Queer NLP: A Critical Survey on Literature Gaps, Biases and Trends

Sabine Weber[1,2,*] **Angelina Wang**[1,3,†] **Ankush Gupta**[1,4,5,†] **Arjun Subramonian**[1,†]
Dennis Ulmer[1,6,†] **Eshaan Tanwar**[1,7,†] **Geetanjali Aich**[1,8,†] **Hannah Devinney**[1,9,†]
Jacob Hobbs[1,†] **Jennifer Mickel**[1,10,†] **Joshua Tint**[1,11,†] **Mae Sosto**[1,12,†]
Ray Groshan[1,13,†] **Simone Astarita**[1,†] **Vagrant Gautam**[1,14,†] **Verena Blaschke**[1,15,16,†]
William Agnew[1,17,†] **Wilson Y Lee**[1,18,†] **Yanan Long**[1,19,†]

[1]Queer in AI [2]University of Bamberg [3]Cornell Tech [4]IIIT-DELHI [5]PayGlocal
[6]ILLC, University of Amsterdam [7]University of Copenhagen [8]University of Massachusetts Amherst
[9]Linköping University [10]EleutherAI [11]Arizona State University
[12]Centrum Wiskunde & Informatica (CWI) [13]University of Maryland, Baltimore County
[14]Heidelberg Institute for Theoretical Studies (HITS) [15]LMU Munich [16]MCML
[17]Carnegie Mellon University [18]HubSpot [19]StickFlux Labs

Abstract

Natural language processing (NLP) technologies are rapidly reshaping how language is created, processed, and analyzed by humans. With current and potential applications in hiring, law, healthcare, and other areas that impact people's lives, understanding and mitigating harms towards marginalized groups is critical. In this survey, we examine NLP research papers that explicitly address the relationship between LGBTQIA+ communities and NLP technologies. We systematically review all such papers published in the ACL Anthology, to answer the following research questions: (1) What are current research trends? (2) What gaps exist in terms of topics and methods? (3) What areas are open for future work? We find that while the number of papers on queer NLP has grown within the last few years, most papers take a reactive rather than a proactive approach, pointing out bias more often than mitigating it, and focusing on shortcomings of existing systems rather than creating new solutions. Our survey uncovers many opportunities for future work, especially regarding stakeholder involvement, intersectionality, interdisciplinarity, and languages other than English. We also offer an outlook from a queer studies perspective, highlighting understudied topics and gaps in the harms addressed in NLP papers. Beyond being a roadmap of what has been done, this survey is a call to action for work towards more just and inclusive NLP technologies.

anonymous/anonymized-github

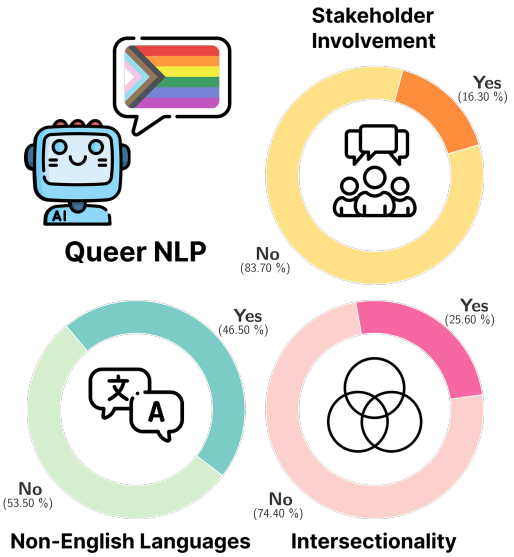

Figure 1: The majority of queer NLP papers published in the ACL Anthology are focused on English, disregard intersectionality and omit stakeholders.

1 Introduction

Natural language processing (NLP) is a field that focuses on the interaction between computers and human language, aiming for machines to analyze, understand, and generate natural language. As NLP technologies become increasingly integrated into everyday life, users' direct interaction with them has expanded dramatically. It is critical that these systems function well, do not reinforce social harms, and maintain the privacy of all users, including queer individuals, who have historically been excluded from technological development and research (Scott et al., 2017; Queer in AI et al., 2023).

In this survey, we examine the state of NLP research as it relates to LGBTQIA+ topics. We

*First author, please send correspondence to sabine.weber@uni-bamberg.de.

†Equal contributions, alphabetical order by first name

use the term *queer* as a shorthand for the broad spectrum of identities represented under the LGBTQIA+ umbrella (see Glossary), while recognizing the diverse and heterogeneous experiences within this community. Although there has been a significant amount of work in NLP on bias more broadly (Stanczak and Augenstein, 2021; Field et al., 2021; Goldfarb-Tarrant et al., 2023), we foreground the queer community and we survey papers that discuss the distinct social impacts queer people face in the context of NLP technologies.

As evidenced by an increasing amount of articles and its first archival event for ACL (Queer in AI @ NAACL 2025 (Pranav et al., 2025)), queer NLP is growing as a topic of study, which necessitates synthesizing and systematizing the knowledge in this subfield to understand current trends, identify gaps, and assess promising future directions. These objectives are critical for NLP to meaningfully engage with queer communities and their linguistic realities, and foster more inclusive and socially responsible systems. Therefore, we systematically survey all queer NLP papers from the ACL Anthology, annotating them for the tasks, types of data, and queer groups and harms they focus on, as well as their motivations, methods, and engagement with intersectionality and stakeholders.

We provide an overview of the findings per NLP topic (§4), and find more broadly:

1. **Trends:** While there is increased interest in queer NLP, approaches are often guided by the affordances of current models, focusing on pronouns and stereotypes in terms of topics (§5.1), and template-based approaches and data augmentation in terms of methods (§5.2).

2. **Gaps:** Approaches rarely involve stakeholders, take an intersectional view, or examine languages other than English (Figure 1; §5.3).

3. **Future directions:** Beyond closing these methodological gaps, we call on the field to address larger, structural gaps, such as the possibility to opt out of NLP systems and dynamic evaluations which allow for temporal and contextual differences (§6).

2 Related Work

Gender and sexual orientation have been examined in NLP in a number of prior studies. Devinney et al. (2022) survey 200 articles on gender bias in NLP to examine how the field conceptualizes gender, Cao and Daumé (2021) review 150 articles on gender in co-reference resolution, and Hobbs (2025) surveys 55 papers on sexuality bias in NLP. While Zhou (2024) offers a broad overview of queer harms and structural barriers in NLP, our work extends this with an in-depth survey of the field.

Outside of queer NLP, several papers surveys research on social groups and issues in NLP, including gender bias (Stanczak and Augenstein, 2021) and racial bias (Field et al., 2021). Others survey specific terms, aiming to investigate how they are used in the field, e.g., "low resource" (Nigatu et al., 2024), "intersectionality" (Ovalle et al., 2023b; Wang et al., 2022), and "democratization" (Subramonian et al., 2024). Many surveys, like ours, set their goal to shape the future of burgeoning subfields of NLP, such as African NLP (Alabi et al., 2025) and interpretability (Mosbach et al., 2024). To our knowledge, our paper is the first to systematically survey all queer NLP papers within the ACL Anthology across diverse NLP tasks.

3 Methodology

3.1 Paper Selection

We, a group of NLP scholars (see Ethics for our positionality statement), curated a collection of papers following the PRISMA methodology (Moher et al., 2015; Shamseer et al., 2015). We searched the ACL Anthology for papers containing a subset of the terms used by Taylor et al. (2024): *queer, aromantic, gender non*, lgbt*, agender, glbt, lesbian, gay, bisexual, transgender*. The ACL Anthology search algorithm returned 3,864 entries. This large number is due to the search algorithm often returning the same paper several times for the same query, and to many papers containing multiple search terms. The search also yielded irrelevant papers, e.g. by authors named Gay, or papers using "gay" as an example of language change rather than engagement with the queer community. Manual filtering for relevant papers lead to a set of 55 papers. Subsequently, the authors collectively expanded the paper list based on our knowledge of queer NLP and new papers found using scholarly search engines such as Semantic Scholar. We also added 19 papers from the ACL conference 2025 which took place after our initial search.

We excluded papers with no explicit mention of the queer community or subsets thereof. If a paper only mentioned sexuality as one of many demo-

graphic categories without explicitly engaging with it, we excluded it. Questions about whether a paper counted as "queer NLP" were resolved through discussion. In refining our collection of papers, we aimed to confidently answer our research questions: describe trends of both what *is* included in queer NLP papers and what tends to *not* be included. Our final survey covers 86 ACL papers.

While we make statistical claims only about the papers contained within the ACL anthology, our search brought up many papers that are preprints or published in other venues, which we cite when their perspective is necessary to gain a full understanding of the field. We present all papers as a living resource in an open-source repository[1].

3.2 Annotation Process

We annotate each paper in our survey along several axes. First, we annotate its NLP task or topic (e.g., question answering), its subcategory (e.g., pronoun-inclusive QA), the motivation and the method. Second, we collect which queer groups are explicitly named and what kinds of harm are addressed (if mentioned). Lastly, we focus on the language (see Figure 2), geographical region, and domain of the data used, whether intersectionality is explicitly addressed,[2] the involvement of any stakeholders and limitations named in the paper. We then group all papers into the 7 categories used in Section 4 using inductive coding. Papers can belong to multiple categories. While each paper is only read and annotated by a single person, to calculate inter-rater reliability we randomly select 11 papers to be annotated by a second person. We report high reliability across our three focus themes of stakeholder involvement, intersectionality, and language diversity, with 90.09% agreement (Cohen's Kappa = .79 and .62 for the first two, respectively) and perfect agreement for language diversity.

4 Survey Findings

We begin with Language Modeling, examining bias evaluations and bias mitigation strategies. We then turn to prominent use cases of NLP: text classification (specifically sentiment analysis, toxicity, and hate speech classification), information extraction and access (specifically, coreference resolution, media analysis, and question answering chatbots), machine translation, and speech processing. Lastly, we survey position papers. We focus on ACL papers (see Figure 3), but also discuss non-ACL papers where their inclusion helps to gain a fuller understanding of the research landscape. We mark all non-ACL papers with an asterisk (*) in this section for a clearer distinction.

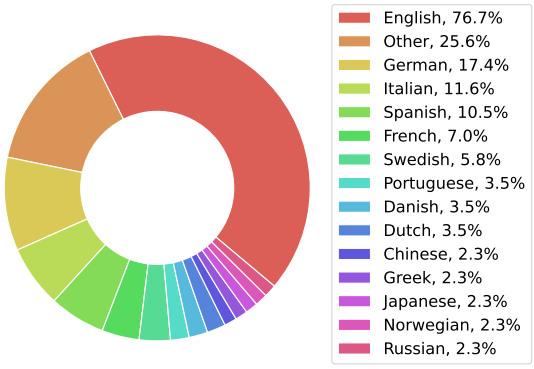

Figure 2: Percent breakdown of the 35 languages represented across papers. The 'Other' category aggregates all 20 languages that appeared in only one paper. The total percent is over 100% due to papers on multiple languages. See Appendix B for a full list of languages along with paper counts.

4.1 Language Models: Evaluation

Large Language Models (LLMs), which increasingly serve as the backbone of modern NLP systems, are trained on massive, mostly uncurated online text corpora. LLMs thus tend to inherit and reproduce social biases from these corpora (Mickel et al., 2025)*. For the queer community, such biases manifest in the form of erasure, stereotyping, misgendering, hate speech and under- or misrepresentation. In addition to being used directly, LLMs frequently serve as the foundation for other NLP or multimodal models that amplify and propagate these harms. A first step to mitigating these harms is to make them visible through evaluation.

Template-based methods. The most widely-used approach to elicit LLM bias, template-based methods involve prompting LLMs with templatic sentences containing subject placeholders, which can be filled with terms representing various social groups, e.g., identity terms relating to gender and sexuality, names, or pronouns (Nozza et al., 2022; Dhingra et al., 2023; Hossain et al., 2023; Sosto

[1] anonymous/anonymized-github

[2] Intersectionality is a critical framework combining inquiry and praxis to understand how intersecting systems of power reinforce social inequalities (Collins and Bilge, 2020). Thus, we also discuss how papers address complexity (e.g., overlapping groups) and contextualize queer bias, harms, and participation with regards to inequality, power, and justice.

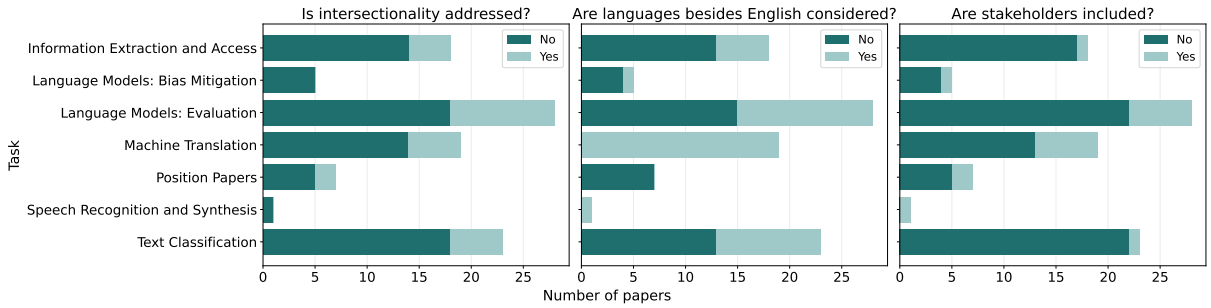

Figure 3: Comparing paper categories by (a) intersectionality, (b) language diversity, and (c) stakeholder inclusion.

and Barrón-Cedeño, 2024*; Ovalle et al., 2023a*).

By analyzing the outputs conditioned on these different terms, the authors report systematic differences that often reflect social biases or imbalances. Notably, Felkner et al. (2023) take a community-centered template-based approach by collaborating with queer annotators to craft stereotype-based factual-counterfactual sentence pairs, such as "LGBTQ people are sinful" versus "Heterosexual people are sinful".

Intrinsic and extrinsic measures. Given the complexity of assessing LLM bias in queer-related contexts, many studies rely on intrinsic bias measures. Building on the template-based methods discussed in the previous paragraph, common approaches include direct statistical comparison between expected and actual answers (Bunzeck and Zarrieß, 2024; Xie et al., 2024), as well as comparison of probabilistic scores such as pseudo-log-likelihood or perplexity for target–attribute pairings within sentence templates (Felkner et al., 2023; Hossain et al., 2023; Bergstrand and Gambäck, 2024; Xie et al., 2024; Ovalle et al., 2023a*). In contrast, other studies perform group-level comparisons using downstream tasks, yielding more extrinsic bias measures. For instance, sentiment analysis has been widely applied to examine polarity (negative, neutral, or positive) in LLMs outputs (Kiritchenko and Mohammad, 2018; Sosto and Barrón-Cedeño, 2024*) while regard analysis (Sheng et al., 2019) has been used to capture perceived respectfulness in generated text (Ovalle et al., 2024; Dhingra et al., 2023). Similarly, harmful language detection tools such as Perspective API[3] are frequently used to detect toxicity and offensive language (Nozza et al., 2022; Sosto and Barrón-Cedeño, 2024*) , alongside lexicon-based approaches such as HONEST (Nozza et al., 2021*), which builds on HurtLex (Bassignana et al., 2018*) to quantify offensiveness. These methods are rarely tailored to capture the specific nuances of queer or intersectional identities, which limits their ability to provide a full picture of harmful biases in LLMs.

[3] https://www.perspectiveapi.com

4.2 Language Models: Bias Mitigation

Discovering bias is a necessary precursor to mitigating it, the focus of two of the surveyed papers. Dhingra et al. (2023) apply a post-hoc method that combines chain-of-thought prompting with SHAP analysis (Lundberg and Lee, 2017*) to identify words that reduce social regard. After detecting low-regard words, they query an LLM for explanations and then re-prompt it to rephrase the sentence, removing those words while preserving meaning. Tang et al. (2024) create a debiasing dataset using Counterfactual Data Augmentation (CDA) to debias biased contexts. They construct anti-biased descriptors e.g., replacing stereotypical associations with counter-stereotypical ones, to generate texts with them, creating a dataset for fine-tuning.

4.3 Text Classification

Text classification is a technique used to assign discrete labels to textual data such as words, sentences, or entire documents. This process typically involves training an algorithm on labeled examples so it can recognize patterns and categorize large collections of unstructured text. In recent years, detecting transphobic and homophobic hate speech has become a prominent focus, with numerous shared tasks addressing this challenge across multiple languages (Chakravarthi et al., 2022, 2023, 2024; Nozza et al., 2023; Bel-Enguix et al., 2023*; Gómez-Adorno et al., 2024*; Jiménez-Zafra et al., 2023*; García-Baena et al., 2024*).

Oversensitivity towards queer language. Several papers show that models regularly misclassify words like "gay," "trans," or "queer" as toxic, even when used neutrally or positively. Some papers show that text classification systems are also biased against language used among queer people (Wu and Hsieh, 2017; Ramesh et al., 2022), and references to queer identities (Tint, 2025; Ungless et al., 2023; Zhang et al., 2020; Zueva et al., 2020).

Biased datasets. The deficiencies of text classification systems can often be traced back to their training data, and better resources are thus one potential solution. For example, Ljubešić et al. (2020) introduce the LiLaH emotion lexicon for Croatian, Dutch, and Slovene, showing more consistent results on LGBTQ topics in emotion classification. Waldis et al. (2024) create a resource for German text classification for seven different tasks, including gender-neutral and gender-fair formulations.

In hate speech and toxicity detection, Engelmann et al. (2024) and Sahoo et al. (2022) construct datasets to detect dehumanizing language, including with respect to queer communities. Dacon et al. (2022) focus entirely on LGBTQIA+ individuals, creating a real-world dataset of harmful online conversational content. Wiegand and Ruppenhofer (2024) provide a dataset and train predictors for more subtle language where a group is portrayed as deviating from a norm. In a similar vein, Lu and Jurgens (2022) introduce a dataset to detect trans-exclusionary language by TERFs (see Glossary).

Non-English hate speech detection. Despite the prevalence of hate speech in a diversity of languages, a relatively small number of papers study languages other than English. Work has been done on Mexican Spanish (Vásquez et al., 2023), Russian (Zueva et al., 2020), Malayalam and Tamil (Sharma et al., 2023*), and Gujarati, Kannada, and Telugu (Kumaresan et al., 2024). Locatelli et al. (2023) study homotransphobic speech on Twitter in seven languages, concluding that current automated moderation systems are insufficient and have large differences in sentiment across languages.

4.4 Information Extraction and Access

NLP techniques are used to facilitate information access by extracting information about entities and events from text, including the tasks of named entity recognition and coreference resolution. While some papers use these tools to study how queer people are discussed in texts, others analyze and mitigate the shortcomings of the tools themselves. In this section, we also include papers on the use of chatbots as information access points for queer users.

Coreference resolution. Papers evaluating coreference resolution highlight the harms of misrecognizing people's genders, disconnecting someone's name from their pronouns, or linking their name to a different pronominal chain. Most papers name trans and non-binary people as the targets of potential harm (Dev et al., 2021), while Cao and Daumé III (2020) state that cis people can also be harmed by misgendering, and Gautam et al. (2024b) argue from a standpoint of universal design without naming specific communities. To counter the effects of the underrepresentation of queer identities in training data, some papers create synthetic data to re-balance representation. For example, Cao and Daumé III (2020) create a new dataset that better represents trans and non-binary people. While most work focuses on gender-neutral language, Gautam et al. (2024a) and Cao and Daumé III (2020) also consider neopronouns.

Queer terms in media. Some works use information extraction tools to evaluate texts written about the queer community. Locatelli et al. (2023) and Andersen et al. (2024) analyze English and Mexican Spanish homotransphobic speech on Twitter, respectively, noting a decrease in the use of many derogatory words over time, while the overall vocabulary for queer terms has increased. CH-Wang and Jurgens (2021) also study Twitter and Reddit data, finding that the use of gender-neutral expressions (e.g., partner, spouse, folks) has risen steadily, with the sharpest uptake among heterosexual, liberal, and higher-socioeconomic-status users, indicating that queer people are becoming increasingly discussed in mainstream social media. Hicks et al. (2016) use Twitter data to find terms for gender expression that do not appear in the US National Transgender Discrimination Survey. In general, these papers demonstrate the potential of NLP tools in analyzing the use of queer terms in society more broadly.

Chatbots for queer users. Research in this domain seeks to uncover biases and gaps in NLP models regarding queer topics. An application of chatbots in the queer health context is their role in offering gender-affirming care and sexual health education. Najafali et al. (2023)* investigate Chat-

GPT's ability to provide evidence-based recommendations for gender affirming surgery according to the standards of care of the World Professional Association for Transgender Health. Although ChatGPT can accurately describe aspects of gender diversity and treatment for gender dysphoria, the study highlighted gaps in completeness, reliability, and explicit reference to authoritative guidelines. Koss et al. (2025)* present a similar user study, showing low trust in ChatGPT's ability to provide gender affirming surgery information. Furthermore, Shaier et al. (2023) find that demographic factors such as gender and sexuality affect outputs in biomedical question answering where such attributes are irrelevant. Motivated by the high engagement of queer youth in social media, Brixey et al. (2017) develop a chatbot for HIV prevention and care for queer youth. They create a corpus of 3,000 question-and-answer pairs from reliable sources and train an NPCEditor chatbot (Leuski and Traum, 2010) which they deploy on Facebook Messenger. However, their work does not analyze their success in reaching queer youth.

Some works study chatbot use by queer individuals for mental health and emotional support. Lissak et al. (2024), Edwards et al. (2021), and Ma et al. (2024) conclude that these systems have positive effects, but are generic and culturally insensitive outside the U.S. While unsuitable to replace human emotional intelligence, the authors argue that a chatbot may be a useful tool for initial support.

4.5 Machine Translation

Machine translation can be challenged by differing conceptualizations of gender and sexual orientation across cultures. Moreover, grammatical systems vary in the extent to which gender is encoded in the surface form of language, e.g., when a source language has non-gendered pronouns, but the target language does not have an equivalent.

Gender-fair language. Gender-fair language comprises both gender-inclusive and gender-neutral language (Paolucci et al., 2023). The level of inclusion of non-binary identities varies across papers, with two papers explicitly leaving non-binary gender expressions to future work (Stewart and Mihalcea, 2024; Robinson et al., 2024). Several papers however focus on neopronouns: Piergentili et al. (2024) introduce NEO-GATE, an extension of the GATE benchmark (Rarrick et al., 2023*), for translation from English to Italian, and Piergentili et al. (2025) adapt chain-of-thought prompting to improve gender-neutral translation. Lauscher et al. (2023) incorporate nounself, emojiself, numberself, and other neopronouns in addition to gender-neutral (they/them) and gendered (he/she) pronouns while evaluating MT models.

Several papers tackle rewriting and post-editing: Lardelli and Gromann (2023) recruit translators to post-edit English to German machine translations into gender-fair language, while Veloso et al. (2023) create a rule-based and a neural-based rewriter for Portuguese. Amrhein et al. (2023) propose a method for "leveraging inherently biased NLP models" to train neural rewriting models.

Anglocentrism. Excluding Veloso et al. (2023), every language pair studied included English. A high proportion of the surveyed works study translation only in one direction, typically with English as the source language. English-German was the most frequently studied language pair, followed by English to Romance languages (French, Spanish, Italian). Only three papers (Lauscher et al., 2023; Robinson et al., 2024; Ghosh and Caliskan, 2023*) include low-resource languages.

4.6 Speech Recognition and Synthesis

Processing *spoken* language is relevant for transgender and non-binary voices, since they challenge normative assumptions about gender (Zimman, 2018*, 2021*). Nevertheless, speech technology research focuses on binary genders and cisgender speech (Sanchez et al., 2024*). Voice also affects the expression and perception of sexual orientation (Levon, 2007*; Simpson and Weirich, 2020*).

Queer voices and topics. Very few speech datasets explicitly contain speech by queer individuals, although such data might be included and simply not labeled. Although this hinders research on the robustness of speech tools, publicly sharing such metadata comes with privacy risks. Siegert et al. (2025)* present a German speech dataset of podcasts and videos by people who publicly identify as queer.[4] An audit of seven major audio corpora finds that queer identities are rarely discussed, and several datasets additionally contain queerphobic content (Agnew et al., 2024*).

[4]Three more corpora without corresponding publications focus on transgender and/or non-binary English voices (German et al., 2022*; Dolquist and Munson, 2023*; Hope, 2023*). Another largely non-queer corpus contains metadata on gender expression and sexual orientation (Weirich et al., 2024*).

Automatic speech recognition (ASR). Only a single project focuses on queer voices in the context of ASR: a recurring shared task on ASR for underrepresented social groups in India including transgender people (Bharathi et al., 2022). However, none of the overview or participant papers address gender, transness or queerness.

Generating queer voices. Three papers build speech synthesis systems for voices perceived as non-binary (Danielescu et al., 2023*; Szekely and Hope, 2024*; Hope and Székely, 2025*) using speech data from non-binary and trans speakers, and asking non-binary people to evaluate the resulting systems. Sigurgeirsson and Ungless (2024)* examine the shortcomings of voice cloning text-to-speech systems in replicating gay voices. They also discuss dual-use issues with queer speech technology, including their use for mockery.

Gaps in queer speech processing. Most publications on queer speech technology are very recent (2022–2025). The special session on *Queer and trans speech science and technology*[5] at Interspeech 2025 also demonstrates a current interest in analyzing and processing queer speech.[6] Given the (non-queer) research on ASR quality differences based on speaker gender (Feng et al., 2024*), one future research direction might be to investigate performance gaps for gender-non-normative voices, including in other tasks with speech input, such as spoken language understanding and speech translation. In speech translation or summarization, queer speakers might be especially prone to being misgendered, which has been acknowledged (Gaido et al., 2020), but not explicitly investigated.

4.7 Position Papers

The last category of surveyed papers is position papers that try to make researchers rethink their approach to queer topics in NLP.

Cisnormative conceptions of gender. An important thread of discussion concerns the treatment of gender in research. Devinney et al. (2022) criticize the field's predominantly binary and cis-normative conceptions of gender. Hovy and Spruit (2016) note that researchers often use pronouns or names as a proxy for gender, and Gautam et al. (2024c) extend this observation to other sociodemographic attributes, describing ethical issues and giving practical advice. Lauscher et al. (2022) study the use of third-person pronouns in English and propose a set of desiderata for pronoun inclusivity in NLP systems. The framing as an ethical issue has already been argued before by Larson (2017), outlining ethical principles about the use of gender, and suggesting that researchers make underlying theory of gender explicit and use gender only if necessary.

Dissect and shape future research. The remaining position papers examine other research trends, dissecting current tendencies to shape future research. Amironesei and Diaz (2023) argue for viewing offensive speech through a relational lens, i.e., "the social relations that inform how language is used and interpreted." McAra-Hunter (2024)* discusses how hype in the field of AI affects queer communities, writing about how harms arise from discriminatory AI technologies, and how hype drives early adoption, obfuscates priorities, drives cisheteronormativity, and neglects human impact.

5 Common Themes

Surveying queer NLP research, we find that several problems have been repeatedly addressed from different angles. The following section looks at these common phenomena, approaches, and gaps.

5.1 Objects of Study

Pronouns. Gender-neutral pronouns and neopronouns feature prominently in our survey. Underrepresentation of specific pronouns (e.g. xier) and pronoun–noun combinations (e.g. "her wife") in training data has been shown to impair performance in coreference resolution and translation tasks. 25 of the queer ALC papers (∼30%) investigate these issues.

Slurs and identity terms. When trained on data stripped of situational context, systems frequently misclassify identity terms as slurs or negative expressions. This leads to problems in hate speech classification, where systems over-classify instances with queer terms as offensive, while not flagging offensive text that lack these words. In sentiment analysis, queer terms are more often associated with negative sentiment. 12 of the queer ALC papers (∼14%) examine these phenomena.

[5] sites.google.com/view/is2025-queer-trans

[6] Outside of NLP, Netzorg et al. (2024)* focused on how speaker identification systems struggle with voice modulation for gender expression and Doukhan et al. (2023)* and McAllister et al. (2025)* on systems that help with voice transition.

Stereotypes and under-representation. Negative stereotypes are also embedded in uncurated training data, leading to harmful outputs in upstream evaluations and downstream applications. However, in more curated data sets, queer data are underrepresented or even filtered out for fear of being offensive, leading to shallow responses from chatbots on queer topics. 13 of the papers surveyed ($\sim$16%) illuminate these problems.

5.2 Methods

Post-hoc bias mitigation. A central motivation in queer NLP is to mitigate anti-queer biases in models. However, this is usually restricted to first revealing the existence of some bias, and only subsequently developing or applying methods that mitigate it. This reactive approach often results in a cat-and-mouse game where papers demonstrate bias in existing models but are unable to develop better models, as debiasing methods are rarely adopted or feasible to operationalize at scale. This pattern has also been observed in other subfields of NLP (Gururaja et al., 2023; Mosbach et al., 2024).

Template-based methods & data augmentation. The prevalence of data augmentation and template-based methods can be seen as a product of the present time, where the affordances of LLMs make these methods expedient. There is acknowledgment that training data are a source of biases, yet there is little focus on data sourcing processes. Templatic methods can reveal egregious biases, and template-based data augmentation can directly address these shortcomings. However, augmenting existing datasets with synthetic language does not represent actual language use, and is not easily portable to non-English languages, further entrenching the Anglocentrism of NLP research.

5.3 Research Gaps

Stakeholder involvement. Although most papers name the impact on the queer community as a motivating factor for their research, very few include queer community members as part of the development or evaluation of their methods. More often, linguistic representations like identity terms, pronouns, and automatically-calculated metrics are used as a stand-in for user feedback. Notable exceptions are Gromann et al. (2023), who conduct a workshop with representatives from the intersex and non-binary community to determine how they would like to be represented in translation, Felkner et al. (2023), who create the WinoQueer Benchmark with text input from the community, and Ungless et al. (2023), who survey trans and non-binary people on their experience with generated images.

Intersectionality. Ovalle et al. (2023b) point out that AI fairness literature mainly engages with intersectionality when applying metrics to demographic subgroups, and avoids other aspects, such as examining power structures. Similarly, Devinney (2025) argues that in order to mitigate the negative effects of stereotyping, NLP researchers need to engage with the social mechanisms behind it. Devinney et al. (2024) addresses these shortcomings by analyzing social biases in LLMs along additive dimensions such as race, gender, transness, and religion, showing that they reflect social power relationships where some groups are seen as default.

Interdisciplinarity. NLP can be thought of as an inherently interdisciplinary field that straddles linguistics and computer science. Especially when dealing with queer language phenomena, previous work in sociolinguistics, gender studies, and queer studies can provide guiding principles for research. One positive example is Amironesei and Diaz (2023), who suggest modeling social relations as a pathway to improved hate speech detection.

Language diversity. The vast majority of papers examine English language phenomena, and even the inherently multilingual field of machine translation considers mostly translations to and from English. Figure 2 shows that the languages considered in the papers we survey are highly skewed towards English and Western European languages, which also leads to a lack of cultural diversity. In Section 6.2 we discuss non-English conferences that show a different path forward.

6 Discussion

Having surveyed common themes in queer NLP, we turn to how we can address the found gaps.

6.1 Representation and Harm in Datasets

Under- and misrepresentation in datasets is a recurring theme in the surveyed papers. However, a neglected point that Zhou (2024) draws attention to is that understandings of queerness vary both across cultures and over time. Given the fluid nature of "queerness" as used by queer people and noted by philosophers (Butler, 2020), models of anti-queer

bias must be equally fluid to account for different dataset contexts. For instance, inoffensive or clinical terms from the past, such as "transvestite," can gain negative associations through pejoration (Finkbeiner et al., 2016).

Representation of queerness in NLP technology would hence require constant rewriting by the affected community, similarly to the slow and constant change of language in different communities of practice, which the affordances of current NLP technologies make hard to imagine. When affected communities are not integrated into the model development process, they cannot contribute data, assess model quality, address shortcomings, or ensure the implementation of debiasing procedures.

Another under-examined problem is "unmarkedness" in language, or how unnatural mentions of privileged identities in templates e.g., "straight man" may yield unreliable results (Blodgett et al., 2021). Here, a lack of engagement with linguistics leads to biased methods. Therefore, we encourage future research to look for alternatives to static models and template-based mitigation strategies.

6.2 Conferences for Non-English Queer NLP

An inherent limitation of our survey is the focus on ACL. Even with its international appeal and prestige, the kind of paper submitted and accepted is still geared towards an English-speaking audience (Lepp and Sarin, 2024), as the target languages in the publications reveal (see Figure 2).

The predominance (if not hegemony) of English decreases when we broaden our scope beyond the Anglosphere. The Brazilian Computer Society has hosted workshops featuring papers on queer NLP about and written in Portuguese (Souza et al., 2022; Santos et al., 2022), and so has the Italian Association of Computational Linguistics, although written in English (Nozza et al., 2023; Draetta et al., 2024). The Spanish Society for Natural Language Processing is particularly prolific, with many queer NLP papers on Spanish and other Iberian languages at their IberLEF workshop (Bel-Enguix et al., 2023; Jiménez-Zafra et al., 2023; Gómez-Adorno et al., 2024; García-Baena et al., 2024).

These workshops address topics that could be explored in queer English NLP papers but are currently absent from them. For example, García-Baena et al. (2023) introduce the task of hope speech detection for queer individuals in Spanish, defined as the detection of "speech that is able to relax a hostile environment and that helps, gives suggestions and inspires for good", which culminated in HOPE, a shared task on hope detection at IberLEF. The first edition in 2023 focused on queer hope speech (Jiménez-Zafra et al., 2023) and the second edition in 2024 still dedicated half of the task to it (García-Baena et al., 2024), while this topic is virtually unexplored in English NLP. Similarly, the study the identification of reclaimed queer slurs in Italian (Draetta et al., 2024) spawned a shared task on multilingual queer slur reclamation for the EVALITA workshop 2026, [7] while the topic remains understudied in English NLP (Zsisku et al., 2024).

Shared tasks play a key role in these events, drawing researchers' attention to the subject. The HOPE tasks from above included 8 papers in 2023 (Jiménez-Zafra et al., 2023) and 16 in 2024 (García-Baena et al., 2024), making it the largest repository of papers on hope detection for the queer community. There are similar numbers for hate detection in Mexican Spanish with the HOMO-MEX task at IberLEF (Bel-Enguix et al., 2023; Gómez-Adorno et al., 2024) and Italian with the HODI task at EVALITA (Nozza et al., 2023). This may indicate that specific challenges help gather momentum for research in non-English queer NLP, which can be especially useful for identities that don't fit the Western mold of "queer" e.g., hijra (see Glossary).

Much like English queer NLP research influences the papers at venues not focused on English, many of their ideas and advances have relevance for languages beyond the one under study. We encourage researchers to broaden their focus accordingly.

6.3 Harms in the Wild vs. in Research

Most queer NLP research has focused on a relatively narrow set of harms. Reliance on rigid taxonomies (e.g., of harm types, speech types, and identities) can flatten the queer experience while disregarding the fluidity of queer language and the necessity for expert stakeholder input (§6.1). Creating these taxonomies is a fundamentally political decision that influences what is visible and invisible within a system (Bowker and Star, 2000; Smith et al., 2018; Guyan, 2025). Most hate speech and toxicity detection systems, for example, categorize (implicitly or explicitly) their targeted speech, but rarely involve the perspectives of queer users who experience hate speech: only 3 out of the 15 papers

[7] https://multipride-evalita.github.io/

surveyed in this category involve queer stakeholders. However, queer slang is often falsely flagged as hateful, due to the use of reclaimed slurs and discussion of sexual themes, (Ramesh et al., 2022; Tint, 2025; Zueva et al., 2020). Meanwhile, so-called "gender-critical" and TERF (see Glossary) rhetoric often avoids such terms, and is typically not detected despite being overtly harmful to queer users (Locatelli et al., 2023; Leto et al., 2025).

Beyond misclassifications, queer users may also reject the premise that systems should categorize them at all. Critical refusal, or opting out of being measured, labeled, or represented in normative ways, is a vital queer methodological position that current NLP systems leave little room for (Abramovich and Ma, 2024). Refusal is not simply a rejection of inclusion, but a rejection of the frameworks that inclusion often demands: that identities be stable, categorizable, and intelligible to institutions (Costanza-Chock, 2020; Crawford, 2021). In practice, this could include actions such as declining to specify a gender on a form, using intentionally ambiguous language online, resisting classification through irony, camp, or coded speech, or, among NLP researchers, refusing to use harmful categorization in research (Gautam et al., 2024c). These modes of resistance challenge NLP systems that depend on fixed typologies and clear labels, exposing a gap between what queer people want from NLP technologies and what those technologies are designed to do. However, refusal is difficult to publish about in NLP, as it definitionally resists detection, annotation, and evaluation. As a result, queer refusal is rendered technically irrelevant even as it remains politically essential.

6.4 Perspectives from Queer Theory

One aim of this survey is to look beyond immediate gaps in queer NLP and to create visions for future work. To understand the complexity of socially-constructed identities and their interactions with sociotechnical systems, we can turn to queer theory approaches. This work examines and challenges power (Foucault, 1978) while confronting the certainty of sex, gender, and sexuality (Butler, 1990). Acknowledging that power plays a role in shaping what it means to be "queer" can help us better conceptualize why queer life is often denigrated or outright erased in NLP systems, providing a path forward for better queer NLP research.

One of queer theory's aims is to question the stability of identity categories, taking particular issue with how these categories are presumed and normalized in everyday life and academic discourse. NLP is currently no less guilty of adhering to strict categorization, particularly when it comes to encoding queerness. In our survey, all papers that specifically study queerness rely on lists of sexual and gender identities, reducing nuanced or fuzzy forms of queer life into simple categorizations. Further, Klipphahn-Karge et al. (2023) argue that AI "posits societal norms within a double bind" by both *reflecting* existing social bias and *producing* heteronormative knowledge that flattens multiplicities. Because of this, some queer theorists claim that data science and its applications, such as NLP, are entirely incompatible with queerness (Keyes, 2019). Others, such as Shah (2023), propose ways to "queer" AI systems by reworking data science building blocks to allow imperfect, non-normative, and community-gathered data points. In either case, in order to break down cis/heteronormative NLP standards, we must expand our scholarship and collaborations from the intersection of computing, statistics, and linguistics to include sociology, gender studies, queer theory, and beyond. With one foot "in the craft work of design" and the other "in the reflexive work of critique" (Agre, 1998), we can build interdisciplinary systems that proactively center those at the margins instead of relying on post-hoc tweaks and retroactive guardrails.

7 Conclusion

This work provides a comprehensive overview of the current state of queer NLP, summarizing findings from a variety of subtopics (§4), as well as contextualizing them (§5–6). Queer NLP echoes many shortcomings of the field as a whole. Despite some recent efforts (§6.2), most work focuses on or around English. Moreover, many works still operate with a binary and cis-normative viewpoint of gender and sexuality, which could be informed better by research from queer theory. Finally, any harm mitigation usually happens post-hoc; this can at least partially be explained by the lack of stakeholder involvement in modern NLP systems.

For these reasons, we advocate for future work to involve affected queer communities (e.g., Bergstrand and Gambäck, 2024) and develop solutions for socially-grounded (rather than assumed) understandings of harm. This could be realized through community-driven data collection, as has

been done for low-resource languages. Participatory action research, which involves stakeholders as research collaborators, can help challenge oppressive epistemologies and methods. Lastly, we also call for the development of NLP systems that not only mitigate harm but actively and meaningfully support queer communities.

References

Shira Abramovich and Anna J. Ma. 2024. Evaluating refusal. Presented at the Workshop on Evaluating Evaluations: Examining Best Practices for Measuring Broader Impacts of Generative AI at NeurIPS 2024, Vancouver, Canada.

William Agnew, Julia Barnett, Annie Chu, Rachel Hong, Michael Feffer, Robin Netzorg, Harry H. Jiang, Ezra Awumey, and Sauvik Das. 2024. Sound check: Auditing audio datasets. *Computing Research Repository*, arXiv:2410.13114.

Philip E. Agre. 1998. Toward a Critical Technical Practice: Lessons Learned in Trying to Reform AI. In Geoffrey Bowker, Geoffrey Bowker, Susan Leigh Star, Les Gasser, and William Turner, editors, *Social Science, Technical Systems, and Cooperative Work*. Psychology Press.

Jesujoba O. Alabi, Michael A. Hedderich, David Ifeoluwa Adelani, and Dietrich Klakow. 2025. Charting the landscape of African NLP: Mapping progress and shaping the road ahead. *Computing Research Repository*, arXiv:2505.21315.

Razvan Amironesei and Mark Diaz. 2023. Relationality and offensive speech: A research agenda. In *The 7th Workshop on Online Abuse and Harms (WOAH)*, pages 85–95, Toronto, Canada. Association for Computational Linguistics.

Chantal Amrhein, Florian Schottmann, Rico Sennrich, and Samuel Läubli. 2023. Exploiting biased models to de-bias text: A gender-fair rewriting model. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4486–4506, Toronto, Canada. Association for Computational Linguistics.

Scott Andersen, Segio-Luis Ojeda-Trueba, Juan Vásquez, and Gemma Bel-Enguix. 2024. The Mexican gayze: A computational analysis of the attitudes towards the LGBT+ population in Mexico on social media across a decade. In *Proceedings of the 8th Workshop on Online Abuse and Harms (WOAH 2024)*, pages 178–200, Mexico City, Mexico. Association for Computational Linguistics.

Elisa Bassignana, Valerio Basile, and Viviana Patti. 2018. Hurtlex: A multilingual lexicon of words to hurt. In *Proceedings of the Fifth Italian Conference on Computational Linguistics (CLiC-it 2018)*, volume 2253 of *CEUR Workshop Proceedings*, Torino, Italy. CEUR-WS.org.

Gemma Bel-Enguix, Helena Gómez-Adorno, Gerardo Sierra, Juan Vásquez, Scott Thomas Andersen, and Sergio Ojeda-Trueba. 2023. Overview of HOMO-MEX at Iberlef 2023: Hate speech detection in online messages directed towards the MEXican Spanish speaking LGBTQ+ population. *Procesamiento del Lenguaje Natural*, 71:361–370.

Selma Bergstrand and Björn Gambäck. 2024. Detecting and mitigating LGBTQIA+ bias in large Norwegian language models. In *Proceedings of the 5th Workshop on Gender Bias in Natural Language Processing (GeBNLP)*, pages 351–364, Bangkok, Thailand. Association for Computational Linguistics.

B Bharathi, Bharathi Raja Chakravarthi, Subalalitha Chinnaudayar Navaneethakrishnan, N Sripriya, Arunaggiri Pandian, and Swetha Valli. 2022. Findings of the shared task on speech recognition for vulnerable individuals in Tamil. In *Proceedings of the Second Workshop on Language Technology for Equality, Diversity and Inclusion*, pages 339–345, Dublin, Ireland. Association for Computational Linguistics.

Su Lin Blodgett, Gilsinia Lopez, Alexandra Olteanu, Robert Sim, and Hanna Wallach. 2021. Stereotyping Norwegian salmon: An inventory of pitfalls in fairness benchmark datasets. 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 1004–1015, Online. Association for Computational Linguistics.

Geoffrey C. Bowker and Susan Leigh Star. 2000. *Sorting things out: Classification and its consequences*. MIT press.

Jacqueline Brixey, Rens Hoegen, Wei Lan, Joshua Rusow, Karan Singla, Xusen Yin, Ron Artstein, and Anton Leuski. 2017. SHIHbot: A Facebook chatbot for sexual health information on HIV/AIDS. In *Proceedings of the 18th Annual*

SIGdial Meeting on Discourse and Dialogue, pages 370–373, Saarbrücken, Germany. Association for Computational Linguistics.

Bastian Bunzeck and Sina Zarrieß. 2024. The SlayQA benchmark of social reasoning: testing gender-inclusive generalization with neopronouns. In *Proceedings of the 2nd GenBench Workshop on Generalisation (Benchmarking) in NLP*, pages 42–53, Miami, Florida, USA. Association for Computational Linguistics.

Judith Butler. 1990. *Gender Trouble: Feminism and the Subversion of Identity*. Thinking Gender. Routledge, New York.

Judith Butler. 2020. Critically queer. In Shane Phelan, editor, *Playing with fire*, pages 11–29. Routledge.

Yang Trista Cao and Hal Daumé III. 2020. Toward gender-inclusive coreference resolution. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4568–4595, Online. Association for Computational Linguistics.

Yang Trista Cao and III Daumé, Hal. 2021. Toward gender-inclusive coreference resolution: An analysis of gender and bias throughout the machine learning lifecycle. *Computational Linguistics*, 47(3):615–661.

Sky CH-Wang and David Jurgens. 2021. Using sociolinguistic variables to reveal changing attitudes towards sexuality and gender. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 9918–9938, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Bharathi Raja Chakravarthi, Prasanna Kumaresan, Ruba Priyadharshini, Paul Buitelaar, Asha Hegde, Hosahalli Shashirekha, Saranya Rajiakodi, Miguel Ángel García, Salud María Jiménez-Zafra, José García-Díaz, Rafael Valencia-García, Kishore Ponnusamy, Poorvi Shetty, and Daniel García-Baena. 2024. Overview of third shared task on homophobia and transphobia detection in social media comments. In *Proceedings of the Fourth Workshop on Language Technology for Equality, Diversity, Inclusion*, pages 124–132, St. Julian's, Malta. Association for Computational Linguistics.

Bharathi Raja Chakravarthi, Rahul Ponnusamy, Malliga S, Paul Buitelaar, Miguel Ángel García-Cumbreras, Salud María Jimenez-Zafra, Jose Antonio Garcia-Diaz, Rafael Valencia-Garcia, and Nitesh Jindal. 2023. Overview of second shared task on homophobia and transphobia detection in social media comments. In *Proceedings of the Third Workshop on Language Technology for Equality, Diversity and Inclusion*, pages 38–46, Varna, Bulgaria. INCOMA Ltd., Shoumen, Bulgaria.

Bharathi Raja Chakravarthi, Ruba Priyadharshini, Thenmozhi Durairaj, John McCrae, Paul Buitelaar, Prasanna Kumaresan, and Rahul Ponnusamy. 2022. Overview of the shared task on homophobia and transphobia detection in social media comments. In *Proceedings of the Second Workshop on Language Technology for Equality, Diversity and Inclusion*, pages 369–377, Dublin, Ireland. Association for Computational Linguistics.

Patricia Hill Collins and Sirma Bilge. 2020. *Intersectionality*. John Wiley & Sons.

Sasha Costanza-Chock. 2020. *Design justice: Community-led practices to build the worlds we need*. The MIT Press.

Kate Crawford. 2021. *The atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press.

Jamell Dacon, Harry Shomer, Shaylynn Crum-Dacon, and Jiliang Tang. 2022. Detecting harmful online conversational content towards LGBTQIA+ individuals. *Computing Research Repository*, arXiv:1503.06733. Presented at the Queer in AI Workshop at NAACL 2022, Seattle, Washington, USA.

Andreea Danielescu, Sharone A Horowit-Hendler, Alexandria Pabst, Kenneth Michael Stewart, Eric M Gallo, and Matthew Peter Aylett. 2023. Creating inclusive voices for the 21st century: A non-binary text-to-speech for conversational assistants. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*, Hamburg, Germany. 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. 2025. Power(ful) associations: Rethinking "stereotype" for NLP. In *Proceedings of the 6th Workshop on Gender Bias in Natural Language Processing (GeBNLP)*, pages 52–58, Vienna, Austria. 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, pages 2083–2102, Seoul, Republic of Korea. Association for Computing Machinery.

Hannah Devinney, Jenny Björklund, and Henrik Björklund. 2024. We don't talk about that: Case studies on intersectional analysis of social bias in large language models. In *Proceedings of the 5th Workshop on Gender Bias in Natural Language Processing (GeBNLP)*, pages 33–44, Bangkok, Thailand. Association for Computational Linguistics.

Harnoor Dhingra, Preetiha Jayashanker, Sayali Moghe, and Emma Strubell. 2023. Queer people are people first: Deconstructing sexual identity stereotypes in large language models. *Computing Research Repository*, arXiv:2307.00101. Presented at the Queer in AI Workshop at ACL 2023, Toronto, Canada.

Devin V. Dolquist and Benjamin Munson. 2023. A palette of transmasculine voices. Data Repository for the University of Minnesota (DRUM).

David Doukhan, Simon Devauchelle, Lucile Girard-Monneron, Mía Chávez Ruz, V. Chaddouk, Isabelle Wagner, and Albert Rilliard. 2023. Voice passing : a non-binary voice gender prediction system for evaluating transgender voice transition. In *Proceedings of Interspeech 2023*, pages 5207–5211, Dublin, Ireland.

Lia Draetta, Chiara Ferrando, Marco Cuccarini, Liam James, and Viviana Patti. 2024. ReCLAIM project: Exploring Italian slurs reappropriation with large language models. In *Proceedings of the Tenth Italian Conference on Computational Linguistics (CLiC-it 2024)*, volume 3878 of *CEUR Workshop Proceedings*, pages 335–342, Pisa, Italy. CEUR-WS.org.

Justin Edwards, Leigh Clark, and Allison Perrone. 2021. LGBTQ-AI? Exploring expressions of gender and sexual orientation in chatbots. In *Proceedings of the 3rd Conference on Conversational User Interfaces*, CUI '21, Bilbao (online), Spain. Association for Computing Machinery.

Paul Engelmann, Peter Trolle, and Christian Hardmeier. 2024. A dataset for the detection of dehumanizing language. In *Proceedings of the Fourth Workshop on Language Technology for Equality, Diversity, Inclusion*, pages 14–20, St. Julian's, Malta. Association for Computational Linguistics.

Virginia Felkner, Ho-Chun Herbert Chang, Eugene Jang, and Jonathan May. 2023. WinoQueer: A community-in-the-loop benchmark for anti-LGBTQ+ bias in large language models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9126–9140, Toronto, Canada. Association for Computational Linguistics.

Siyuan Feng, Bence Mark Halpern, Olya Kudina, and Odette Scharenborg. 2024. Towards inclusive automatic speech recognition. *Computer Speech & Language*, 84:101567.

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.

Rita Finkbeiner, Jörg Meibauer, and Heike Wiese. 2016. What is pejoration, and how can it be expressed in language? In Rita Finkbeiner, Jörg Meibauer, and Heike Wiese, editors, *Pejoration*, pages 1–18. John Benjamins Publishing Company.

Michel Foucault. 1978. *The History of Sexuality, Volume 1: An Introduction*. Pantheon Books, New York.

Marco Gaido, Beatrice Savoldi, Luisa Bentivogli, Matteo Negri, and Marco Turchi. 2020. Breeding gender-aware direct speech translation systems. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 3951–3964, Barcelona, Spain (Online). International Committee on Computational Linguistics.

Daniel García-Baena, Miguel Ángel García-Cumbreras, Salud María Jiménez-Zafra, José Antonio García-Díaz, and Rafael Valencia-García. 2023. Hope speech detection in Spanish: The LGBT case. *Language Resources and Evaluation*, 57:1487–1514.

Daniel García-Baena, Fazlourrahman Balouchzahi, Sabur Butt, Miguel Ángel García-Cumbreras, Atnafu Lambebo Tonja, José Antonio García-Díaz, Selen Bozkurt, Bharathi Raja Chakravarthi, Hector G. Ceballos, Rafael Valencia-García, Grigori Sidorov, L. Alfonso Ureña-López, Alexander Gelbukh, and Salud María Jiménez-Zafra. 2024. Overview of HOPE at IberLEF 2024: Approaching hope speech detection in social media from two perspectives, for equality, diversity and inclusion and as expectations. *Procesamiento del Lenguaje Natural*, 73:407–419.

Vagrant Gautam, Eileen Bingert, Dawei Zhu, Anne Lauscher, and Dietrich Klakow. 2024a. Robust pronoun fidelity with English LLMs: Are they reasoning, repeating, or just biased? *Transactions of the Association for Computational Linguistics*, 12:1755–1779.

Vagrant Gautam, Julius Steuer, Eileen Bingert, Ray Johns, Anne Lauscher, and Dietrich Klakow. 2024b. WinoPron: Revisiting English Winogender schemas for consistency, coverage, and grammatical case. In *Proceedings of the Seventh Workshop on Computational Models of Reference, Anaphora and Coreference*, pages 52–66, Miami, Florida, USA. Association for Computational Linguistics.

Vagrant Gautam, Arjun Subramonian, Anne Lauscher, and Os Keyes. 2024c. Stop! In the name of flaws: Disentangling personal names and sociodemographic attributes in NLP. In *Proceedings of the 5th Workshop on Gender Bias in Natural Language Processing (GeBNLP)*, pages 323–337, Bangkok, Thailand. Association for Computational Linguistics.

James Sneed German, Maria Candea, LeAnn Brown, Timothy Mahrt, and Oriana Reid-Collins. 2022. Gender spectrum speech corpus. ORTOLANG (Open Resources and TOols for LANGuage).

Sourojit Ghosh and Aylin Caliskan. 2023. ChatGPT perpetuates gender bias in machine translation and ignores non-gendered pronouns: Findings across Bengali and five other low-resource languages. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '23, pages 901–912. ACM.

Ina Goel. 2016. Hijra communities of Delhi. *Sexualities*, 19(5-6):535–546.

Seraphina Goldfarb-Tarrant, Eddie Ungless, Esma Balkir, and Su Lin Blodgett. 2023. This prompt is measuring <mask>: evaluating bias evaluation in language models. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 2209–2225, Toronto, Canada. Association for Computational Linguistics.

Dagmar Gromann, Manuel Lardelli, Katta Spiel, Sabrina Burtscher, Lukas Daniel Klausner, Arthur Mettinger, Igor Miladinovic, Sigrid Schefer-Wenzl, Daniela Duh, and Katharina Bühn. 2023. Participatory research as a path to community-informed, gender-fair machine translation. In *Proceedings of the First Workshop on Gender-Inclusive Translation Technologies*, pages 49–59, Tampere, Finland. European Association for Machine Translation.

Sireesh Gururaja, Amanda Bertsch, Clara Na, David Widder, and Emma Strubell. 2023. To build our future, we must know our past: Contextualizing paradigm shifts in natural language processing. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 13310–13325, Singapore. Association for Computational Linguistics.

Kevin Guyan. 2025. *Rainbow Trap: Queer Lives, Classifications and the Dangers of Inclusion*. Bloomsbury Academic Publishing.

Helena Gómez-Adorno, Gemma Bel-Enguix, Hiram Calvo, Sergio Ojeda-Trueba, Scott Thomas Andersen, Juan Vásquez, Tania Alcántara, Miguel Soto, and Cesar Macias. 2024. Overview of HOMO-MEX at IberLEF 2024: Hate speech detection towards the Mexican Spanish speaking LGBT+ population. *Procesamiento del Lenguaje Natural*, 73:393–405.

Donna Haraway. 1988. Situated knowledges: The science question in feminism and the privilege of partial perspective. *Feminist Studies*, 14(3):575–599.

Amanda Hicks, Michael Rutherford, Christiane Fellbaum, and Jiang Bian. 2016. An analysis of WordNet's coverage of gender identity using Twitter and the national transgender discrimination survey. In *Proceedings of the 8th Global WordNet Conference (GWC)*, pages 123–130, Bucharest, Romania. Global Wordnet Association.

Jacob Hobbs. 2025. Theories of "sexuality" in natural language processing bias research. *Computing Research Repository*, arXiv:2506.22481.

Maxwell Hope. 2023. Mid-Atlantic gender expansive speech (MAGES) corpus. Online.

Maxwell Hope and Éva Székely. 2025. Voices of 'cyborg awesomeness': Posthuman embodiment of nonbinary gender expression in AI speech technologies. In *Proceedings of Interspeech 2025*, pages 689–693, Rotterdam, The Netherlands.

Adnan Hossain. 2017. The paradox of recognition: hijra, third gender and sexual rights in Bangladesh. *Culture, Health & Sexuality*, 19(12):1418–1431.

Tamanna Hossain, Sunipa Dev, and Sameer Singh. 2023. MISGENDERED: Limits of large language models in understanding pronouns. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5352–5367, Toronto, Canada. Association for Computational Linguistics.

Dirk Hovy and Shannon L. Spruit. 2016. The social impact of natural language processing. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 591–598, Berlin, Germany. Association for Computational Linguistics.

Human Rights Campaign. 2025. Glossary of terms. Retrivied October 28th, 2025.

Salud María Jiménez-Zafra, Miguel Ángel Garcia-Cumbreras, Daniel García-Baena, José Antonio Garcia-Díaz, Bharathi Raja Chakravarthi, Rafael Valencia-García, and Luis Alfonso Ureña-López. 2023. Overview of HOPE at IberLEF 2023: Multilingual hope speech detection. *Procesamiento del Lenguaje Natural*, 71:371–381.

Rosemary A. Joyce. 2015. Heteronormativity. In *The International Encyclopedia of Human Sexuality*, pages 501–581. John Wiley & Sons.

Os Keyes. 2019. Counting the countless. Published in the online magazine *Real Life*. Retrivied October 28th, 2025.

Svetlana Kiritchenko and Saif Mohammad. 2018. Examining gender and race bias in two hundred sentiment analysis systems. In *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics*, pages 43–53, New Orleans, Louisiana. Association for Computational Linguistics.

Michael Klipphahn-Karge, Ann-Kathrin Koster, and Sara Morais dos Santos Bruss. 2023. Introduction: Queer AI. In *Queer Reflections on AI*, pages 1–19. Routledge.

Mitchell R. Koss, Matthew McLaughlin, Kayla Switalla, Israel Falade, and Esther Kim. 2025. Exploring the role of ChatGPT in decision making for gender-affirming surgery. *Artificial Intelligence Surgery*, 5(3):116–125.

Prasanna Kumar Kumaresan, Rahul Ponnusamy, Dhruv Sharma, Paul Buitelaar, and Bharathi Raja Chakravarthi. 2024. Dataset for identification of homophobia and transphobia for Telugu, Kannada, and Gujarati. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 4404–4411, Torino, Italia. European Language Resources Association and International Committee on Computational Linguistics.

Manuel Lardelli and Dagmar Gromann. 2023. Gender-fair post-editing: A case study beyond the binary. In *Proceedings of the 24th Annual Conference of the European Association for Machine Translation*, pages 251–260, Tampere, Finland. European Association for Machine Translation.

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.

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.

Anne Lauscher, Debora Nozza, Ehm Miltersen, Archie Crowley, and Dirk Hovy. 2023. What about "em"? How commercial machine translation fails to handle (neo-)pronouns. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 377–392, Toronto, Canada. Association for Computational Linguistics.

Haley Lepp and Parth Sarin. 2024. A global AI community requires language-diverse publishing. *Computing Research Repository*, arXiv:2408.14772. Presented at the Global AI Cultures workshop at ICLR 2024, Vienna, Austria.

Alexandria Leto, Juan Vásquez, Alexis Palmer, and Maria Leonor Pacheco. 2025. Dehumanization of LGBTQ+ groups in sexual interactions with ChatGPT. In *Proceedings of the Queer in AI Workshop at NAACL 2025*, pages 17–25, Albuquerque, New Mexico, USA (Hybrid). Association for Computational Linguistics.

Anton Leuski and David Traum. 2010. NPCEditor: A tool for building question-answering characters. In *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10)*, Valletta, Malta. European Language Resources Association.

Erez Levon. 2007. Sexuality in context: Variation and the sociolinguistic perception of identity. *Language in Society*, 36(4):533–554.

Shir Lissak, Nitay Calderon, Geva Shenkman, Yaakov Ophir, Eyal Fruchter, Anat Brunstein Klomek, and Roi Reichart. 2024. The colorful future of LLMs: Evaluating and improving LLMs as emotional supporters for queer youth. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 2040–2079, Mexico City, Mexico. Association for Computational Linguistics.

Nikola Ljubešić, Ilia Markov, Darja Fišer, and Walter Daelemans. 2020. The LiLaH emotion lexicon of Croatian, Dutch and Slovene. In *Proceedings of the Third Workshop on Computational Modeling of People's Opinions, Personality, and Emotion's in Social Media*, pages 153–157, Barcelona, Spain (Online). Association for Computational Linguistics.

Davide Locatelli, Greta Damo, and Debora Nozza. 2023. A cross-lingual study of homotransphobia on Twitter. In *Proceedings of the First Workshop on Cross-Cultural Considerations in NLP (C3NLP)*, pages 16–24, Dubrovnik, Croatia. Association for Computational Linguistics.

Christina Lu and David Jurgens. 2022. The subtle language of exclusion: Identifying the toxic speech of trans-exclusionary radical feminists. In *Proceedings of the Sixth Workshop on Online Abuse and Harms (WOAH)*, pages 79–91, Seattle, Washington (Hybrid). Association for Computational Linguistics.

Scott M. Lundberg and Su-In Lee. 2017. A unified approach to interpreting model predictions. In *Proceedings of the 31st International Conference on Neural Information Processing Systems (NIPS'17)*, pages 4768–4777, Long Beach, California, USA. Curran Associates Inc.

Zilin Ma, Yiyang Mei, Yinru Long, Zhaoyuan Su, and Krzysztof Z. Gajos. 2024. Evaluating the experience of LGBTQ+ people using large language model based chatbots for mental health support. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*

(CHI '24), Honolulu, Hawaii, USA. Association for Computing Machinery.

Tara McAllister, Collin Eagen, Yi Shan, Peter Traver, Daphna Harel, Tae Hong Park, and Vesna Novak. 2025. Web-based application for real-time biofeedback of vocal resonance in gender-affirming voice training: Design and usability evaluation. In *Proceedings of Interspeech 2025*, pages 664–668, Rotterdam, The Netherlands.

Dawn McAra-Hunter. 2024. How AI hype impacts the LGBTQ + community. *AI and Ethics*, 4(3):771–790.

Jennifer Mickel, Maria De-Arteaga, Leqi Liu, and Kevin Tian. 2025. More of the same: Persistent representational harms under increased representation. *Computing Research Repository*, arXiv:2503.00333.

David Moher, Larissa Shamseer, Mike Clarke, Davina Ghersi, Alessandro Liberati, Mark Petticrew, Paul Shekelle, Lesley A. Stewart, and PRISMA-P. Group. 2015. Preferred reporting items for systematic review and meta-analysis protocols (PRISMA-P) 2015 statement. *Systematic Reviews*, 4(1).

Marius Mosbach, Vagrant Gautam, Tomás Vergara Browne, Dietrich Klakow, and Mor Geva. 2024. From insights to actions: The impact of interpretability and analysis research on NLP. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 3078–3105, Miami, Florida, USA. Association for Computational Linguistics.

Daniel Najafali, Chandler Hinson, Justin M. Camacho, Logan G. Galbraith, Tannon L. Tople, Danielle Eble, Brielle Weinstein, Loren S. Schechter, Amir H. Dorafshar, and Shane D. Morrison. 2023. Artificial intelligence knowledge of evidence-based recommendations in gender affirmation surgery and gender identity: is ChatGPT aware of WPATH recommendations? *European Journal of Plastic Surgery*, 46(6):1169–1176.

Robin Netzorg, Alyssa Cote, Sumi Koshin, Klo Vivienne Garoute, and Gopala Krishna Anumanchipalli. 2024. Speech after gender: A trans-feminine perspective on next steps for speech science and technology. In *Proceedings of Interspeech 2024*, pages 3075–3079, Kos, Greece.

Hellina Hailu Nigatu, Atnafu Lambebo Tonja, Benjamin Rosman, Thamar Solorio, and Monojit Choudhury. 2024. The Zeno's paradox of 'low-resource' languages. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 17753–17774, Miami, Florida, USA. Association for Computational Linguistics.

Nonbinary Wiki. 2025a. Cissexism. Retrivied October 28th, 2025.

Nonbinary Wiki. 2025b. LGBT. Retrivied October 28th, 2025.

Debora Nozza, Federico Bianchi, and Dirk Hovy. 2021. HONEST: Measuring hurtful sentence completion in language models. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2398–2406, Online. Association for Computational Linguistics.

Debora Nozza, Federico Bianchi, Anne Lauscher, and Dirk Hovy. 2022. Measuring harmful sentence completion in language models for LGBTQIA+ individuals. In *Proceedings of the Second Workshop on Language Technology for Equality, Diversity and Inclusion*, pages 26–34, Dublin, Ireland. Association for Computational Linguistics.

Debora Nozza, Alessandra Teresa Cignarella, Greta Damo, Tommaso Caselli, and Viviana Patti. 2023. HODI at EVALITA 2023: Overview of the first shared task on homotransphobia detection in Italian. In *Proceedings of the Eighth Evaluation Campaign of Natural Language Processing and Speech Tools for Italian. Final Workshop (EVALITA 2023)*, volume 3473 of *CEUR Workshop Proceedings*. CEUR-WS.org.

Anaelia Ovalle, Palash Goyal, Jwala Dhamala, Zachary Jaggers, Kai-Wei Chang, Aram Galstyan, Richard Zemel, and Rahul Gupta. 2023a. "I'm fully who I am": Towards centering transgender and non-binary voices to measure biases in open language generation. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '23, pages 1246–1266, Chicago, Illinois, USA. Association for Computing Machinery.

Anaelia Ovalle, Krunoslav Lehman Pavasovic, Louis Martin, Luke Zettlemoyer, Eric Michael Smith, Kai-Wei Chang, Adina Williams, and Levent Sagun. 2024. The root shapes the fruit: On the persistence of gender-exclusive harms in aligned language models. *Computing Research Repository*, arXiv:2411.03700.

Anaelia Ovalle, Arjun Subramonian, Vagrant Gautam, Gilbert Gee, and Kai-Wei Chang. 2023b. 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, pages 496–511, Montréal, Canada. Association for Computing Machinery.

Angela Balducci Paolucci, Manuel Lardelli, and Dagmar Gromann. 2023. Gender-fair language in translation: A case study. In *Proceedings of the First Workshop on Gender-Inclusive Translation Technologies*, pages 13–23, Tampere, Finland. European Association for Machine Translation.

Andrea Piergentili, Beatrice Savoldi, Matteo Negri, and Luisa Bentivogli. 2024. Enhancing gender-inclusive machine translation with neomorphemes and large language models. In *Proceedings of the 25th Annual Conference of the European Association for Machine Translation (Volume 1)*, pages 300–314, Sheffield, UK. European Association for Machine Translation (EAMT).

Andrea Piergentili, Beatrice Savoldi, Matteo Negri, and Luisa Bentivogli. 2025. An LLM-as-a-judge approach for scalable gender-neutral translation evaluation. In *Proceedings of the 3rd Workshop on Gender-Inclusive Translation Technologies (GITT 2025)*, pages 46–63, Geneva, Switzerland. European Association for Machine Translation.

A Pranav, Alissa Valentine, Shaily Bhatt, Yanan Long, Arjun Subramonian, Amanda Bertsch, Anne Lauscher, and Ankush Gupta, editors. 2025. *Proceedings of the Queer in AI Workshop*. Association for Computational Linguistics, Hybrid format (in-person and virtual).

Organizers of Queer in AI, Anaelia Ovalle, Arjun Subramonian, Ashwin Singh, Claas Voelcker, Danica J. Sutherland, Davide Locatelli, Eva Breznik, Filip Klubicka, Hang Yuan, Hetvi J, Huan Zhang, Jaidev Shriram, Kruno Lehman, Luca Soldaini, Maarten Sap, Marc Peter Deisenroth, Maria Leonor Pacheco, Maria Ryskina, and 31 others. 2023. Queer in AI: A case study in community-led participatory ai. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency (FAccT '23)*, pages 1882–1895, Chicago, Illinois, USA. Association for Computing Machinery.

Krithika Ramesh, Sumeet Kumar, and Ashiqur Khudabukhsh. 2022. Revisiting queer minorities in lexicons. In *Proceedings of the Sixth Workshop on Online Abuse and Harms (WOAH)*, pages 245–251, Seattle, Washington (Hybrid). Association for Computational Linguistics.

Spencer Rarrick, Ranjita Naik, Varun Mathur, Sundar Poudel, and Vishal Chowdhary. 2023. GATE: A challenge set for gender-ambiguous translation examples. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '23, pages 845–854, Montréal, Canada. Association for Computing Machinery.

Kevin Robinson, Sneha Kudugunta, Romina Stella, Sunipa Dev, and Jasmijn Bastings. 2024. MiTTenS: A dataset for evaluating gender mistranslation. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 4115–4124, Miami, Florida, USA. Association for Computational Linguistics.

Nihar Sahoo, Himanshu Gupta, and Pushpak Bhattacharyya. 2022. Detecting unintended social bias in toxic language datasets. In *Proceedings of the 26th Conference on Computational Natural Language Learning (CoNLL)*, pages 132–143, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.

Ariadna Sanchez, Alice Ross, and Nina Markl. 2024. Beyond the binary: Limitations and possibilities of gender-related speech technology research. In *2024 IEEE Spoken Language Technology Workshop (SLT)*, pages 526–532.

Vinícius Santos, Felipe Henriques, and Gustavo Guedes. 2022. O discurso de ódio homofóbico no Twitter a partir da análise de dados. In *Anais do XI Brazilian Workshop on Social Network Analysis and Mining (BRASNAM)*, pages 109–120, Porto Alegre, RS, Brasil. Sociedade Brasileira de Computação (SBC).

Allison Scott, Freada Kapor Klein, and Uriridiakoghene Onovakpuri. 2017. TECH LEAVERS STUDY. Technical report, Kapor Center for Social Impact.

Nishant Shah. 2023. I spy, with my little AI: How queer bodies are made dirty for digital technologies to claim cleanness. In *Queer Reflections on AI*, pages 57–62. Routledge.

Sagi Shaier, Kevin Bennett, Lawrence Hunter, and Katharina Kann. 2023. Emerging challenges in personalized medicine: Assessing demographic effects on biomedical question answering systems. In *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 540–550, Nusa Dua, Bali. Association for Computational Linguistics.

Larissa Shamseer, David Moher, Mike Clarke, Davina Ghersi, Alessandro Liberati, Mark Petticrew, Paul Shekelle, Lesley A Stewart, and PRISMA-P Group. 2015. Preferred reporting items for systematic review and meta-analysis protocols (PRISMA-P) 2015: elaboration and explanation. *BMJ*, 350:g7647.

Deepawali Sharma, Vedika Gupta, and Vivek Kumar Singh. 2023. Detection of homophobia & transphobia in Malayalam and Tamil: Exploring deep learning methods. In *Advanced Network Technologies and Intelligent Computing: Part II*, volume 1798 of *Communications in Computer and Information Science (CCIS)*, pages 217–226, Varanasi, India. Springer Nature Switzerland.

Emily Sheng, Kai-Wei Chang, Premkumar Natarajan, and Nanyun Peng. 2019. The woman worked as a babysitter: On biases in language generation. 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 3407–3412, Hong Kong, China. Association for Computational Linguistics.

Ingo Siegert, Jan Marquenie, and Sven Grawunder. 2025. Queer Waves: a German speech dataset capturing gender and sexual diversity from podcasts and YouTube. In *Proceedings of Interspeech 2025*, pages 679–683, Rotterdam, The Netherlands.

Atli Sigurgeirsson and Eddie L. Ungless. 2024. Just because we camp, doesn't mean we should: The ethics of modelling queer voices. In *Proceedings of Interspeech 2024*, pages 3050–3054, Kos, Greece.

Adrian P. Simpson and Melanie Weirich. 2020. Phonetic correlates of sex, gender and sexual orientation. In *Oxford Research Encyclopedia of Linguistics*. Oxford University Press.

Warren Smith, Matthew Higgins, George Kokkinidis, and Martin Parker. 2018. Becoming invisible: The ethics and politics of imperceptibility. *Culture and Organization*, 24(1):54–73.

Mae Sosto and Alberto Barrón-Cedeño. 2024. Queerbench: Quantifying discrimination in language models toward queer identities. *Computing Research Repository*, arXiv:2406.12399.

Andrey Souza, Eduardo Nakamura, and Fabíola Nakamura. 2022. Detecção de discurso de ódio: Homofobia. In *Anais do XVI Brazilian e-Science Workshop*, pages 73–80, Porto Alegre, RS, Brasil. Sociedade Brasileira de Computação (SBC).

Karolina Stanczak and Isabelle Augenstein. 2021. A survey on gender bias in natural language processing. *Computing Research Repository*, arXiv:2112.14168.

Ian Stewart and Rada Mihalcea. 2024. Whose wife is it anyway? Assessing bias against same-gender relationships in machine translation. In *Proceedings of the 5th Workshop on Gender Bias in Natural Language Processing (GeBNLP)*, pages 365–375, Bangkok, Thailand. Association for Computational Linguistics.

Arjun Subramonian, Vagrant Gautam, Dietrich Klakow, and Zeerak Talat. 2024. Understanding "democratization" in NLP and ML research. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 3151–3166, Miami, Florida, USA. Association for Computational Linguistics.

Eva Szekely and Maxwell Hope. 2024. An inclusive approach to creating a palette of synthetic voices for gender diversity. In *Proceedings of Interspeech 2024*, pages 3070–3074, Kos, Greece.

Kunsheng Tang, Wenbo Zhou, Jie Zhang, Aishan Liu, Gelei Deng, Shuai Li, Peigui Qi, Weiming Zhang, Tianwei Zhang, and NengHai Yu. 2024. GenderCARE: A comprehensive framework for assessing and reducing gender bias in large language models. In *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security*, CCS '24, pages 1196–1210, Salt Lake City, Utah, USA. Association for Computing Machinery.

Jordan Taylor, Ellen Simpson, Anh-Ton Tran, Jed R. Brubaker, Sarah E Fox, and Haiyi Zhu. 2024. Cruising queer HCI on the DL: A literature review of LGBTQ+ people in HCI. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*, Honolulu, Hawaii, USA. Association for Computing Machinery.

Joshua Tint. 2025. Guardrails, not guidance: Understanding responses to LGBTQ+ language in large language models. In *Proceedings of the Queer in AI Workshop at NAACL 2025*, pages 6–16, Albuquerque, New Mexico, USA (Hybrid). Association for Computational Linguistics.

Eddie Ungless, Bjorn Ross, and Anne Lauscher. 2023. Stereotypes and smut: The (mis)representation of non-cisgender identities by text-to-image models. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 7919–7942, Toronto, Canada. Association for Computational Linguistics.

Juan Vásquez, Scott Andersen, Gemma Bel-enguix, Helena Gómez-adorno, and Sergio-luis Ojeda-trueba. 2023. HOMO-MEX: A Mexican Spanish annotated corpus for LGBT+phobia detection on Twitter. In *The 7th Workshop on Online Abuse and Harms (WOAH)*, pages 202–214, Toronto, Canada. Association for Computational Linguistics.

Leonor Veloso, Luisa Coheur, and Rui Ribeiro. 2023. A rewriting approach for gender inclusivity in Portuguese. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 8747–8759, Singapore. Association for Computational Linguistics.

Andreas Waldis, Joel Birrer, Anne Lauscher, and Iryna Gurevych. 2024. The Lou dataset - exploring the impact of gender-fair language in German text classification. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 10604–10624, Miami, Florida, USA. Association for Computational Linguistics.

Angelina Wang, Vikram V Ramaswamy, and Olga Russakovsky. 2022. Towards intersectionality in machine learning: Including more identities, handling underrepresentation, and performing evaluation. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '22, pages 336–349, Seoul, Republic of Korea. Association for Computing Machinery.

Melanie Weirich, Daniel Duran, and Stefanie Jannedy. 2024. Gender and age based f0-variation in the German Plapper corpus. In *Proceedings of Interspeech 2024*, pages 1565–1569, Kos, Greece.

Michael Wiegand and Josef Ruppenhofer. 2024. Oddballs and misfits: Detecting implicit abuse in which identity groups are depicted as deviating from the norm. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 2200–2218, Miami, Florida, USA. Association for Computational Linguistics.

Hsiao-Han Wu and Shu-Kai Hsieh. 2017. Exploring lavender tongue from social media texts[in Chinese]. In *Proceedings of the 29th Conference on Computational Linguistics and Speech Processing (ROCLING 2017)*, pages 68–80, Taipei, Taiwan. The Association for Computational Linguistics and Chinese Language Processing.

Sean Xie, Saeed Hassanpour, and Soroush Vosoughi. 2024. Addressing healthcare-related racial and LGBTQ+ biases in pretrained language models. In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 4451–4464, Mexico City, Mexico. Association for Computational Linguistics.

Guanhua Zhang, Bing Bai, Junqi Zhang, Kun Bai, Conghui Zhu, and Tiejun Zhao. 2020. Demographics should not be the reason of toxicity: Mitigating discrimination in text classifications with instance weighting. In *Proceedings of the 58th Annual Meeting of the Association for Com-*

putational Linguistics, pages 4134–4145, Online. Association for Computational Linguistics.

Amy (Azure) Zhou. 2024. Queer bias in natural language processing: Towards more expansive frameworks of gender and sexuality in NLP bias research. *GRACE: Global Review of AI Community Ethics*, 2(1).

Lal Zimman. 2018. Transgender voices: Insights on identity, embodiment, and the gender of the voice. *Language and Linguistics Compass*, 12(8):e12284.

Lal Zimman. 2021. Gender diversity and the voice. In Jo Angouri and Judith Baxter, editors, *The Routledge Handbook of Language, Gender, and Sexuality*, pages 69–90. Routledge.

Eszter Zsisku, Arkaitz Zubiaga, and Haim Dubossarsky. 2024. Hate speech detection and reclaimed language: Mitigating false positives and compounded discrimination. In *Proceedings of the 16th ACM Web Science Conference*, WEBSCI '24, pages 241–249, Stuttgart, Germany. Association for Computing Machinery.

Nadezhda Zueva, Madina Kabirova, and Pavel Kalaidin. 2020. Reducing unintended identity bias in Russian hate speech detection. In *Proceedings of the Fourth Workshop on Online Abuse and Harms*, pages 65–69, Online. Association for Computational Linguistics.

A Glossary

Here, we define some terms that we use throughout the paper:

- **cisnormativity**: an ideology that promotes cisgender individuals as valid while painting transgender individuals as "unnatural or dangerous" (Nonbinary Wiki, 2025a)
- **heteronormativity**: an ideology that promotes only heterosexual individuals as natural and normal (Joyce, 2015)
- **hijra**: a third gender of "male-bodied feminine-identified people" (Hossain, 2017) rooted in the social, cultural and religious framework of India and several other South Asian countries (Goel, 2016)
- **LGBTQIA+**: a collective term for gender, sex and sexuality minorities. It is an acronym for **L**esbians, **G**ay, **B**isexual, **T**ransgender, **Q**ueer/**Q**uestioning, **I**ntersexual and **A**romatic/**A**sexual, with the plus sign (**+**) including any other relevant minority identity not hereby covered (Human Rights Campaign, 2025; Nonbinary Wiki, 2025b)
- **non-binary**: a category of individuals who do not fully identify with a binary gender i.e., individuals who are neither man or woman
- **trans-exclusionary radical feminists (TERFS)**: individuals who advocate for "feminism" with the view that transgender women are infringing on the rights of cisgender women

B Language Breakdown

This is a breakdown of all the languages covered in the examined papers:

Language	#	Singletons (1 each)	
Chinese	2	Arabic	Polish
Danish	3	Croatian	Romanian
Dutch	3	Catalan	Slovene
English	66	Farsi	Tagalog
French	6	Finnish	Tamil
German	15	Georgian	Telugu
Greek	2	Gujarati	Ukrainian
Italian	10	Hindi	Vietnamese
Japanese	2	Icelandic	
Norwegian	2	Indonesian	
Portuguese	3	Kannada	
Russian	2	Korean	
Spanish	9	Lithuanian	
Swedish	5	Persian	

Table 1: Counts of appearances of languages across papers. Languages with only one paper are grouped under "Singletons."

Limitations

In our efforts to survey the field of NLP for works regarding queer topics, we acknowledge two main limitations: Firstly, while we include papers from other venues such as Interspeech, FAccT, CHI, etc. in Section 4 and non-English conferences in Section 6.2, our work is still centered on research published in *ACL conferences. Secondly, our initial paper search utilized a list of queer identity keywords. Both factors might lead to a selection that misses publications from other venues, other non-English venues and queer aspects that are not captured by our chosen keywords.

Ethics Statement

Positionality The authors of this paper hold many different standpoints afforded to us by our diverse identities, cultures, geographic locations, and personal histories. Most relevantly for this work, we draw from our positions both as experts in various NLP fields, and as members and allies of queer communities. Our perspective on the issues we identify and discuss does not aspire to be an objective "view from nowhere" (Haraway, 1988), but instead is rooted in our local knowledge and understandings. In particular, this paper is largely shaped by North American and Eurocentric perspectives. We draw together partial perspectives to synthesize a collective sense of the current state of queer NLP, its implications and impacts, and possible futures.