to be balanced rather than merely representative of the (skewed) population distribution (Suresh and Guttag, 2021; Barocas et al., 2019)¹².

Big (speech and language) technology companies do not tend to have publicly available officially declared language policies. However, as alluded to above, just because there is no official document outlining a language policy, it does not mean that there is no policy in place. Some language policy scholars such as Schiffman (1996) and Shohamy (2006) distinguish between *de jure* and *de facto* language policies. Even in the absence of the former, *de facto* policies can still arise, often on the basis of what people in a particular context find to be sensible, convenient or common sense. In this context, beliefs about language (i.e. language ideologies) can be particularly influential (Shohamy, 2006). A key aspect of language management is the selection of a particular language variety to be used in a particular context. In the context of speech and language technologies, this selection process includes the choice of a particular variety to train and test a system on, and consequently, develop for. For example, Benjamin (2019) quotes a former Apple speech technology researcher working on Apple's voice assistant Siri asking their supervisor in 2015 why AAE was not a priority while support for other varieties of English such as Singaporean English was being developed. The response: "Well, Apple products are for the premium market." (Benjamin, 2019, 15). This statement expresses a language ideology held by (at least a part of) the corporation: AAE is not spoken by "the premium market" and AAE speakers do not (or cannot afford to) buy "premium products". Assuming that Apple's main goal is to attract (and keep) the "premium market" as is implicit in the quote above, only developing "premium" linguistic varieties is a good investment. This ideology is the company's *de facto* language policy: AAE is not supported by the company. By applying this economic reasoning to language varieties (and their speakers), Apple also reinforces existing "linguistic markets" (Bourdieu, 1977). It's perhaps not surprising that Koenecke et al. (2020) found the racial gap in predictive errors to be largest, and overall performance on AAE to be worst for Siri (as compared to other systems tested). More broadly, selecting language varieties based on their perceived value on the (linguistic) market means that varieties spoken by marginalised or small communities are less likely to be supported. Dif-

¹²As Hooker (2021), notes, the fact that most "real-world" data have skewed distribution is why it's important to focus on mitigating bias through model choice too.

ferences in language policy between corporations are also reflected in the different sets of languages they select. Google has the largest range of language varieties, including national varieties for languages like Arabic, Urdu, English and Spanish¹³. While smaller national and regional languages spoken in Europe (like Macedonian and Basque) are supported, the same can only be said for languages with larger speaker populations outwith Europe like Uzbek, Zulu, Amharic, and Gujarati, highlighting a general global skew in speech technology availability. Similarly, Apple's Siri is offered in US Spanish and two post-colonial English varieties (India & Singapore) but does not support any languages indigenous to Africa, the Americas, Oceania or the Indian subcontinent. These choices do not just impact current and future customers of these technology corporations: Apple, Google and Microsoft sell their speech recognition services to third parties, and their choices (of data and algorithms) likely impact the way smaller companies act.

4.1.2. Open-source: Crowdsourcing

The most obvious alternative to this purely market-driven model of technology development already in use today are open-source and crowdsourced technologies, such as Mozilla's DeepSpeech ASR system and CommonVoice collection of crowdsourced speech datasets¹⁴. The latter currently covers 76 languages. Volunteers contribute by reading out sentences which are recorded via an interactive interface and validated by other volunteers. All contributors can optionally provide information about their gender, age and accent. CommonVoice does not appear to have a top-down policy for selecting language varieties. Volunteers can request the initiation of a corpus for a new language. The accent labels available for volunteers seem to be selected by community members¹⁵, with Spanish varieties defined in geographic terms while German varieties are defined as national varieties (eliding variation within nation states). Similarly, the English corpus contains "Scottish English" and "England English" alongside a very broad "US English" making comparisons of sampling bias very difficult. Mozilla is currently in the process of replacing this apparent "null policy" with a declared "languages and accent strategy" which has at least in part been crowdsourced in discussion with community members on a public Mozilla discussion forum (and seems to have also been informed by discussion with linguists) (Mozilla Common Voice: Discourse, 2019; Mozilla Common Voice: Discourse, 2020; Mozilla Common Voice: Community Playbook, ; Mozilla Common Voice, 2022).

For smaller or marginalised speech communities and/or those in the Global South in particular, this participatory framework of crowdsourcing both language and

¹³<https://cloud.google.com/speech-to-text/docs/languages>

¹⁴<https://commonvoice.mozilla.org/en>

¹⁵<https://discourse.mozilla.org/t/spanish-accent/35638>

language policy appears a better strategy for speech and language technology development than relying on large for-profit corporations. Speakers can engage in “conscious data contribution” (Vincent et al., 2021), and (within limits) directly shape what kind of language(s) DeepSpeech will support. For some varieties, like Kabyle (a Berber language with 7 million speakers) or Kinyarwanda (a Niger-Congo language with 12 million speakers) this approach also appears successful as they have sizeable validated corpora. However, some varieties (regardless of speaker numbers) have only very small CommonVoice corpora, or corpora which are very unbalanced across varieties (most notably Arabic, which has a large number of distinct dialects spoken in different regions but is currently only represented by Standard Arabic), as well as age and gender groups. The majority of the contributors to the English CommonVoice corpus, for example, did not provide any information about their accent and only 15% identified themselves as female. Notably, in the context of existing research on bias in ASR, CommonVoice does not collect information on race or ethnicity, and “African American English” is not one of the possible “native accents”. This lack of documentation makes evaluation of data bias difficult. Overall, while crowdsourcing can alleviate some of the data bias issues we see in commercial ASR, especially when done with an explicit focus on accent diversity, many representation issues persist. Recall Martin and Tang (2020) show that it performs worse for AAE speakers. As has also been discussed in the context of Wikipedia, who contributes to crowdsourced projects depends on many factors such as availability of free time, technical skills, access to digital technology and the culture of the crowdsourced project (Hargittai and Shaw, 2015; Tripodi, 2021). Crowdsourcing also places the onus to create data on potentially already marginalised speech communities who might furthermore disagree about how (and if) their language should be represented in these systems (e.g. which accents or writing systems) and how they would like any finished system to be used.

4.2. Testing

In speech and language technologies (and machine learning more broadly), benchmark datasets are used to evaluate the performance of new algorithmic systems (Schlangen, 2021). While this focus on benchmarks has recently become the subject of critique (Bowman and Dahl, 2021; Denton et al., 2020; Koch et al., 2021; Raji et al., 2021), they are still central to the way the field defines “progress”. In the following section we explore how language ideologies shape the well-established academic benchmark corpora TIMIT (English) (Garofolo et al., 1993), Switchboard (Godfrey and Holliman, 1993) and CallHome (American English) (Canavan, Alexandra et al., 1997).

4.2.1. Data bias in academic corpora

TIMIT (English) (Garofolo et al., 1993), Switchboard (Godfrey and Holliman, 1993) and CallHome (American English) (Canavan, Alexandra et al., 1997) are well-established licensed speech corpora which were collected in the late 1980s and early 1990s and are held by the Linguistic Data Consortium. TIMIT (English) was collected by MIT, SRI International and Texas Instruments to be used in speech technology development and acoustic-phonetic research. It features recordings of 630 speakers of 8 “major dialects of American English”, each reading 10 phonetically rich sentences which have been phonetically transcribed and aligned. Switchboard contains 2,400 two-sided telephone conversations between 543 US American strangers on one of 70 pre-selected topics collected by Texas Instruments. CallHome features 120 unscripted 30-minute telephone conversations between friends or family members (all “native speakers of English” who grew up in the United States) and was collected by the Linguistic Data Consortium.

While all three corpora were carefully designed to capture some regional dialectal variation in US English, they are not balanced across gender groups. Further, most speakers appear to be White, though race is only recorded in the documentation of TIMIT. In the case of TIMIT this is perhaps due to convenience sampling of participants: most of the speakers were employees of Texas Instruments in Dallas which collected the corpus¹⁶. Demographic imbalances are potentially more critical for Switchboard, where only the topic of conversation, not the speech style was constrained and for CallHome where speech styles could also vary widely, and women are over-represented. As noted in 3, this gender imbalance could be indicative of a speech style imbalance. A recent analysis by Martin (2021) further confirms that Switchboard and TIMIT under-represent AAE.

4.2.2. Evaluation bias & biased benchmarks

Systems trained on biased datasets can exhibit predictive bias. But training is not the only context in which harms and biases can be introduced in the development and implementation of a machine learning system. Suresh and Gutttag (2021) use the term “evaluation bias” to describe the bias which occurs when there’s a mismatch between the benchmark data used for a particular task and the intended use population. As outlined above, some established benchmarks are unrepresentative of the potential user base of English language ASR, which include second language speakers, speakers of “non-standard” regional dialects and ethnolects and speakers who frequently code-switch between several varieties. These benchmarks are also in some ways misaligned to current ASR applications (Szymański et al., 2020). Today, ASR is widely used to transcribe conversational speech which is notoriously challenging

¹⁶<https://nvlpubs.nist.gov/nistpubs/Legacy/IR/nistir4930.pdf>

for systems designed to recognise simple commands for virtual agents in human-computer directed speech.

Particular evaluation strategies can exacerbate this kind of bias (Suresh and Guttag, 2021). Computing an aggregate word error rate across these homogeneous and/or unrepresentative test sets hides predictive bias. If Koenecke et al. (2020), for example, had computed word error rate over all speakers, the overall higher than state-of-the-art word error rate would have perhaps been attributed to the conversational nature of the recordings, rather than significant difference by speaker race. As discussed in 2, and as the CORAAL (Kendall and Farrington, 2021) recordings used by Koenecke et al. (2020) illustrate, race and gender interact in language variation. This is reflective of the concept of intersectionality originating in Black feminist thought (Crenshaw, 1991; Cooper, 2016), which recognises that interacting social categories (and axes of oppression) such as race and gender cannot be considered separately. Intersectional evaluation, then, is mindful of these interactions and can capture the differences in life experiences and linguistic behaviours between, for example, Black women and White women, rather than considering either only race or only gender. Within machine learning, this type of approach to evaluation has also been successfully applied in the context of facial analysis (Buolamwini and Gebru, 2018).

It is difficult to ascertain how much language ideologies influenced the collection of these licensed corpora in the 1980s and 1990s. At the time, they were created for a relatively narrow purpose (to research speech technologies, particularly in an academic context). It is unlikely that the researchers designing the data collection expected these resources to still be used to benchmark state-of-the-art speech recognition systems thirty years later. While incorporating some regional dialectal variation was clearly a priority, ethnic diversity or the inclusion of African American English wasn't.

The decision to use these datasets as benchmarks in the 2020s despite these limitations is, however, a choice that constitutes language policy. Just as particular language varieties or datasets are "selected" in training, they are also selected in testing. And just as training is shaped by language policy, so is testing. At first glance, Switchboard, TIMIT and CallHome fulfil the primary function of a benchmark: to allow comparison with other systems. Following Schlangen (2021)'s definition of a benchmark, they should, however, also "exemplify" the overall task of interest. A mismatch between benchmark and real-world application is therefore undesirable. More importantly, a mismatch is unexpected, as there is an implied relationship between benchmark and real-life application. The selection of an unrepresentative benchmark is shaped by beliefs about what kind of speech (and by extension, what kind of speakers) speech recognition should (be expected to) work for. Due to the evaluation bias this application of benchmarks produces, these ideologies are then fur-

ther reinforced. Failure to perform accurately on underrepresented speech not only goes undetected, but, perhaps more troublingly, is not penalised. Of course, the benchmark doesn't have to be representative of all application contexts if we choose to only use it to compare new systems to older systems. But nevertheless, the picture benchmarks provide are always partial and potentially very misleading, especially since they are almost never described in detail in the papers that use them to evaluate (Szymański et al., 2020).

5. Towards better practices

As we tried to highlight in this paper, both the curation and the use of particular speech datasets constitutes a form of language management, itself influenced by beliefs and ideologies surrounding language variation. Given the potentially far-reaching consequences of their decisions, practitioners working with speech datasets could be considered "language policy arbiters": individuals who "[wield] a disproportionate amount of power in how a policy gets created, interpreted, appropriated, or instantiated relative to other individuals in the same context" (Johnson, 2013, 100). Who gets to select which data is used in training and testing obviously depends on the broader institutional context. In a commercial context, language policy appears to be primarily driven by (linguistic) markets, and may be decided by business strategists, rather than technologists. But even in commercial contexts, researchers can reflect critically on those policies and, as work in language policy highlights, often have some leeway in the way they implement them (Hornberger and Johnson, 2007). This is not to say that we should only rely on individuals' sense of justice in the face of structural oppression, but instead to note that agentive and reflective work by technologists has potential to spark or enable broader discussion and change. In this final section, we also echo other critical work in machine learning (Paullada et al., 2021; Hutchinson et al., 2021) and argue that understanding (speech) datasets as increasingly important infrastructure is useful. It allows us to reframe the task of speech technology development from one primarily done by corporations for markets to one done by a wider range of actors for speech communities.

5.1. Speech technology design as civic design

A central obstacle to minimising predictive bias in commercial ASR systems appears to be a lack of incentive for corporations to do so. Smaller and more marginalised speech communities are unlikely to be seen as desirable markets by big technology companies, and curating very large datasets could be challenging and relatively expensive. Where proprietary datasets derived from user-data do exist, evaluating data bias is potentially difficult. It's unlikely that a technology company would (or should) be able to document or reliably infer important demographic information (such as accent, age, gender, race) about the

speakers whose data is used to create a balanced dataset (Andrus et al., 2021). Curated licensed corpora could be combined to train complex systems (as was done by Microsoft in: Xiong et al. (2017)) but since current well-established corpora only represent a small section of all English speakers, new corpora would have to be collected for this purpose. Speech technology companies could, of course, do this themselves, for example by offering payment to users (or crowd-workers) who complete a survey about their demographic background and provide speech recordings of read or naturalistic speech (see Facebook AI’s Hazirbas et al. (2021) for one of the first attempts at this method). Ultimately, however, this approach would not solve the fundamental issues arising from designing for markets.

Alternatively, we could reframe speech technology as a kind of public infrastructure and its design as civic design. Mugar and Gordon (2020, 25) define “civic design” as an approach to design that “creates the conditions for a plurality of voices and interests to be represented, accounted for, and involved in shaping the outputs and effects of public life.” Civic design is design for and with publics, rather than markets (Mugar and Gordon, 2020, 53). The notion of a “public” as a collective of people which emerges through discursive circulation of shared interests with the purpose of influencing decision-making (Mugar and Gordon, 2020, 66), has also been taken up in the analysis of language users (Muehlmann, 2014; Gal and Woolard, 1995).

Some linguistic publics intersect with the public of a nation (state), such as the Icelandic-speaking public or the Estonian-speaking public. In those cases, a (national) government (a traditional actor in language policy) shares the public’s interest in the development of speech technologies which it understands as a type of infrastructure. It can steer (and pay for) corpus development. The governments of Iceland and Estonia have both overseen design and development of open-source speech and language technology resources (corpora and models) by private and public partners (Nikulásdóttir et al., 2020). Similarly, the Welsh government has prioritised speech and language technology development and is working with universities and private sector businesses to deliver it (Welsh Language Division, 2018; Welsh Language Division, 2020). Organisations like CLARIN¹⁷ maintaining access to resources created in different contexts are also important here.

A civic design approach can also be useful for other kinds of diverse linguistic publics which do not necessarily form a “viable market”. As digital devices are becoming crucial gateways to accessing public services, jobs, and media and predictive bias could exclude many people from using them. Civic design as something that is done by a public for a public also has the potential to resolve some of the current issues with crowd-sourcing speech datasets. By carefully (and meaningfully) engaging speakers, not just as anonymous data

sources, but as co-designers who can shape the technology development process, following, for example, principles of design justice (Costanza-Chock, 2020), technology developers (private or public) would likely be able to create more representative and ultimately more useful technologies, and move away from colonial frames inherent in many drives to “spread language technologies” (Bird, 2020). With the proliferation of open-source speech technology toolkits and cheap(er) cloud computing, some publics may be able to build or modify these technologies without much or any support from governments of corporations (e.g. Masakhane¹⁸, Khandelwal et al. (2020)). Mugar and Gordon (2020) also emphasise that, in their vision, the aim of civic design by and for publics is care rather than innovation, and space for meaningful interaction between people. These values run counter to the ethos currently driving commercial technology development, but they are excellent principles in the context of technology designed fundamentally to facilitate communication.

5.2. Speech datasets as infrastructure

Whether speech technologies are approached from a perspective of civic design or not, speech datasets are, like all datasets in machine learning, infrastructure. As Hutchinson et al. (2021) point out, curation and maintenance of this infrastructure is undervalued in the machine learning community and as a result, datasets are often poorly documented and precariously stored. Fundamentally, careful curation (following a civic design model or any other model) and good documentation of speech corpora is tractable, due to the comparatively smaller size of datasets compared to, for example, large language models (Bender et al., 2021). Documentation is essential in mitigating (or even simply anticipating) predictive bias. Speech datasets (like other language datasets) need not be static but rather, like physical infrastructure, require maintenance and updating. As language (both in use and form) continuously changes, static datasets will deprecate over time and an approach in which practitioners can add or remove data from training sets in deployment may be more useful (assuming any changes are documented).

6. Conclusion

Predictive bias in speech recognition technologies is an increasingly important problem as speech recognition systems get embedded into complex algorithmic systems, with harms disproportionately falling on already marginalised speech communities. We believe that language policy is a lens that can empower technologists to mitigate data bias and recognise potential harms of biased technologies. We want to encourage practitioners to adopt this reflexive approach to better understand how language ideologies affect speech technologies and their users, and to use this understanding to build better speech technologies.

¹⁷<https://www.clarin.eu/>

¹⁸<https://www.masakhane.io/home>

7. Acknowledgements

This work was supported in part by the UKRI Centre for Doctoral Training in Natural Language Processing, funded by the UKRI (grant EP/S022481/1) and the University of Edinburgh, School of Informatics and School of Philosophy, Psychology & Language Sciences. This work was also partially made possible by an ESRC studentship (grant ES/P000681/1). We'd also like to thank Catherine Lai, Lauren Hall-Lew, Bonnie Webber, John Joseph, Joseph Gafaranga and three reviewers for their feedback.

8. Bibliographical References

- Adda-Decker, M. and Lamel, L. (2005). Do speech recognizers prefer female speakers? *9th European Conference on Speech Communication and Technology*, (January 2005):2205–2208.
- Anderson, R. C., Klobstad, C. A., Mayew, W. J., and Venkatachalam, M. (2014). Vocal fry may undermine the success of young women in the labor market. *PLoS one*, 9(5):e97506–e97506.
- Andrus, M., Spitzer, E., Brown, J., and Xiang, A. (2021). What we can't measure, we can't understand: Challenges to demographic data procurement in the pursuit of fairness. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '21, page 249–260, New York, NY, USA. Association for Computing Machinery.
- Baratta, A. (2017). Accent and linguistic prejudice within british teacher training accent and linguistic prejudice within british teacher training. *Journal of Language, Identity & Education*, 16(6):416–423. Publisher: Routledge.
- Barocas, S., Hardt, M., and Narayanan, A. (2019). *Fairness and Machine Learning*. fairmlbook.org. <http://www.fairmlbook.org>.
- Bender, E. M., Gebru, T., McMillan-Major, A., and 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*, FAccT '21, page 610–623, New York, NY, USA. Association for Computing Machinery.
- Benjamin, R. (2019). *Race after technology : abolitionist tools for the New Jim Code*. Polity Press, Newark.
- Benzeghiba, M., De Mori, R., Deroo, O., Dupont, S., Erbes, T., Jouvét, D., Fissore, L., Laface, P., Mertins, A., Ris, C., Rose, R., Tyagi, V., and Wellekens, C. (2007). Automatic speech recognition and speech variability: A review. *Speech Communication*, 49(10-11):763–786.
- Bird, S. (2020). Decolonising speech and language technology. In *Proceedings of the 28th international conference on computational linguistics*, pages 3504–3519, Barcelona, Spain (Online), December. International Committee on Computational Linguistics.
- Blodgett, S. L., Barocas, S., Daumé III, H., and Wallach, H. (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, July. Association for Computational Linguistics.
- Blommaert, J., Kelly-Holmes, H., Lane, P., Leppänen, S., Moriarty, M., Pietikäinen, S., and Piirainen-Marsh, A. (2009). Media, multilingualism and language policing: an introduction. *Language Policy*, 8(3):203–207.
- Bourdieu, P. (1977). The economics of linguistic exchanges. *Social Science Information*, 16(6):645–668.
- Bowman, S. R. and Dahl, G. (2021). What will it take to fix benchmarking in natural language understanding? In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4843–4855, Online, June. Association for Computational Linguistics.
- Buolamwini, J. and Gebru, T. (2018). Gender shades: Intersectional accuracy disparities in commercial gender classification. In Sorelle A. Friedler et al., editors, *Proceedings of the 1st conference on fairness, accountability and transparency*, volume 81 of *Proceedings of machine learning research*, pages 77–91, New York, NY, USA, February. PMLR.
- Chao, M. and Bursten, J. R. S. (2021). Girl talk: Understanding negative reactions to female vocal fry. *Hypatia—a Journal of Feminist Philosophy*, 36(1):42–59.
- Chiu, C., Sainath, T. N., Wu, Y., Prabhavalkar, R., Nguyen, P., Chen, Z., Kannan, A., Weiss, R. J., Rao, K., Gonina, E., Jaitly, N., Li, B., Chorowski, J., and Bacchiani, M. (2018). State-of-the-art speech recognition with sequence-to-sequence models. In *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 4774–4778.
- Cooper, B. (2016). Intersectionality. In Lisa Disch et al., editors, *The Oxford Handbook of Feminist Theory*, volume 1. Oxford University Press.
- Costanza-Chock, S. (2020). *Design Justice*. MIT Press.
- Crenshaw, K. (1991). Mapping the margins: Intersectionality, identity politics, and violence against women of color. *Stanford Law Review*, 43(6):1241–1299.
- De Russis, L. and Corno, F. (2019). On the impact of dysarthric speech on contemporary ASR cloud platforms. *Journal of Reliable Intelligent Environments*, 5(3):163–172.
- Denton, E. L., Hanna, A., Amironesei, R., Smart, A., Nicole, H., and Scheuerman, M. K. (2020). Bringing the people back in: Contesting benchmark machine learning datasets. *ArXiv*, abs/2007.07399.

- Eckert, P. (2012). Three Waves of Variation Study: The Emergence of Meaning in the Study of Sociolinguistic Variation. *Annual Review of Anthropology*, pages 87–100.
- Gal, S. and Woolard, K. A. (1995). Constructing languages and publics. *Pragmatics. Quarterly Publication of the International Pragmatics Association (IPrA)*, 5(2):129–138, June.
- Garnerin, M., Rossato, S., and Besacier, L. (2021). Investigating the Impact of Gender Representation in ASR Training Data: A Case Study on Librispeech. In *Proceedings of the 3rd Workshop on Gender Bias in Natural Language Processing*, pages 86–92. Association for Computational Linguistics.
- Glasser, A. (2019). Automatic speech recognition services: Deaf and hard-of-hearing usability. In *Extended abstracts of the 2019 CHI conference on human factors in computing systems*, pages 1–6, New York, NY, USA. Association for Computing Machinery.
- Green, L. J. (2002). *African american english: A linguistic introduction*. Cambridge University Press, Cambridge.
- Hargittai, E. and Shaw, A. (2015). Mind the skills gap: The role of Internet know-how and gender in differentiated contributions to Wikipedia. *Information, Communication & Society*, 18(4):424–442.
- Hazirbas, C., Bitton, J., Dolhansky, B., Pan, J., Gordo, A., and Ferrer, C. C. (2021). Towards measuring fairness in AI: the Casual Conversations dataset. April. arXiv: <http://arxiv.org/abs/2104.02821v1> [cs.CV].
- Herk, G. (2018). *What is sociolinguistics*. John Wiley & Sons, Inc, Hoboken, NJ.
- Hill Collins, P. (2000). *Black Feminist Thought: Knowledge, Consciousness, and the Politics of Empowerment*. Routledge, second edition.
- Hooker, S. (2021). Moving beyond “algorithmic bias is a data problem”. *Patterns*, 2(4):100241, April.
- Hornberger, N. H. and Johnson, D. C. (2007). Slicing the onion ethnographically: Layers and spaces in multilingual language education policy and practice. *TESOL Quarterly: A Journal for Teachers of English to Speakers of Other Languages and of Standard English as a Second Dialect*, 41(3):509–532.
- Hosoda, M. and Stone-Romero, E. (2010). The effects of foreign accents on employment-related decisions. *Journal of Managerial Psychology*, 25(2):113–132, February. Publisher: Emerald.
- Hutchinson, B., Smart, A., Hanna, A., Denton, E., Greer, C., Kjartansson, O., Barnes, P., and Mitchell, M. (2021). Towards accountability for machine learning datasets: Practices from software engineering and infrastructure. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, FAccT '21*, page 560–575, New York, NY, USA. Association for Computing Machinery.
- Irvine, J. T. and Gal, S. (2000). Language ideology and linguistic differentiation. In P. V. Kroskrity, editor, *Regimes of language: Ideologies, politics, and identities*, pages 35–84. School of American Research Press, Santa Fe.
- Johnson, D. C. and Johnson, E. J. (2014). Power and agency in language policy appropriation. *Language Policy*, 14(3):221–243.
- Johnson, D. C. (2013). *Language policy*. Palgrave Macmillan, Houndmills, Basingstoke.
- Joshi, P., Santy, S., Budhiraja, A., Bali, K., and Choudhury, M. (2020). The State and Fate of Linguistic Diversity and Inclusion in the NLP World. pages 6282–6293.
- Khandelwal, K., Jyothi, P., Awasthi, A., and Sarawagi, S. (2020). Black-Box Adaptation of ASR for Accented Speech. In *Proc. Interspeech 2020*, pages 1281–1285.
- Koch, B., Denton, E., Hanna, A., and Foster, J. (2021). Reduced, Reused and Recycled: The Life of a Dataset in Machine Learning Research. *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, 1.
- Koenecke, A., Nam, A., Lake, E., Nudell, J., Quartey, M., Mengesha, Z., Toups, C., Rickford, J. R., Jurafsky, D., and Goel, S. (2020). Racial disparities in automated speech recognition. *Proceedings of the National Academy of Sciences*, 117(14):7684–7689, April.
- Labov, W. (1990). The intersection of sex and social class in the course of linguistic change. *Language Variation and Change*, 2(2):205–254.
- Lee, D. (2021). The next Big Tech battle: Amazon’s bet on healthcare begins to take shape. *Financial Times*.
- Lev-Ari, S. and Keysar, B. (2010). Why don’t we believe non-native speakers? The influence of accent on credibility. *Journal of Experimental Social Psychology*, 46(6):1093–1096.
- Lippi-Green, R. (2012). *English with an accent language, ideology, and discrimination in the United States*. Routledge, London ; New York, 2nd ed.. edition.
- Martin, J. L. and Tang, K. (2020). Understanding racial disparities in automatic speech recognition: The case of habitual “be”. pages 626–630.
- Martin, J. L. (2021). Spoken corpora data, automatic speech recognition, and bias against african american language: The case of habitual ‘be’. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency, FAccT '21*, page 284, New York, NY, USA. Association for Computing Machinery. Number of pages: 1 Place: Virtual Event, Canada.
- Mozilla Common Voice: Community Playbook.). Community guidance for languages and variants. Accessed: 24/02/2022.
- Mozilla Common Voice: Discourse. (2019). Feedback

- needed: Languages and accents strategy. Accessed: 19/04/2022.
- Mozilla Common Voice: Discourse. (2020). Common voice languages and accent strategy v5. Accessed: 19/04/2022.
- Mozilla Common Voice. (2022). How we're making common voice even more linguistically inclusive. Accessed: 24/02/2022.
- Muehlmann, S. (2014). The speech community and beyond. In N.J. Enfield, et al., editors, *The cambridge handbook of linguistic anthropology*, pages 577–598. Cambridge University Press.
- Mugar, G. and Gordon, E. (2020). *Meaningful inefficiencies: Civic design in an age of digital expediency*. Oxford University Press, New York.
- Nikulásdóttir, A. B., Guðnason, J., Ingason, A. K., Loftsson, H., Rögnvaldsson, E., Sigurðsson, E. F., and Steingrímsson, S. (2020). Language technology programme for Icelandic 2019-2023. *LREC 2020 - 12th International Conference on Language Resources and Evaluation, Conference Proceedings*, (May):3414–3422.
- Paullada, A., Raji, I. D., Bender, E. M., Denton, E., and Hanna, A. (2021). Data and its (dis)contents: A survey of dataset development and use in machine learning research. *Patterns*, 2(11):100336.
- Pratap, V., Sriram, A., Tomasello, P., Hannun, A., Liptchinsky, V., Synnaeve, G., and Collobert, R. (2020). Massively multilingual ASR: 50 languages, 1 model, 1 billion parameters. arXiv: 2007.03001 [eess.AS].
- Raghavan, M., Barocas, S., Kleinberg, J., and Levy, K. (2020). Mitigating bias in algorithmic hiring: Evaluating claims and practices. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, FAT* '20, page 469–481, New York, NY, USA. Association for Computing Machinery.
- Raji, D., Denton, E., Bender, E. M., Hanna, A., and Paullada, A. (2021). AI and the Everything in the Whole Wide World Benchmark. *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, 1.
- Schiffman, H. (1996). *Linguistic Culture and Language Policy*. Taylor & Francis, London, United States.
- Schlangen, D. (2021). Targeting the benchmark: On methodology in current natural language processing research. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 670–674, Online, August. Association for Computational Linguistics.
- Shah, D. S., Schwartz, H. A., and Hovy, D. (2020). Predictive biases in natural language processing models: A conceptual framework and overview. In *Proceedings of the 58th annual meeting of the association for computational linguistics*, pages 5248–5264, Online, July. Association for Computational Linguistics.
- Shohamy, E. (2006). *Language policy : hidden agendas and new approaches*. Routledge, London.
- Spolsky, B. (2004). *Language Policy*. Cambridge University Press, Cambridge.
- Suresh, H. and Gutttag, J. V. (2021). A framework for understanding unintended consequences of machine learning. *CoRR*, abs/1901.10002v4.
- Szymański, P., Żelasko, P., Morzy, M., Szymczak, A., Żyła-Hoppe, M., Banaszczak, J., Augustyniak, L., Mizgajski, J., and Carmiel, Y. (2020). WER we are and WER we think we are. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3290–3295, Online, November. Association for Computational Linguistics.
- Tatman, R. and Kasten, C. (2017). Effects of talker dialect, gender & race on accuracy of bing speech and youtube automatic captions. *Proceedings of the Annual Conference of the International Speech Communication Association, INTER-SPEECH*, 2017-Augus:934–938.
- Tatman, R. (2017). Gender and dialect bias in YouTube's automatic captions. In *Proceedings of the First ACL Workshop on Ethics in Natural Language Processing*, pages 53–59, Valencia, Spain, April. Association for Computational Linguistics.
- Timming, A. R. (2016). The effect of foreign accent on employability: a study of the aural dimensions of aesthetic labour in customer-facing and non-customer-facing jobs. *Work, Employment and Society*, 31(3):409–428, April.
- Tollefson, J. W. (1991). *Planning language, planning inequality : language policy in the community*. Longman, London.
- Tripodi, F. (2021). Ms. Categorized: Gender, notability, and inequality on Wikipedia. *New Media & Society*.
- Vincent, N., Li, H., Tilly, N., Chancellor, S., and Hecht, B. (2021). Data leverage: A framework for empowering the public in its relationship with technology companies. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '21, page 215–227, New York, NY, USA. Association for Computing Machinery.
- Weinreich, U., Herzog, M., Labov, W., and Lehmann, W., (1968). *Empirical foundations for a theory of language change*, page 95–188. University of Texas.
- Welsh Language Division. (2018). Welsh language technology action plan. Technical report, Welsh Government.
- Welsh Language Division. (2020). Welsh language technology action plan: Progress report 2020. Technical report, Welsh Government.
- Wiley, T. (2012). A brief history and assessment of language rights in the united states. In W. Tollefson, J., editor, *Language Policies in Education: Critical*

Issues, pages 61–90. Routledge, New York, 2nd edition.

Xiong, W., Wu, L., Allewa, F., Droppo, J., Huang, X., and Stolcke, A. (2017). The microsoft 2017 conversational speech recognition system. *CoRR*, abs/1708.06073.

Young, V. and Mihailidis, A. (2010). Difficulties in automatic speech recognition of dysarthric speakers and implications for speech-based applications used by the elderly: A literature review. *Assistive Technology*, 22(2):99–112.

9. Language Resource References

Canavan, Alexandra and Graff, David and Zipperlen, George. (1997). *CALLHOME american english speech*. distributed via Linguistic Data Consortium, ISLRN 952-976-147-406-5.

John S. Garofolo and Lori F. Lamel and William M. Fisher and Jonathan G. Fiscus and David S. Pallett and Nancy L. Dahlgren and Victor Zue. (1993). *TIMIT Acoustic-Phonetic Continuous Speech Corpus LDC93S1*. distributed via Linguistic Data Consortium, ISLRN 664-033-662-630-6.

John J. Godfrey and Edward Holliman. (1993). *Switchboard-1 Release 2 LDC97S62*. Linguistic Data Consortium, ISLRN 988-076-156-109-5.

Kendall, Tyler and Charlie Farrington. (2021). *The Corpus of Regional African American Language*. Eugene, OR: The Online Resources for African American Language Project.