

On ideology and labor in language technologies

JUSTINE ZHANG, University of Michigan, USA

ACM Reference Format:

Justine Zhang. 2023. On ideology and labor in language technologies. 1, 1 (August 2023), 5 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

In this paper, I describe some preliminary intuitions on how thinking about labor could help us think about language technologies. In particular, I explore connections I see between two strands of scholarship, that turn from analyzing work or language towards analyzing labor or ideology. I suggest that such ideas can shed light on whether language technologies work and how they should be applied, because they attend to the terms on which linguistic competency and linguistic work are defined and debated over.

2 BACKGROUND: LANGUAGE TECHNOLOGIES AND THEIR SOCIAL PROBLEMS

For the past few years, I've been doing research in the field of natural language processing. From my vantage point, I've noticed a few controversies that NLP practitioners often encounter and struggle to address.

First, much talk around NLP has focused on whether our computational models of language actually work. Many researchers recognize that language technologies have a hopelessly partial view of the world—they deal with language, not with the rest of society, even though language and society are inextricably linked together [2, 3]. Thus, the argument goes, language models do not really understand language nor people, and this lack of understanding has dire consequences for society [1]. Nonetheless, practitioners tend to press on, often acknowledging that language models, particularly their more recent iterations, are “actually pretty good”—at modeling language, and increasingly at generating it too. This suggests a frustrating contradiction between theory and practice. If language models have such an obviously incomplete view of language and the world, then why are they so good?—or, at least, under what conditions do people act as if they are?

Talk has also turned towards the question of who we work with or for. In NLP's efforts towards increasingly big-data views of language use, it's helpful to have the support of entities with the means to collect sufficiently large datasets. The space between our intellectual pursuits and these entities' interests—what exactly are they trying to do with language and computational models of it?—can be uncomfortably thin. Recently, I collaborated with a crisis counseling service that provided us with transcripts of over a million counseling conversations. Implicitly, our computational study of conversation would help them enact their vision for how to make counseling better, premised on training and managing counselors with data-driven insights. Whether such an approach—undeniably well-intentioned—leads to better mental health outcomes is a difficult question. But when the organization's data practices—ranging from arguably defensible to definitely objectionable—became publicized [12, 16], and when some counselors and service-users objected to them, it became conspicuous to me that the vision undergirding our

Author's address: Justine Zhang, tisjune@umich.edu, University of Michigan, School of Information, Ann Arbor, USA.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2023 Association for Computing Machinery.

XXXX-XXXX/2023/8-ART \$15.00

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

valuable research dataset was a *particular* one, that those the service was supposed to support did not necessarily share. In trying to make work better, whose vision of work are we working with?

Such questions—of language use and of work—have become entwined in increasingly public ways. Consider the recent Writer’s Guild of America strike, where AI is a salient point of contention between writers and studios [22]. Can AIs replace writers in generating screenplays?—and or, what does it say about screenwriting and media production as imagined by studios, that this is even a question? Alternatively, what does it say about how screenwriters imagine their work, such that such a proposition is so objectionable to them as to be contested in labor action?—and or, what does it say about their increasingly precarious industry, centered in a city that has become increasingly unaffordable to live in?

Or, consider the National Eating Disorders Association’s recent decision to deploy a chatbot, days after letting go of the employees staffing their helpline [24]. The chatbot promptly started giving harmful advice, and was hastily taken offline [14]. Here, it was interesting that the organization, and the chatbot’s designers, framed the issue as mainly technical: the chatbot had been programmed with more guardrails in place that would’ve constrained its behavior had the wrong model not been rolled out; perhaps the implication was that a better and safer model would’ve avoided the fallout. Such explanations belie a particular vision of what mental healthcare entails, what sorts of difficulties are involved (and could be engineered through), and how mental healthcare workers should be compensated for their efforts. These were possibly all concerns for the former employees, who had, incidentally, unionized days before getting fired.

3 RELATED WORK: LANGUAGE IDEOLOGIES AND LABOR

In these events, as well as in the debates taking place within NLP, we observe that different parties have different views of how language works, and how people use language to do work. We also note that these different views imply different political positions—that are interested in making a living wage, in doing fulfilling creative or care work, in more efficiently and profitably producing text. Enacting these visions of language and of work would result in different outcomes for different people—as the acrimonious nature of the writers’ strike makes clear. Examining these multiple views and their consequences is therefore of utmost importance for people who develop technologies—and hence enact particular views—of work and of language.

As linguistic anthropologists Irvine and Gal put it [8], “Statements about language are never only about language—and they are never only statements.” This quote summarizes a recent turn in scholarship on language and society towards examining *ideologies* of language [6, 18]. The term “language ideology” has multiple definitions in the literature, but Irvine’s is particularly apt—ideologies are partial and politically interested views [10].

Scholars focused on language ideologies aim to expose the social, economic and political conditions surrounding particular views of language—as structured code, as valuable resource, as marker of identity—and their users—as more or less competent, intelligent or human. Works oriented towards this goal note that assessments of linguistic (in)competence are often premised on racialized speakers and white listeners [7]; that linguistic diversity is a both problem and an asset—for nationalist projects [13] or for expanding transnational markets [5]; that conversations are a core business practice—for particular models of capital, labor and corporate governance [15]. Such work also unsettles assumptions underlying existing linguistics scholarship, which focuses on modeling grammar and syntax, or on analyzing language use in particular social situations. Here, scholars do not oppose studies of form and of use so much as urge that they be properly contextualized [18]. What’s left out, and what are the consequences of that, when linguistic codes or language use situations are carved out of the broader social world, for academic study or technological development?

To me, this turn towards language ideologies rhymes with Greenbaum’s proposal for CSCW scholarship to focus on labor beyond work [9]. The field of CSCW has thought a lot about designing ways of improving work, cooperative or otherwise. Perhaps a labor-oriented framing would find such studies incomplete if they did not

also attend to how units of work are identified for improvement in the first place, and the social conditions in which work happens, at multiple scales. Greenbaum calls for a critical examination of when and why CSCW concepts happen to be “in the right place at the right time,” relative to certain economic conditions or models of management. We could extend that suggestion to linguistic concepts and technologies, too.

4 EXPLORATORY SKETCHES

I now revisit the controversies around NLP I described above. In two sketches, I explore ways that attending to labor and ideology can shed light on these controversies, by clarifying what we mean when we debate over whether language technologies are competent, or whether they’ll make work better.

4.1 Labor markets for competent language users

First, I return to the broad sense that language models, despite their obvious theoretical shortcomings, are “pretty good.” Here, I focus on products like ChatGPT that generate text: increasingly, the output seems to resemble coherent language, and is read as meaningful interpersonal engagement (despite experts cautioning against that). The increasing prominence of generative AI underlines the importance of clarifying what we mean when we say they’re good at using language. Linguistic and philosophical accounts [2]—that try to map competence to constructs like meaning and intent—only go so far: they do not adequately describe the broad range of practical situations where such technologies apparently succeed or fail.

For scholars attentive to ideology, situating assessments of linguistic competence is crucial. For instance, dominant theories of competence in sociolinguistics presume the existence of homogeneous speech communities and implicitly treat linguistic differences as deficient [7]. In parallel with contemporaneous social projects, these theories suggest that everyone can be recognized as a legitimate language user, provided they learn appropriate language practices. Left open are the questions of who decides which language practices are appropriate, and who is doing the recognizing. Accordingly, in their research on educational institutions, Flores and Rosa [7] find that particular types of students are often classified as linguistically lacking, despite regularly demonstrating linguistic skill. The problem, this scholarship suggests, resides in the “power dynamics that lead to the perpetual perception of their linguistic practices as deficient:” the students are members of racialized populations being assessed by white listeners, in institutions that reinforce existing social hierarchies. (We might also imagine an inverse situation, in which faulty language-using entities are perpetually perceived as competent by listeners who are somehow invested in their proliferation.)

In what types of situations do we imagine language-generating AIs? Recent work in linguistic anthropology has tracked how language is increasingly cast as economically valuable, locating such discourses in a broader political economy shaped by processes of late capitalism [5]. For instance, there has been a rise in the number of tertiary sector jobs, where language is central not just as the process but as the product of work. While such scholarship often focuses on such sites as transnational call centres and tourism, we may also think to paradigmatically corporate jobs characterized by the production of linguistic artefacts: here I consider the marketing materials, memos, emails and endless slide decks that my friends in consulting complain about, and would probably appreciate AI help in generating. Depending on the job, the language use that’s desired resembles dominant varieties like radio English, or varieties iconic of some form of ethnicness, or what former tech startup employee Anna Wiener describes as “garbage language [21].” Linguistic anthropologists Duchêne and Heller [5] suggest that the “new centrality of linguistic form and practice in economic production” may affect our ideas about language. Indeed, it would be interesting to examine how judgements of ChatGPT’s quality, or NLP researchers’ conceptualizations of linguistic competence, are tied to the present landscape of linguistic labor.

4.2 Managing linguistic work

Next, I return to the question of whose vision of work we're helping to improve. In trying to build computational models that captured ineffable qualities of crisis counseling conversations, I was increasingly struck by the coarseness of my tools, relative to the intricacies of the domain. Good counselors (as the training materials emphasized) come to empathetic understandings of each individual calling in. My big-data approach, by contrast, was supposed to quantify some quality over all conversations; unlike an empathetic listener, this view depended on presumptions. I wondered if a better approach—one that, incidentally, might not have constituted a gross violation of informed consent—would've been to directly ask counselors how they worked. As Suchman [20] puts it, "Why do computer scientists go about making up all these typologies of interaction? Aren't the typologies used by practitioners themselves before we go to work, as designers, good enough?"

The immediate object of Suchman's critique was a system called The Coordinator that aimed to improve cooperative work by enforcing communicative clarity: workers wishing to talk with each other would use the system to explicitly label the speech act performed as one of a small number of categories. The Coordinator's designers emphasized the theoretical underpinnings of their use of speech act theory—such a system would better reflect that language use is about doing things with each other, beyond manipulating symbols [23]. Suchman, however, notes that the theory of language and action that's operationalized is premised on categorization and systematization, rather than on the irreducibly contingent nature of social interaction. As such, it "carries with it an agenda of discipline and control over organization members' actions."

Like other products and work practices, language is subject to taylorist management techniques [5]. Strategies for decomposing language use into standardized units to maximize efficiency involve specifying or benchmarking conduct across multiple linguistic dimensions, e.g., prosodic, lexical and discursive. Boutet [4] surveys examples across the history of linguistic work management, from early 20th-century telephone operators to modern-day call center workers; in each case, professional conduct was structured by various scripts, routines and time constraints—and constantly monitored by higher-ups. Duchêne and Heller [5] suggest that in its push towards standardization, linguistic taylorism decouples language from identity, moving it towards "being understood as a skill that we don't even need a human to be able to produce." This naturally prompts questions about whether and how NLP and CSCW efforts can be read as taylorizing talk, and how such management practices further feed into conceptualizations about AI.

Language use, however, is arguably harder to manage and standardize than steelmaking. This leads to some odd work configurations, as humans inevitably fill in the missing mass [11]—for instance, a recent Guardian article documents the writer's experience as a human propping up a conversational AI designed to field real estate questions in a humanlike manner [17]. The attendant judgements about what is or isn't automatable are revealing. In an ethnographic study of research teams aiming to engineer speech-recognition systems to detect mental illness, Semel [19] notes that diagnostic procedures are seen as sufficiently mechanical, skillless and "custodial" as to be automated, in contrast to the essential humanness of actual psychiatric treatment; of course, it was humans who still needed to generate the training data. It would be interesting to examine how such judgements about labor are made and operationalized in other sites of language technology development.

5 CONCLUSION

This paper has explored how ideas on language ideologies and labor could be knit together, and used to enrich thinking around CSCW and NLP. Future work could much more substantively elaborate on the theoretical concepts I've roughly brought together, or empirically study the ways in which language and technology development relate to labor conditions. Such efforts would contribute valuable understandings of how language technologies—and scholars of language and technology—feature in the ways that labor is "divided [and] coordinated back together"[9], and in the ways that language is carved out of society and put back in.

REFERENCES

- [1] Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 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*. 610–623.
- [2] Emily M Bender and Alexander Koller. 2020. Climbing towards NLU: On meaning, form, and understanding in the age of data. In *Proceedings of the 58th annual meeting of the association for computational linguistics*. 5185–5198.
- [3] Yonatan Bisk, Ari Holtzman, Jesse Thomason, Jacob Andreas, Yoshua Bengio, Joyce Chai, Mirella Lapata, Angeliki Lazaridou, Jonathan May, Aleksandr Nisnevich, et al. 2020. Experience Grounds Language. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 8718–8735.
- [4] Josiane Boutet. 2012. Language Workers: Emblematic Figures of Late Capitalism. *Language in late capitalism: Pride and profit* 1 (2012), 207.
- [5] Alexandre Duchêne and Monica Heller. 2012. *Language in late capitalism: Pride and profit*. Vol. 1. Routledge.
- [6] Alessandro Duranti. 2003. Language as culture in US anthropology: Three paradigms. *Current anthropology* 44, 3 (2003), 323–347.
- [7] Nelson Flores and Jonathan Rosa. 2022. Undoing competence: Coloniality, homogeneity, and the overrepresentation of whiteness in applied linguistics. *Language learning* (2022).
- [8] Susan Gal and Judith T Irvine. 2019. *Signs of difference: Language and ideology in social life*. Cambridge University Press.
- [9] Joan Greenbaum. 1996. Back to Labor: Returning to labor process discussions in the study of work. In *Proceedings of the 1996 ACM conference on Computer supported cooperative work*. 229–237.
- [10] Judith T Irvine. 2022. Revisiting theory and method in language ideology research. *Journal of Linguistic Anthropology* 32, 1 (2022), 222–236.
- [11] Bruno Latour. 1992. Where are the missing masses? The sociology of a few mundane artifacts. *Shaping technology/building society: Studies in sociotechnical change* 1 (1992), 225–258.
- [12] Alexandra S Levine. 2022. Suicide hotline shares data with for-profit spinoff, raising ethical questions. *Politico* (2022). <https://www.politico.com/news/2022/01/28/suicide-hotline-silicon-valley-privacy-debates-00002617>
- [13] Stephen May. 2017. Language, imperialism, and the modern nation-state system. *The Oxford handbook of language and society* (2017), 35–53.
- [14] Lauren McCarthy. 2023. EA Wellness Chatbot Is Offline After Its ‘Harmful’ Focus on Weight Loss. *The New York Times* (2023). <https://www.nytimes.com/2023/06/08/us/ai-chatbot-tessa-eating-disorders-association.html>
- [15] Bonnie McElhinny. 2012. Silicon Valley Sociolinguistics?: Analyzing Language, Gender and Communities of Practice in the New Knowledge Economy 1. In *Language in late capitalism*. Routledge, 230–260.
- [16] Joanne McNeil. 2022. Crisis Text Line and the Silicon Valleyfication of Everything. *Motherboard* (2022). <https://www.vice.com/en/article/wxdpym/crisis-text-line-and-the-silicon-valleyfication-of-everything>
- [17] Laura Preston. 2022. Becoming a chatbot: my life as a real estate AI’s human backup. *The Guardian* (2022). <https://www.theguardian.com/technology/2022/dec/13/becoming-a-chatbot-my-life-as-a-real-estate-ais-human-backup>
- [18] Jonathan Rosa and Christa Burdick. 2017. Language ideologies. *The Oxford handbook of language and society* (2017), 103–123.
- [19] Beth Michelle Semel. 2019. *Speech, signal, symptom: machine listening and the remaking of psychiatric assessment*. Ph.D. Dissertation. Massachusetts Institute of Technology.
- [20] Lucy Suchman. 1993. Do categories have politics? The language/action perspective reconsidered. *Computer supported cooperative work (CSCW)* 2 (1993), 177–190.
- [21] Anna Wiener. 2020. *Uncanny valley: A memoir*. MCD.
- [22] Alissa Wilkinson. 2023. Hollywood’s writers are on strike. Here’s why that matters. *Vox* (2023). <https://www.vox.com/culture/23696617/writers-strike-wga-2023-explained-residuals-streaming-ai>
- [23] Terry Winograd. 1986. A language/action perspective on the design of cooperative work. In *Proceedings of the 1986 ACM conference on Computer-supported cooperative work*. 203–220.
- [24] Chloe Xiang. 2023. Eating Disorder Helpline Fires Staff, Transitions to Chatbot After Unionization. *Motherboard* (2023). <https://www.vice.com/en/article/n7ezkm/eating-disorder-helpline-fires-staff-transitions-to-chatbot-after-unionization>