

**UCLA**

**UCLA Electronic Theses and Dissertations**

**Title**

Learning from the Outliers: On Centering Underrepresented Communities to Build Inclusive and Socially-Grounded Language Technologies

**Permalink**

<https://escholarship.org/uc/item/6zq745qx>

**Author**

OVALLE, ANAELIA

**Publication Date**

2024

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA

Los Angeles

Learning from the Outliers: On Centering Underrepresented Communities to Build Inclusive and  
Socially-Grounded Language Technologies

A dissertation submitted in partial satisfaction  
of the requirements for the degree  
Doctor of Philosophy in Computer Science

by

Anaelia Altagracia Ovalle

2024

© Copyright by  
Anaelia Altagracia Ovalle  
2024

## ABSTRACT OF THE DISSERTATION

Learning from the Outliers: On Centering Underrepresented Communities to Build Inclusive and Socially-Grounded Language Technologies

by

Anaelia Altagracia Ovalle

Doctor of Philosophy in Computer Science

University of California, Los Angeles, 2024

Professor Kai-Wei Chang, Chair

Large scale deployment of chat-based large language models (LLM) require careful evaluations to ensure these systems operate in an inclusive manner across diverse sociocultural contexts. Prior research has found that AI-driven systems can replicate and amplify existing social inequalities, such as ascribing a person who uses the pronoun *she* as less likely to be a doctor and more likely to be a homemaker [BCZ16]. Historically marginalized communities, such as transgender and non-binary (TGNB) individuals, are particularly susceptible to these harms, as algorithmic systems often fail to represent identities that diverge from binary gender conventions.

This dissertation demonstrates the interdependence of technical and social considerations in the development of inclusive language models. In the first part, we systematically investigate the representational harms LLMs can inflict on TGNB identities. We introduce TANGO, a benchmark dataset designed to evaluate gender-inclusive competencies such as pronoun congruence and gender disclosure. Our findings reveal high misgendering rates and severe data-resource limitations, leading to poor handling of gender-diverse pronouns. To address these challenges, we propose novel mitigation techniques which center tokenization and low-resource methods, leading to sig-

nificant improvements in LLM gender inclusivity.

In the second part, we uncover fundamental limitations within existing gender bias evaluation frameworks, highlighting the sociotechnical consequences of limited construct validity. Through contextually grounded evaluations based on lived TGNB experiences, we demonstrate that even LLMs explicitly aligned for safety can propagate harmful biases that go undetected by conventional evaluation frameworks. By involving the TGNB community in dataset creation and evaluation, we showcase how participatory methods can ensure that marginalized voices guide the development of more inclusive AI systems. Finally, we present SLOGAN, a framework for detecting local biases in clinical prediction tasks, illustrating how these contextually grounded techniques can address biases in various domains. Together, these findings collectively highlight promising directions for tackling LLM harms through community-informed technical and systemic mitigation strategies.

The dissertation of Anaelia Altagracia Ovalle is approved.

Wei Wang

Keith C. Norris

Yizhou Sun

Kai-Wei Chang, Committee Chair

University of California, Los Angeles

2024

*For my family y él que anda guayando yuca.*

*“I lack imagination you say*

*No. I lack language.*

*The language to clarify  
my resistance to the literate.*

*Words are a war to me.*

*They threaten my family.*

*To gain the word*

*to describe the loss*

*I risk losing everything.*

*I may create a monster*

*the word’s length and body*

*swelling up colorful and thrilling*

*looming over my mother, characterized.*

*Her voice in the distance*

*unintelligible illiterate.*

*These are the monster’s words.”*

— *“It’s the Poverty,” Cherríe Moraga [Mor83]*



## TABLE OF CONTENTS

|          |  |           |
|----------|--|-----------|
| <b>1</b> | <b>Introduction</b>  | <b>1</b>  |
| 1.1      | Motivation and Background  | 1         |
| 1.2      | Research Objectives  | 3         |
| 1.3      | Contributions  | 3         |
| 1.4      | Thesis Structure   | 4         |
| <b>2</b> | <b>Background</b>  | <b>7</b>  |
| 2.1      | Gender in Natural Language Processing  | 7         |
| 2.2      | Historical Context & Existing LGBTQIA+ Harms   | 9         |
| 2.3      | Approaches to Gender Bias Mitigation in Language Technologies                            | 10        |
| <b>I</b> | <b>Building Gender-Inclusive Language Models</b>   | <b>13</b> |
| <b>3</b> | <b>Gender-Diverse Erasure in Open Language Generation</b>                                | <b>14</b> |
| 3.1      | Introduction   | 14        |
| 3.2      | Related Work   | 16        |
| 3.3      | TANGO - Misgendering Dataset   | 17        |
| 3.4      | Experiments  | 19        |
| 3.5      | Discussion   | 30        |
| <b>4</b> | <b>Mapping Misrepresentation: Understanding Representational Skew in Language Models</b> | <b>32</b> |
| 4.1      | Introduction   | 32        |

|           |  |           |
|-----------|--|-----------|
| 4.2       | Related Works . . . . .  | 34        |
| 4.3       | Methodology . . . . .  | 35        |
| 4.4       | Results . . . . .  | 38        |
| 4.5       | Discussion . . . . .   | 45        |
| <b>5</b>  | <b>From Skew to Erasure: The Role of Tokenization in Gender-Diverse Bias Propagation and Mitigation Strategies . . . . .</b> | <b>47</b> |
| 5.1       | Introduction . . . . .   | 47        |
| 5.2       | Background & Related Works . . . . .   | 50        |
| 5.3       | Low-Resource Challenges for BPE . . . . .  | 51        |
| 5.4       | Tracing LLM Misgendering to Grammatical Deficiencies . . . . .   | 52        |
| 5.5       | Improving LLM Neopronoun Proficiency . . . . .   | 56        |
| 5.6       | Experimental Setup . . . . .   | 60        |
| 5.7       | Results . . . . .  | 62        |
| 5.8       | Discussion . . . . .   | 66        |
| <b>II</b> | <b>Technical Choices Meet Social Consequences</b>  | <b>68</b> |
| <b>6</b>  | <b>Beyond the Binary: Refining Conceptual Models of Gender-Inclusive Bias Evaluation and Mitigation . . . . .</b>            | <b>69</b> |
| 6.1       | Introduction . . . . .   | 69        |
| 6.2       | Intersectionality on the Ground: A Framework for Social Grounding in AI Fairness   | 70        |
| 6.3       | Intersectionality Illuminates Gaps in Gender Bias Benchmark Construct Validity .   | 71        |
| 6.4       | Persistence of Conceptual Gaps in LLM Gender Bias Evaluation and Mitigation . .  | 73        |

|          |   |           |
|----------|---|-----------|
| 6.5      | On Social Consequences Behind Statistical Assumptions in Common AI Fairness Mitigations . . . . .                     | 75        |
| 6.6      | Discussion . . . . .  | 81        |
| <b>7</b> | <b>In Action: Employing Socially Grounded Bias Assessments for Pretrained and Preference-finetuned LLMs . . . . .</b> | <b>82</b> |
| 7.1      | Introduction . . . . .  | 82        |
| 7.2      | Related Works . . . . .   | 84        |
| 7.3      | Developing the TANGO - Disclosure Dataset . . . . .   | 84        |
| 7.4      | Pretrained Language Model Evaluations . . . . .   | 89        |
| 7.4.1    | Experimental Setup . . . . .  | 89        |
| 7.4.2    | Results . . . . .   | 91        |
| 7.5      | Evaluating Chat-based LLMs with Human Feedback . . . . .  | 94        |
| 7.5.1    | Preference Fine-tuning Overview . . . . .   | 94        |
| 7.5.2    | Experimental Setup . . . . .  | 95        |
| 7.5.3    | Results . . . . .   | 95        |
| 7.6      | Discussion . . . . .  | 97        |
| <b>8</b> | <b>Socially Grounded Bias Detection in Other Domains: Clinical NLP . . . . .</b>                                      | <b>98</b> |
| 8.1      | Introduction . . . . .  | 98        |
| 8.2      | Background and Related Work . . . . .   | 99        |
| 8.3      | Methodology . . . . .   | 101       |
| 8.4      | Experimental Setup . . . . .  | 102       |
| 8.5      | Results . . . . .   | 103       |
| 8.6      | Discussion . . . . .  | 107       |

**9 Conclusion . . . . . 109**

## LIST OF FIGURES

|     |   |    |
|-----|---|----|
| 1.1 | Grammarly flags the gender-diverse pronoun 'xe' as incorrect in the sentence above and suggests replacing it with the binary pronoun 'he', demonstrating automated writing tools' failure to recognize gender-diverse language. . . . .   | 2  |
| 3.1 | Our template-based misgendering evaluation framework. Templates are gathered from Nonbinary Wiki and populated with various referent forms and pronouns, then fed to an LLM. The resulting generated text is evaluated for misgendering. . . . .  | 16 |
| 3.2 | Distribution of pronoun consistency (left) and perplexity (right) across 9 models. Templates with binary pronouns consistently result in the least misgendering across model sizes. . . . .   | 24 |
| 3.3 | Pronoun Template vs Pronouns in Generations. From left to right: GPT2, GPT-Neo, OPT, All . . . . .  | 27 |
| 3.4 | Pronoun Template Type vs Errors in Generations. From left to right: GPT2, GPT-Neo, OPT, All . . . . .   | 28 |
| 3.5 | Pronouns generated using respective pronoun template types when using only non-binary names or distal antecedents. From left to right: GPT2, GPT-Neo, OPT, ChatGPT . . . . .  | 30 |
| 4.1 | PCA Components for each Gender Subspace, indicating binary gender takes the least amount of components to represent, and is therefore less complex to model. . . . .  | 45 |
| 5.1 | Byte-Pair Encoding (BPE) tokenization disproportionately fragments neopronouns compared to binary pronouns due to their infrequency in the training corpus. Our paper reveals that this overfragmentation leads to syntactic difficulties for LLMs, which are tied to their propensity to misgender data-scarce pronouns. . . . . | 48 |

|     |  |    |
|-----|--|----|
| 5.2 | Evaluation. We determine our method’s efficacy in reducing LLM misgendering using a constrained decoding approach across 3 metrics. . . . .  | 54 |
| 5.3 | Overview. We (1) tokenize neopronouns using PTP for a given LLM, (2) either fully finetune or only finetune the LLM lexical layer with data containing neopronouns, and (3) determine our method’s efficacy in reducing LLM misgendering using a constrained decoding approach across 3 metrics. . . . . | 59 |
| 5.4 | 70M model pronoun consistency for each pronoun family across 10-50% data resource levels and model variants. <i>Takeaway: PTP sustains improvements in neopronoun consistency across data resource levels.</i> . . . . .   | 62 |
| 5.5 | Results across all models at data resource level=10. The uncertainty estimates are 95% confidence intervals computed from 10k bootstrap iterations. <i>Takeaway: Across model size, variants of PTP consistently improve neopronoun consistency over models employed with standard BPE.</i> . . . . .    | 65 |
| 6.1 | Bias benchmarks employed by top 15 performing preference-tuned LLMs reported by Chatbot Arena Leaderboard across socially-relevant categories. Evaluations fully cover binary gender bias, with limited evaluation for gender-diverse minorities and other socially-salient dimensions. . . . .          | 75 |
| 7.1 | Collection of gender disclosure prompts. We locate intro sections of TGNB identities from Nonbinary Wiki. Then we extract the first description of a person’s gender and convert it to a gender disclosure template. . . . .   | 85 |
| 7.2 | Proportion of toxic generations based on Perspective API toxicity reported across models and in aggregate. . . . .   | 91 |
| 7.3 | Differences in toxicity scores between static and dynamic gender disclosures across TGNB and binary genders. Dots left of the dotted black line indicate toxicity scores are <i>lower</i> for dynamic disclosures than static disclosure forms. . . . .  | 93 |

|     |  |     |
|-----|--|-----|
| 7.4 | <i>Left:</i> Difference in percent of texts classified as negative regard (TGNB-Binary), with 95% confidence intervals included over 10k bootstrap iterations. TGNB bias amplification (red) from baseline seen in majority of models with SFT+DPO, while DPO alone typically reduced amplification (blue). Black bold is significantly ( $\rho < 0.05$ ) different than base model. <i>Right:</i> % of texts labeled as negative regard across gender groups, textual disclosure forms, and model alignment stages. . . . . | 95  |
| 7.5 | Example of negative regard amplified from 5% to $\geq 90\%$ after DPO, prompt is bold. . . . .   | 96  |
| 8.1 | t-SNE results with circled most biased cluster for <b>HAS DIABETES</b> attribute . . . . .   | 104 |
| 8.2 | Performance differences for <b>HAS DIABETES</b> attribute. Furthest right red box shows global bias, while SLOGAN finds a local area of much higher bias at cluster 4. . . . .   | 107 |

## LIST OF TABLES

|      |   |    |
|------|---|----|
| 3.1  | Pronouns and pronoun types split across prompts . . . . .   | 18 |
| 3.2  | Misgendering Prompt Set Statistics (N=2,400). . . . .   | 20 |
| 3.3  | Consistency metrics for the AMT experiments and automatic tool. Pronoun consistency, relevance, coherence, and type-token ratio are reported based on AMT experiments. Bold values indicate the highest score in each category per model. . . . . | 23 |
| 3.4  | Differences in misgendering and perplexity across antecedents with varying social contexts. $\Delta$ reflects the absolute difference between Named and Distal antecedent forms. . . . .  | 26 |
| 4.1  | Set of unpleasant and pleasant words . . . . .  | 36 |
| 4.2  | Word set definitions for binary and non-binary concepts . . . . .   | 37 |
| 4.3  | Frequency of Gender-related pronouns (left) and terms (right) for English Wikipedia, reported per billion. Frequencies reflect skew towards binary gender-related content. . . . .  | 38 |
| 4.4  | Example sentences containing nonbinary pronouns . . . . .   | 39 |
| 4.5  | Nearest neighbor words in GloVe for binary and non-binary pronouns and their possessive forms. . . . .  | 40 |
| 4.6  | Ten Nearest neighbors of non-binary terms highlighting derogatory Terms . . . . .   | 42 |
| 4.7  | WEAT Scores (vs. pleasant and unpleasant attributes) . . . . .  | 43 |
| 4.8  | Average cosine similarity between occupations and nominative pronouns. . . . .  | 44 |
| 4.9  | Cosine similarity between gendered words and common occupations. . . . .  | 44 |
| 4.10 | Binary and Nonbinary Pronoun Sets for PCA. . . . .  | 45 |



|     |   |     |
|-----|---|-----|
| 5.1 | BPE-tokenized Binary Pronouns and Neopronouns across pronoun forms. $\zeta$ = Fertility. The closer fertility is to 1, the more the tokenizer kept pronoun tokens fully intact. <b>Bold</b> = neopronoun tokenization that does not follow binary pronoun forms. . . . .  | 51  |
| 5.2 | Out-of-the-box evaluations on Pythia, a GPTNeo-X based model across sizes. Uncertainty estimates are 95% confidence intervals computed from 10k bootstrap iterations. <i>Takeaway: Markedly higher grammatical error rates for neopronoun vs. binary pronouns.</i> . . . . .  | 57  |
| 5.3 | 70M-parameter model results at 10% data resource level. $T_{\text{ORIG}}$ = original BPE tokenizer, $T_{\text{PTP}}$ = tokenizer with PTP, $M_{\text{BASE}}$ = original model (no finetuning) $M_{\text{FULL}}$ = full finetuning. Uncertainty estimates are 95% confidence intervals computed from 10k bootstrap iterations. . . . . | 63  |
| 5.4 | Pythia-410M model generations across finetuning regimes. <i>Italics</i> are input prompts and generations are performed with nucleus sampling ( $\text{TOP-P}=0.95$ ). . . . .  | 66  |
| 6.1 | Cosine similarity between gendered words and common occupations. . . . .  | 69  |
| 6.2 | Bias evaluation modalities for top 15 performing LLM families reported by Chatbot Arena Leaderboard. . . . .  | 74  |
| 7.1 | Top 10 most frequently identified TGNB Identities from Nonbinary Wiki . . . . .   | 86  |
| 7.2 | Gender Disclosure Prompt Set Statistics (N=1,422,720). . . . .  | 88  |
| 7.3 | Model generations with the highest proportion of toxic text. Template: <i>[referent] is [gender identity]</i> . . . . .   | 92  |
| 8.1 | Percent of attribute in the MIMIC-3 data . . . . .  | 102 |

8.2 Average values for 12 MIMIC-III attributes across models and evaluation metrics. SCR, SIR, and |Bias| in %. |Bias| is the average absolute model performance difference in biased clusters. Bold is the best performance per row. Right-most column is number of MIMIC-III attributes where SLOGAN performs best. Arrows indicate desired direction of a number. . . . . 104

8.3 Bias detection (%) for in-hospital mortality task. Global indicates global bias. “Yes” indicates patient with diabetes. |Bias| is the max absolute model performance difference in biased clusters. SLOGAN identifies local biases greater than global bias observed in the data (bold). . . . . 105

8.4 Comparison under diabetes attribute. SCR and SIR are respectively the % of biased clusters and % of biased instances. |Bias|(%) is the average absolute bias score for the biased clusters. SLOGAN finds the largest bias (bold). . . . . 105

8.5 Top 20 topic words in the most and least biased clusters using SLOGAN for **HAS DIABETES** attribute. Number is the bias score (%) of that cluster. . . . . 107

## ACKNOWLEDGMENTS

First and foremost, Professor Kai-Wei Chang, thank you for your unwavering support throughout my PhD - your feedback, wisdom, and sense of humor were key to making it happen! Thank you to my committee members, Professor Wei Wang, Professor Yizhou Sun, and Dr. Keith Norris. Thank you to Professor Majid Sarrafzadeh for his support! Thank you to Dr. Bitu Amani, Dr. Chandra Ford, and Dr. Lisa Bowleg for serving as such strong examples of critical scholarship. Dr. Ellesse Akre and Dr. Paris Adkins-Jackson, thank you for opening my eyes to everything beyond a p-value. Arjun Subramonian, William Agnew, Vagrant Gautam, Ashwin Singh, Kruno Lehman, and other friends at Queer in AI, thank you for being such pillars of resistance and reimagination for queer futures. Sunipa Dev, Remi Denton, Zeerak Talat - thank you for your mentorship throughout my PhD. Shout out to brilliant leaders in this field I've had the honor of working with like Levent Sagun, Rahul Gupta, Ninareh Mehrabi, and Palash Goyal. Thank you to wonderful collaborators and labmates Evan Czyzycki, Rosa Garza, Orpaz Goldstein, Mohammad Kachuee, Davina Zamanzede, Shayan Fazeli, Christina Chance, Elaine Wan, Ashima Suvarna, Hritik Bansal, Tanmay Parekh, Jieyu Zhao, Di Wu, Harold Li, Masoud Monjatipour, and many others! Marie-Therese Png and Shakir Mohamed, thank you for being such a source of presence and kindness in this field. Gracias a la Dra. Paula Ricaurte y fA+Ir América Latina y Caribe Hub - su confianza y las oportunidades que me han dado han sido fundamentales para mi desarrollo académico, profesional, personal y espiritual. Thank you to my friends and family for their support and humor throughout this journey. Issy, Lia, Emily, Amanda, Charlene, Laura, Romi, Jason, Diana, Adaku, Renata, Lil, Nassor, Alexandria - WE DID IT, AHHH!!! Shout out to Issy for always looking out for my academic future & connecting me with opportunities for growth! Thank you to my mentors that helped me literally get to UCLA, Diane Woodbridge, Nathaniel Stevens, Sampsa Jaatinen, Peter Wais, Sophie Engle. And a very special shout out to Joseph Brown and the UCLA CS Grad Office. Joseph, thank you for picking up each and every time I called you to figure out my PhD application and how to even do any of this! Juliana Alvarez and Madelen Hem - thank you for all your guidance! Edna Todd, thank you for being such a source of light.

## VITA

- 2017            B.S Data Science, University of San Francisco
- 2021            M.S. Computer Science, University of California, Los Angeles

## SELECT PUBLICATIONS

**Ovalle A**, Pavasovic KL, Martin L, Zettlemoyer L, Smith EM, Williams A, Sagun L. The Root Shapes the Fruit: On the Persistence of Gender-Exclusive Harms in Aligned Language Models. 2024 Neurips Queer in AI Workshop. 2024 Nov 6.

**Ovalle A**, Mehrabi N, Goyal P, Dhamala J, Chang KW, Zemel R, Galstyan A, Pinter Y, Gupta R. Tokenization matters: Navigating data-scarce tokenization for gender inclusive language technologies. In Findings of the Association for Computational Linguistics: NAACL 2024 2024 Jun (pp. 1739-1756).

**Ovalle A**, Liang D, Boyd A. Should they? Mobile Biometrics and Technopolicy Meet Queer Community Considerations. In Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization 2023 Oct 30 (pp. 1-10).

Dennler N, **Ovalle A**, Singh A, Soldaini L, Subramonian A, Tu H, Agnew W, Ghosh A, Yee K, Paradejordi IF, Talat Z. Bound by the Bounty: Collaboratively Shaping Evaluation Processes for Queer AI Harms. In Proceedings of the 2023 AAI/ACM Conference on AI, Ethics, and Society 2023 Aug 8 (pp. 375-386).

**Ovalle A**, Goyal P, Dhamala J, Jaggars Z, Chang KW, Galstyan A, Zemel R, Gupta R. “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 2023 Jun 12 (pp. 1246-1266).

Queerinaï OO, **Ovalle A**, Subramonian A, Singh A, Voelcker C, Sutherland DJ, Locatelli D, Breznik E, Klubicka F, Yuan H, Zhang H. Queer in AI: A case study in community-led participatory AI. In Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency 2023 Jun 12 (pp. 1882-1895).

**Ovalle A**, Subramonian A, Gautam V, Gee G, Chang KW. 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 2023 Aug 8 (pp. 496-511).

**Ovalle A**, Dev S, Zhao J, Sarrafzadeh M, Chang KW. Auditing algorithmic fairness in machine learning for health with severity-based LOGAN. In International Workshop on Health Intelligence 2023 Feb 13 (pp. 123-136). Cham: Springer Nature Switzerland.

**Ovalle A**, Dev S, Zhao J, Sarrafzadeh M, Chang KW. Auditing algorithmic fairness in machine learning for health with severity-based LOGAN. In International Workshop on Health Intelligence 2023 Feb 13 (pp. 123-136). Cham: Springer Nature Switzerland.

Dev S, Monajatipoor M, **Ovalle A**, Subramonian A, Phillips J, Chang KW. 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 2021 Nov (pp. 1968-1994).

# CHAPTER 1

## Introduction

### 1.1 Motivation and Background

AI-driven language models are rapidly becoming the backbone of everyday technologies, influencing how we communicate, access information, and make decisions. Large language models (LLM) like ChatGPT<sup>1</sup> and Claude<sup>2</sup> now offer virtual assistants [CLL23, KYW24] and creative content generation tools [LS22, ASO23], introducing functionalities previously considered inaccessible. However, despite their success in generating human-like text, significant challenges remain in ensuring these systems are fair, inclusive, and respectful of all individuals. LLMs learn patterns from large datasets, but these datasets are often imperfect and fail to reflect the full diversity of human experience [ZWY18a]. As a result, these models can perpetuate and amplify harmful stereotypes [BCZ16], reflect stigmatizing language [HPD20b], and exclude underrepresented identities [DMO21]. Historically marginalized groups, such as transgender and non-binary individuals, are particularly vulnerable to these exclusions, as AI systems struggle to represent identities that diverge from binary gender conventions.

The transgender and non-binary (TGNB) community disproportionately faces discrimination and exclusion in daily life [PAR23], and AI systems exacerbate these harms when they fail to account for diverse gender identities. Prior research has documented cases of directed toxic language [QSA21, NBL22, QOS23] and the overfiltering of queer individuals in content moderation

---

<sup>1</sup><https://chat.openai.com/>

<sup>2</sup><https://claude.ai/>

Alex went to the market because xe wanted to buy apples.

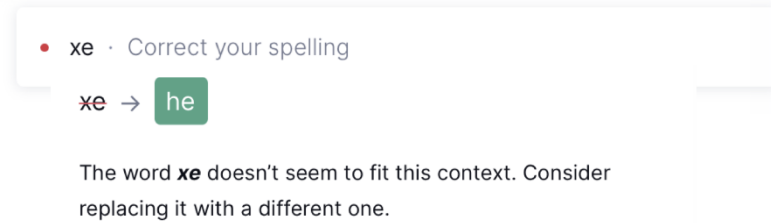


Figure 1.1: Grammarly flags the gender-diverse pronoun 'xe' as incorrect in the sentence above and suggests replacing it with the binary pronoun 'he', demonstrating automated writing tools' failure to recognize gender-diverse language.

systems [FCJ22, WGU21].

NLP-based systems can exhibit systematic failures in handling gender diversity, perpetuating harmful biases through their technical limitations. A key example can be seen in named entity recognition (NER) systems, which may fail to classify non-binary chosen names correctly or even misclassify them as objects rather than people [DMO21]. Even mainstream tools demonstrate this problem - for instance, as shown in Figure 1.1, Grammarly fails to recognize gender-diverse pronouns like 'xe', and instead wants to autocorrect to the binary pronoun 'he'. These technical limitations have real consequences: they de-legitimize transgender and non-binary identities and can lead to real-world discrimination when automated systems don't recognize someone as a person simply because they don't fit into traditional gender categories [RD19a].

This dissertation explores how biases in AI-driven language models are formed and sustained, both within the models themselves and across the broader AI ecosystem in which they are situated. Through centering transgender and non-binary experiences, we develop generalizable approaches for understanding and addressing algorithmic bias. Two key research questions guide this work: **RQ1:** What specific technical challenges do large language models face in fairly representing transgender and non-binary identities, and how can we develop more inclusive approaches that

better serve TGNB communities? **RQ2:** What broader gaps in AI fairness frameworks allow for the exclusion of marginalized groups, and how can these frameworks be revised to promote more just outcomes? By investigating these challenges, this research aims to develop new techniques and sociotechnical frameworks that enable AI researchers to develop LLMs which are more inclusive and human-centric in their design and real-world application, thereby better able to represent diverse sociocultural identities and effectively serve a wide range of users.

The following next sections outline the research objectives, contributions, and structure of this dissertation.

## 1.2 Research Objectives

The primary objectives of this thesis are:

1. **Understand Biases in Gender-Diverse Representation:** Investigate representational harms and exclusionary patterns in LLMs which affect TGNB identities and other gender-diverse persons.
2. **Address Gender-Diverse Harms:** Identify mechanisms driving harmful LLM bias propagation and develop technical interventions to mitigate these unwanted behaviors.
3. **Develop Community-Informed Benchmarks and Evaluation Frameworks:** Demonstrate how community-centered practices can be applied to detect, understand, and mitigate bias in LLMs.

## 1.3 Contributions

This thesis makes the following key contributions:

- **Gender-Diverse Biases Identified in Foundational and Preference-Aligned LLMs:** We identify ways in which both foundational and preference fine-tuned LLMs can perpetuate undesirable social biases reflective of real-world gender-diverse stigma.



- **Tokenization Contributes to the Propagation of Gender-Diverse Bias in LLMs:** We discover how Byte-Pair Encoding (BPE) tokenization, the most popular form of LLM tokenization, can propagate representational erasure of TGNB persons in low-resource settings and propose novel mitigation strategies to address these issues.
- **Adopting Hegemonic Bias Evaluations Identified as Barrier to Operationalizing Gender-Inclusive LLMs:** We identify fundamental limitations in popularly employed gender bias evaluations, revealing their inability to capture harms against TGNB identities. We find a concerning feedback loop where limited evaluation frameworks constrain mitigation efforts, ultimately reinforcing existing biases. Intersectionality illuminates novel approaches towards both technical and systemic mitigations against algorithmic bias.
- **Community-Centric Benchmarking Overcomes Barriers** We introduce the Towards Centering Transgender and Non-Binary Voices to Measure Biases in Open Language Generation (TANGO) dataset, comprising two subsets focused on misgendering and gender disclosure. This dataset serves as a benchmark for evaluating AI systems’ handling of Gender Non-Affirmative language. Demonstrating the value in social-grounding, we also introduce SLOGAN, a framework for effective local bias detection for clinical NLP tasks informed by patient medical and social histories.

## 1.4 Thesis Structure

This dissertation bridges technical and systemic approaches to addressing algorithmic bias in language technologies. Following an introduction to key concepts in Chapter 2, the first section (Chapters 3-5) examines how language models systematically erase and misrepresent transgender and non-binary identities. Here, we propose novel technical interventions, particularly through tokenization and low-resource NLP strategies, to address these harms. The second section (Chapters 6-8) shifts focus to analyze how existing bias evaluation frameworks themselves can perpetuate harm. We develop more socially grounded approaches to bias evaluation, demonstrating their effectiveness through case studies in both gender representation and clinical applications.

In **Chapter 3**, we introduce the TANGO dataset, one of the key contributions of this dissertation. Specifically, we develop the misgendering subset of TANGO, which serves as a benchmark for evaluating AI systems’ ability to handle TGNB language. This chapter presents an in-depth analysis of how existing LLMs fail in terms of misgendering TGNB individuals, highlighting the representational harms that arise from biased model outputs.

**Chapter 4** investigates the role of data infrequency and its relation to biased language representation, embeddings, and tokenization. This is another core contribution, where we find that Byte-Pair Encoding (BPE) tokenization fragments neopronouns due to their low frequency in training data, leading to misrepresentations and reduced understanding of gender-diverse pronouns.

Building on this, **Chapter 5** addresses these data scarcity issues by proposing techniques such as pronoun tokenization parity and low resource NLP considerations, demonstrating significant improvements in how LLMs handle gender-diverse language and neopronouns.

The second part of this dissertation begins with **Chapter 6**, which critically examines existing LLM gender bias evaluation frameworks. Through an intersectional lens, we reveal fundamental limitations in how these frameworks conceptualize and measure gender bias, including their failure to capture TGNB harms. Moving towards evaluating *situated harms*, we demonstrate how traditional evaluation methods can miss crucial context inform LLM harm propagation. We draw on intersectionality as an analytical framework to develop more inclusive evaluation frameworks that move beyond binary gender assumptions and better reflect the lived experiences of gender-diverse individuals.

**Chapter 7** introduces the community-sourced gender disclosure subset of the TANGO dataset, providing an evaluation of biased associations for LLMs with respect to gender non-affirmative language. With this, we uncover ways harmful and stigmatizing outputs are encoded in both foundational and preference fine-tuned LLMs, aspects traditional bias evaluations are unable to detect.

Finally, in **Chapter 8**, we demonstrate the broader applicability of socially-grounded techniques beyond gender bias by introducing SLOGAN, a framework designed to identify local biases in

clinical prediction tasks, shedding light on how machine learning biases can reinforce existing healthcare disparities. We conclude our work in *Chapter 9*.

# CHAPTER 2

## Background

This chapter establishes the social and technical context needed to understand current challenges in gender-inclusive language technologies, laying groundwork for the research presented in subsequent chapters. We begin by covering how concepts of gender identity and expression manifest in language model representations and their outputs. We then detail how AI-propagated harms against gender-diverse communities take place, highlighting patterns of erasure and experienced marginalization. Finally, we discuss common approaches to measuring and mitigating gender bias in these systems.

### 2.1 Gender in Natural Language Processing

#### 2.1.0.1 Understanding Gender Identity

Modeling gender in natural language systems requires careful distinction between gender identity, expression, and biological sex. **Gender identity** represents an individual’s internal self-conceptualization of gender. **Gender expression** encompasses observable characteristics including presentation, mannerisms, and social behaviors [RD19b]. **Biological sex** comprises physical sex characteristics (i.e., primary and secondary) [RD19b]. Current NLP systems often conflate gender identity, expression, and sex as a single variable. However, empirical studies demonstrate these are independent features [Ser07], requiring distinct representations in computational models to avoid systematic bias. This independence poses fundamental challenges for NLP systems that often implicitly assume correlations between these variables, leading to systematic erasure in tasks like coreference resolution and machine translation [CD20]. This distinction is crucial, as con-

flating these concepts in NLP result in varying sociotechnical translations that ripple into model design and bias evaluation - aspects we cover throughout this dissertation.

Reflecting Western conceptualizations of gender, current NLP systems restrict gender representation to a binary encoding: male/female gender identity, and alignment with birth-assigned gender (cisgender/transgender) [RD19b]. However, this binary encoding fails to capture the full spectrum of gender identities observed globally. Non-Western cultures recognize diverse gender identities, including the Jogappas (Karnataka), Muxes (Oaxaca), and Mahuwahines (Hawai'i) [Des18, Mir16, Cla19], which resist classification within Western binary schemas [Mir16, TYB19]. Similarly, persons identifying as genderfluid do not identify with a single gender, nonbinary may encompass all identities outside the binary framework, while agender individuals do not subscribe to gender at all [RD19b] [RD19b]. These varied experiences of gender, which can shift with time, fundamentally challenge NLP systems' reliance on static binary representations [Web19].

### **2.1.0.2 Linguistic Gender Markers**

Languages frequently mark gender through linguistic features [CD20]. In English, pronouns serve as primary gender markers but lack clear one-to-one mappings with gender identity. English pronouns include traditional binary forms (he/she), singular they which spans both binary and nonbinary usage, and gender-neutral neopronouns (e.g., xe/xem, ze/hir) [Cla19, Fer16]. Neopronouns are a special example of evolving language usage patterns, where pronouns span multiple gender categories and a single individual may use multiple pronoun sets. As pronoun usage demonstrates both multiplicity and context-dependence: individuals may use different pronouns across social contexts, multiple pronouns simultaneously (e.g., both she/her and they/them), specify acceptance of all pronouns or none at all (e.g. only use names) [Fer16]. These flexible usage patterns present significant challenges for traditional NLP approaches to gender that assume static mappings between pronouns and gender[Gau21].

Similarly, NLP systems typically enforce binary gender classification of names, though this approach fails to capture naming practices in gender-diverse communities [Ros20]. Individuals often

choose names that align with their gender identity, departing from traditional naming conventions through the use of nature words or common nouns [Nonnd]. These diverse naming patterns further demonstrate the limitations of binary gender classification in current NLP approaches.

## 2.2 Historical Context & Existing LGBTQIA+ Harms

### 2.2.0.1 Systemic Biases and Algorithmic Harms to the Queer Community

Despite increased visibility of gender diversity, algorithmic systems continue to perpetuate systemic biases against gender-diverse communities [Key18, TMK21]. These biases manifest through systems trained on data reflecting cisnormative assumptions, leading to both representational and allocative harms [Bey21, Sha15]. Language technologies can demonstrate several forms of systematic bias. First, content moderation systems disproportionately flag and filter gender-diverse content [aR20, DBS20], effectively amplifying cisnormative perspectives while suppressing gender-diverse voices [TMK21]. Second, training data reflecting societal biases leads to models that perpetuate discriminatory patterns, creating feedback loops that reinforce existing inequities [PAR23]. These systemic biases manifest particularly in NLP systems that encode and perpetuate binary gender assumptions. A survey by [DMO21] identifies several high-risk application areas where such biases can cause direct harm to users. Below, we examine specific NLP tasks and applications where binary gender encoding creates demonstrable adverse impacts.

### 2.2.0.2 Harms as Misgendering, Erasure, and Biased Associations

**Misgendering** Misgendering occurs when systems incorrectly assign or reference gender identity, manifesting in both systemic and instance-level failures [Spa15]. In language technologies, this harm appears through structural constraints, such as systems that enforce binary gender selection in user interfaces, effectively forcing non-binary users to select inaccurate gender categories [Key18, SKB19]. Language models perpetuate misgendering through multiple mechanisms: defaulting to binary pronouns when gender information is ambiguous [CD20], applying stereotypical gender assumptions in generation tasks [SCN19], and failing to maintain consistent gender refer-

ence even when pronouns are explicitly provided [DMO21]. These misgendering patterns, documented across computer vision [Key18] and human-computer interaction [KHB21], contribute to adverse mental health impacts among gender-diverse users.

**Erasure** Erasure describes the systemic invalidation or obscuring of non-binary gender identities in language technologies [Ser07, RD19b] and manifests through multiple technical failures across such systems. For instance, Named Entity Recognition (NER) systems can fail to recognize non-binary chosen names as referring to persons, particularly names that deviate from conventional patterns (e.g., single-letter names) [DMO21]. Similarly, coreference resolution systems struggle with neopronouns and singular "they," either treating them as unknown tokens or failing to maintain consistent reference [CD19, DMO21].

Systems like Genderify reflect systematic erasure by attempting to classify gender as binary based solely on names and usernames, fundamentally erasing non-binary identities [Lau20, SKB19]. In practice, these technologies can perpetuate erasure through feedback loops [HSN18, Sap21]: language models trained on binary-gendered corpora reflect societal biases [Lak, Fis93], which influences content creation [FVB16], further amplifying non-binary erasure. These cyclical harms stem from model and dataset biases, including tainted examples, limited features, and sample size disparities [WZY19, BHN19], rooted in broader societal non-recognition of gender diversity [MAP16, RD19b]. While recent work has introduced gender-inclusive datasets and expanded bias metrics [CD20, RNL18], significant challenges persist in mitigating gender-related harms. The following section reviews current mitigation approaches, which we extend toward more gender-inclusive frameworks in subsequent chapters.

## 2.3 Approaches to Gender Bias Mitigation in Language Technologies

NLP systems can encode and amplify gender biases from their training data. Current mitigation approaches target different stages of the pipeline: pre-processing methods modify training data, in-processing approaches adjust model training, and post-processing techniques calibrate outputs. In the following subsections, we review these methods to contextualize our contributions towards

gender-inclusive NLP.

**Pre-processing** Pre-processing approaches target bias mitigation by modifying training data before model training, aiming to create more balanced and representative datasets [SGT19]. Data augmentation represents a major category of pre-processing methods. Approaches like Gender-Swap [ZWY18a] and counterfactual data substitution [MGC19] create counterfactual examples by exchanging gendered terms. Other variants consider name interventions [MGC19]. Researchers have also explored a combination of both counterfactual data augmentation and dropout techniques to reduce unintended correlations with gender [WRA20]. Template-based approaches are also commonly employed to provide intersectional and consider large variation in cultural context [SGQ19, MWB19].

**In-processing** In-processing methods mitigate bias during model training through modifications to learning objectives and model architectures [ZLM18]. Common approaches include constrained optimization, where models must satisfy fairness criteria (e.g., equal performance across gender groups) while optimizing for task performance [ZWY19]. Recent work has expanded these methods to handle multiple fairness constraints simultaneously [SFG21] and developed techniques for balancing competing objectives [WSK21]. Researchers have also proposed methods for fairness aware online learning [ZMW24] and techniques for handling a family of fairness constraints at once [CHK19]. Another strategy uses adversarial debiasing through information bottlenecks to ensure model predictions remain independent of gender attributes [EG18, GBC22]. Debiasing representation spaces, particularly embeddings, represents another significant approach. While early work focused on removing gender bias from static word embeddings [BCZ16, ZZL18], recent approaches address contextual embeddings [ZWY19] and intersectional bias [LGP20].

**Post-processing** Post-processing techniques mitigate bias after model training by adjusting model outputs or learned representations without modifying the original model or training process. These approaches are particularly valuable in the context of rapid deployment needs or when retraining large models is impractical. One common approach focuses on prediction calibration to satisfy



different fairness constraints. Equalized odds [HPS16] requires that predictions have equal true positive rates and false positive rates across protected groups, achieved through post-hoc threshold adjustments for each group. In contrast, demographic parity requires the proportion of positive predictions to be equal across groups, regardless of the ground truth [DHP12]. Other forms of post-processing also consider bias-aware decoding strategies in generative LLMs. These approaches modify generation probabilities during inference to reduce stereotypical or harmful associations [SCN20]. Plug and play decoding techniques [DML19] and self-debiasing methods[SUS21] have also been employed for controlled text generation via LLM decoding.

**Part I**

**Building Gender-Inclusive Language  
Models**

## CHAPTER 3

# Gender-Diverse Erasure in Open Language Generation

A significant body of research on gender bias in language models has centered on binary gender and the stereotypes associated with masculine and feminine attributes [BCZ16, WRA18, DLP20]. While these studies have improved our understanding of binary gender bias in tasks like coreference resolution and machine translation [MGM20, ZWY18b, SSZ19a], they often neglect nonbinary and other non-cisnormative identities. This chapter examines how large language models can perpetuate biases in open language generation, with a focus on the challenges faced by transgender and nonbinary (TGNB) individuals. We identify the gaps that perpetuate TGNB bias and discuss their implications for LLM mitigation. This chapter is based on our works [DMO21] and [OGD23].

### 3.1 Introduction

Large language models (LLM) are being increasingly utilized for open language generation (OLG) for content creation (e.g., story creation) and conversational AI (e.g., voice assistants, voice user interfaces). However, recent studies demonstrate how LLMs may propagate or even amplify existing societal biases in the form of harmful, toxic, and unwanted associations [WGU21, SCN21, SCN19]. Historically marginalized communities, including but not limited to the *LGBTQIA+*<sup>1</sup> community, disproportionately experience discrimination and exclusion from social, political and economic dimensions of daily life [Hewnd]. Creating more inclusive LLMs must sufficiently include those at the highest risk for harm. Therefore in this paper, we illuminate ways in which harms

---

<sup>1</sup>All italicized words are defined in [https://nonbinary.wiki/wiki/Glossary\\_of\\_English\\_gender\\_and\\_sex\\_terminology](https://nonbinary.wiki/wiki/Glossary_of_English_gender_and_sex_terminology)

may manifest in OLG for members of the *queer*<sup>2</sup> community, specifically those who identify as *transgender* and *nonbinary*.

Varying works in natural language fairness research examine differences in possible representational and allocational harms [BHN22] present in LLMs for TGNB persons. In NLP, studies have explored misgendering with pronouns<sup>3</sup> [DMO21, AH13], directed toxic language [QSA21, NBL22], and the overfiltering content by and for queer individuals [WGU21, FCJ22]. However, in NLG, only a few works (e.g., [SCN20, SQX20, NBL22]) have focused on understanding how LLM harms appear for the TGNB community.

In open language generation, one way to evaluate potential harms is by prompting a model with a set of seed words to generate text and then analyzing the resulting generations for unwanted behavior [DSK21, WGU21]. Likewise, we can assess this for our context by giving models prompts and evaluating their generated text for misgendering using pronouns (Figure 3.1). We ground our work in natural human-written text from the Nonbinary Wiki<sup>4</sup>. Specifically, we make the following contributions:

- (1) Provided the specified harms experienced by the TGNB community, we release *TANGO - Misgender*<sup>5</sup>, a dataset (T)ow(A)rds centering tra(N)s(G)ender and nonbinary voices to evaluate gender non-affirmation in (O)LG consisting of a misgendering evaluation set of 2,880 prompts to assess pronoun consistency<sup>6</sup> across various pronouns, including those commonly used by the TGNB community along with binary pronouns<sup>7</sup>.

---

<sup>2</sup>We use the terms LGBTQIA+ and queer interchangeably. We acknowledge that queer is a reclaimed word and an umbrella term for identities that are not heterosexual or not cisgender. Given these identities' interlocking experiences and facets, we do not claim this work to be an exhaustive overview of the queer experience.

<sup>3</sup>The act of intentionally or unintentionally addressing someone (oneself or others) using a gendered term that does not match their gender identity.

<sup>4</sup><https://nonbinary.wiki/>

<sup>5</sup><https://github.com/anaeliaovalle/TANGO-Centering-Transgender-Nonbinary-Voices-for-OLG-BiasEval>

<sup>6</sup>Addressing someone using a pronoun that *does* match their gender identity. Being consistent in pronoun usage is the opposite of misgendering.

<sup>7</sup>In this work we use this term to refer to gender-specific pronouns he and she which are typically associated to the

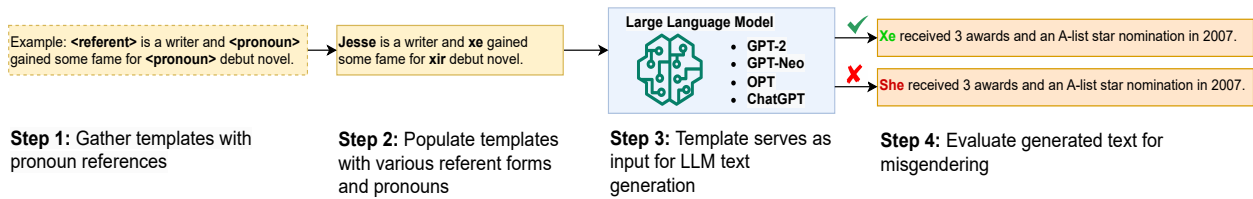


Figure 3.1: Our template-based misgendering evaluation framework. Templates are gathered from Nonbinary Wiki and populated with various referent forms and pronouns, then fed to an LLM. The resulting generated text is evaluated for misgendering.

- (2) Guided by interdisciplinary literature, we create an automatic misgendering evaluation tool and translational experiments to evaluate and analyze the extent to which gender non-affirmation is present across four popular large language models: GPT-2, GPT-Neo, OPT, and ChatGPT using our dataset.
- (3) With these findings, we provide constructive suggestions for creating more gender-inclusive LLMs in each OLG experiment.

We find that misgendering most occurs with pronouns used by the TGNB community across all models of various sizes. LLMs misgender most when prompted with subjects that use neopronouns (e.g., *ey*, *xe*, *fae*), followed by singular they pronouns (§3.4.0.2). When examining the behavior further, some models struggle to follow grammatical rules for neopronouns, hinting at possible challenges in identifying their pronoun-hood (§3.4.0.4). Furthermore, we observe a reflection of binary gender<sup>8</sup> norms within the models. Results reflect more robust pronoun consistency for binary pronouns (§3.4.0.3) and the usage of generic masculine language during OLG (§3.4.0.4).

## 3.2 Related Work

**TGNB Harm Evaluations in LLMs** Gender bias evaluation methods include toxicity measurements and word co-occurrence in OLG [SCN19, SCN21, DSK21, LWW20, DFW19, LB21].

genders man and woman respectively, but acknowledge that TGNB may also use these pronouns.

<sup>8</sup>We use this term to describe two genders, *man* and *woman*, which normatively describes the gender binary.

Expanding into work that explicitly looks at TGNB harms, [DMO21] assessed misgendering in BERT, with [LCH22] elaborating on desiderata for pronoun inclusivity. While we also measure misgendering, we assess such behavior in an NLG context using both human and automatic evaluations. [NBH21, NBL22, BLV21] created evaluations on the LGBTQIA+ community via model prompting, then measuring differences in lexicon presence or perceived toxicity by the Perspective API.

**LGBTQIA+ Datasets** Many datasets exist in NLP to assess binary gender inclusivity, including Winogender and the GAP dataset. In NLG, [DSK21] create a dataset of prompts to assess for harms in OLG across various domains (e.g., politics, occupation) using Wikipedia. However, gender-inclusive LLM evaluation requires gender-inclusive datasets. [FCJ22] released WinoQueer, a set of prompts extracted from Tweets by the queer community to assess queer harms with BERT. Similar to our work, [BLV21] created a dataset of Reddit prompts to assess LGBTQIA+ harms across identity terms in a masked language modeling task. [NBL22] build off this by adding more gender identity terms and neopronouns. Our work differs from these in that our dataset contains prompts to measure misgendering and model responses to gender disclosure.

### 3.3 TANGO - Misgendering Dataset

In this work, we propose *TANGO - Misgenderer*, a dataset for assessing gender non-affirmation of TGNB identities, focusing on examining the extent to which the undesired behavior of misgendering is present in open language generation. Moving forward, throughout this work we use the abbreviated version, *TANGO*.

Misgendering<sup>9</sup> is a form of gender non-affirmation experienced by the TGNB population that results in stigmatization and psychological distress [McL18]. To determine if this behavior persists in LLMs, we create a dataset to evaluate misgendering in OLG. In English grammar, pronouns should agree in number, person, and *gender* with their antecedents (i.e., a person, place, thing, or clause which a pronoun can represent), called pronoun-antecedent agreement [Colnd]. Therefore,

---

<sup>9</sup><https://nonbinary.wiki/wiki/Misgendering>

| Pronoun | # Prompts | Nominative | Accusative | Genitive<br>(Attributive) | Genitive<br>(Predicative) | Reflexive  |
|---------|-----------|------------|------------|---------------------------|---------------------------|------------|
| She     | 480       | She        | Her        | Her                       | Hers                      | Herself    |
| He      | 480       | He         | Him        | His                       | His                       | Himself    |
| They    | 480       | They       | Them       | Their                     | Theirs                    | Themselves |
| Ey      | 480       | Ey         | Em         | Eir                       | Eirs                      | Emselves   |
| Xe      | 480       | Xe         | Xir        | Xir                       | Xirs                      | Xirself    |
| Fae     | 480       | Fae        | Faer       | Faer                      | Faers                     | Faerselves |

Table 3.1: Pronouns and pronoun types split across prompts

we create a set of prompts consisting of various antecedents and pronouns to measure this expected agreement – which we call *pronoun consistency* – in the model’s generated text. Pronouns measured included *she*, *he*, *they*, *xe*, *ey*, and *fae* (Table 3.1). An example prompt is the following:

*[Casey] is an author of children’s fantasy, best known for [their] book that won several awards.*

The antecedent is the name **[Casey]**, who uses the pronoun **[their]**. If this prompt were followed by text referring to Casey as *he*, this would be a case of misgendering. Assessing pronoun-antecedent agreement with *named antecedents* is one way to measure misgendering [DMO21]. However, sociolinguistic works have also investigated other methods of measuring pronoun inclusivity in the TGNC community. For example, socially distant subjects, rather than names, called a *distal antecedent*, can also be used to analyze differences in misgendering behavior [Bjo17]. In our example, we may then replace **[Casey]** with a distal antecedent such as **[The man down the street]** and measure changes in LLM misgendering.

**Curation Setup** To create the templates, we randomly sampled sentences from the Nonbinary Wiki. In order to rule out sentences with ambiguous or multiple antecedent references, we only proceeded with sentences that included an antecedent later, followed by a pronoun referring to that same antecedent. Sentences that began with the subject were collected and replaced with either a

name or a distal antecedent. Distal antecedents were handcrafted to reflect distant social contexts. Common distal forms include naming someone by occupation [Bjo17]. We only used occupations that do not reflect a particular gender (e.g., salesperson, cellist, auditor). For named antecedents, we gather gendered and nongendered popular names. We collected a sample of nongendered names from the Nonbinary Wiki and cross-referenced their popularity using [Flo15]. Common names stereotypically associated with binary genders (i.e., masculine names for a man, feminine names for a woman) were collected from the social security administration [Adm22].

Following our motivating example, we replace the pronoun **their** with other pronouns common to the TGNB community. Based on the Nonbinary Wiki and US Gender Census, we created prompts including singular they and neopronouns *xe*, *ey*, *fae* (TGNB pronouns). We also include *he* and *she* (binary pronouns) to experiment with how inclusive behavior may differ across these pronouns. Finally, we note that there are several variations of neopronouns. For example, *ey* can also take on the Spivak pronoun form, *e*<sup>10</sup>. However, in this study, we only focus on the more popularly used pronouns and their respective forms (i.e. nominative, accusative, genitive, reflexive), though it remains of future interest to expand this work with more pronoun variations (Table 3.1).

**Curation Results** We created 2,880 templates for misgendering evaluation and reported the breakdown in Table 3.2. Our dataset includes 480 prompts for each pronoun family of *she*, *he*, *they*, *xe*, *ey*, and *fae*. It also includes 720 prompts for each antecedent form, including distal antecedents and stereotypically masculine, feminine, and neutral names.

### 3.4 Experiments

In this section, we conduct OLG experiments that explore if and how models misgender individuals in text. First, we create templates detailed in § 3.3 for misgendering evaluation. Next,

---

<sup>10</sup>[https://nonbinary.miraheze.org/wiki/English\\_neutral\\_pronouns#E\\_\(Spivak\\_pronouns\)](https://nonbinary.miraheze.org/wiki/English_neutral_pronouns#E_(Spivak_pronouns))



| Antecedent Type    | # Prompts | Example Prompts  |
|--------------------|-----------|--|
| Nongendered Name   | 720       | Casey is an American actor and they are known for their roles in film.                   |
| Feminine Names     | 720       | Charlotte is a musician and dancer and they currently live somewhere nearby.             |
| Masculine Names    | 720       | James is a writer and they gained some fame for their debut novel.                       |
| Distal Antecedents | 720       | The cellist in my orchestra is a writer and they gained some fame for their debut novel. |

Table 3.2: Misgendering Prompt Set Statistics (N=2,400).

we propose an automatic metric to capture these instances and validate its utility with Amazon Mechanical Turk. Informed by sociolinguistic literature, we later ground further experiments in creating prompts to test how such gaps in pronoun consistency occur, analyze such results through both a technical and sociotechnical lens, and finish by providing constructive suggestions for future works.

### 3.4.0.1 Models Evaluated

We assess possible non-affirmation of TGNB identities across multiple large language models. Each model is triggered to generate text conditioned on prompts from one of our evaluation sets in TANGO. We describe the models in this paper below, with each size described in their respective experimental setup. We choose these models because they are open-source and allow our experiments to be reproducible. We also perform a case study with ChatGPT, with model details and results described in §3.4.0.5.

**GPT-2** Generative Pre-trained Transformer 2 (GPT-2) is a self-supervised transformer model with a decoder-only architecture. In particular, the model is trained with a causal modeling objective of predicting the next word given previous words on Webtext data, a dataset consisting of over 40GB of text [RWC19].

**GPT-Neo** GPT-Neo is an open-source alternative to GPT-3 that maintains a similar architecture

to GPT-2 [BGW21]. In a slightly modified approach, GPT-Neo uses local attention in every other layer for causal language modeling. The model was trained on the PILE dataset, consisting of over 800 GB of diverse text [GGB20].

**OPT** Open Pre-trained Transformer (OPT) is an open-source pre-trained large language model intended to replicate GPT-3 results with similar parameters size [ZRG22]. The multi-shot performance of OPT is comparable to GPT-3. Unlike GPT-2, it uses a BART decoder and is trained on a concatenated dataset of data used for training RoBERTa [LOG19], the PushShift.io Dataset [BZK20], and the PILE [GGB20]

### 3.4.0.2 Misgendering Measured by Automatic Tool and Human Evaluation

To assess LLMs for misgendering behavior in OLG, we create an automatic misgendering evaluation tool. Given a prompt with a referent and their pronoun (Figure 3.1), it measures how consistently a model uses correct pronouns for the referent in the generated text. We expect to find that models generate high-quality text which correctly uses a referent’s pronouns across binary, singular they, and neopronoun examples.

**Automatic Misgendering Evaluation** To automatically measure misgendering, one can compare the subject’s pronoun in the template to the subject’s pronoun provided in the model generation. To locate the subject’s pronoun in the model’s text generation, we initially tried coreference resolution tools from AllenNLP [Allnd] and HuggingFace [Hugnd]. However, coreference tools have been found to have bias with respect to TGNB pronouns often used by the community (e.g. singular they, neopronouns). They may be unable to consistently recall them to a subject in text [CD21]. We find this to be consistent in our evaluations of each tool and provide our assessment in [OGD23]. While ongoing work explores these challenges, we avoid this recall erasure with a simple yet effective tool. Given that the dataset contains only one set of pronouns per prompt, we measure the consistency between the subject’s pronoun in the provided prompt and the first pronoun observed in model generation. While the tool cannot be used with multiple referents, it is a good starting point for OLG misgendering assessments.

**Setup** We evaluate a random sample of 1200 generations for misgendering behavior across the 3 models. First, we run our automatic evaluation tool on all generations. Then we compare our results to human annotations via Amazon Mechanical Turk (AMT). Provided prompts, each model generation is assessed for pronoun consistency and text quality by 3 human annotators. We provide a rubric to annotators and ask them to rate generation coherence and relevance on a 5-point Likert scale [JKC15]. Next, we measure lexical diversity by measuring each text’s type-token ratio (TTR), where more varied vocabulary results in a higher TTR [Tem57]. A majority vote for pronoun consistency labels provides a final label. Then, we calculate Spearman’s rank correlation coefficient,  $\rho$ , between our automatic tool and AMT annotators to assess the correlation in misgendering measurements. We also use Krippendorff’s  $\alpha$  to assess inter-annotator agreement across the 3 annotators for text quality. Finally, we examine behavior across model sizes since the literature points to strong language capabilities even on small LLMs [SS20]. We report our findings on GPT-2 (125M), GPT-Neo (1.3B), and OPT (350M) and repeat evaluations across 3 approximate sizes for each model: 125M, 350M, 1.5B. Huggingface was used to generate the texts for GPT2, GPT-Neo, and OPT, generated 100 tokens with nucleus sampling.

To provide fair compensation, we based payout on 12 USD per hour and the average time taken, then set the payment for each annotation accordingly. There were 3 annotators per task, with 269 unique annotators in total. Since the task consists of English prompts and gender norms vary by location, we restrict the pool of workers to one geography, the United States. For consistent labeling quality, we only included annotators with a hit acceptance rate greater than 95%. To protect worker privacy, we refrain from collecting any demographic information.

While conducting AMT experiments with minimal user error is ideal, we do not expect annotators to have in-depth knowledge of TGNB pronouns. Instead, we first examine the user error in identifying pronoun consistency in a compensated AMT prescreening task consisting of a small batch of our pronoun consistency questions. Then we provide an educational task to decrease the error as best we can before running the full AMT experiment. After our educational task, we found

|         | Accuracy | Recall | Precision | F1    | Spearman $\rho$ ( $p_i$ 0.001) |
|---------|----------|--------|-----------|-------|--------------------------------|
| GPT-2   | 0.851    | 0.726  | 0.746     | 0.735 | 0.546                          |
| GPT-Neo | 0.888    | 0.796  | 0.670     | 0.716 | 0.558                          |
| OPT     | 0.945    | 1.000  | 0.908     | 0.951 | 0.837                          |

| Model   | Pronoun Consistency |       |       | Relevance   |      |      | Coherence   |      |      | Type-Token Ratio |       |              |
|---------|---------------------|-------|-------|-------------|------|------|-------------|------|------|------------------|-------|--------------|
|         | Binary              | They  | Neo   | Binary      | They | Neo  | Binary      | They | Neo  | Binary           | They  | Neo          |
| GPT-2   | <b>0.818</b>        | 0.460 | 0.101 | <b>3.73</b> | 3.38 | 3.40 | <b>4.00</b> | 3.60 | 3.83 | <b>0.761</b>     | 0.728 | 0.753        |
| GPT-Neo | <b>0.839</b>        | 0.365 | 0.166 | <b>4.11</b> | 3.88 | 3.54 | <b>4.14</b> | 4.04 | 3.75 | <b>0.693</b>     | 0.659 | 0.674        |
| OPT     | <b>0.937</b>        | 0.467 | 0.608 | <b>3.24</b> | 2.61 | 2.68 | <b>2.61</b> | 2.45 | 2.61 | 0.338            | 0.418 | <b>0.423</b> |

Table 3.3: Consistency metrics for the AMT experiments and automatic tool. Pronoun consistency, relevance, coherence, and type-token ratio are reported based on AMT experiments. Bold values indicate the highest score in each category per model.

that error rates for neopronoun<sup>11</sup> labeling decreased from 45% to 17%. We invited annotators who took the educational task in the initial screen to annotate the full task. We detail our educational task in the appendix [OGD23].

**Results** We discuss our AMT evaluation results and pronoun evaluation alignment with our automatic tool in Table 3.3. We observe a moderately strong correlation between our automatic metric and AMT across GPT-2, GPT-Neo, and OPT ( $\rho = 0.55, 0.56, 0.84$ , respectively). Across all models, we found pronouns most consistently generated when a referent used binary pronouns. We observed a substantial drop in pronoun consistency across most models when referent prompts used singular they. Drops were even more substantial when referent prompts took on neopronouns. OPT misgendered referents using TGNB pronouns (e.g., singular they, neopronouns) the least overall, though, upon further examination, multiple instances of its generated text consisted of the initial prompt. Therefore, we additionally reported text generation quality following this analysis. After OPT, GPT-Neo misgendered referents with neopronouns the least, though GPT-2 reflected the highest pronoun consistency for TGNB pronouns overall (Binary: 0.82, They: 0.46,

<sup>11</sup>Moving forward, we use *neo* as a reporting shorthand.




Figure 3.2: Distribution of pronoun consistency (left) and perplexity (right) across 9 models. Templates with binary pronouns consistently result in the least misgendering across model sizes. Neo: 0.10, Mann-Whitney p-value < 0.001).

We observed a moderate level of inter-annotator agreement ( $\alpha=0.53$ ). All models’ relevance and coherence were highest in generated text prompted by referents with binary pronouns (Relevance: Binary Pronoun Means GPT-2: 3.7, GPT-Neo: 4.1, OPT: 3.2, Kruskal Wallis p-value < 0.001. Coherence: Binary Pronoun Means GPT-2: 4.0, GPT-Neo: 4.1, OPT: 2.6, Kruskal Wallis p-value < 0.001). Across most models, lexical diversity was highest in generated text prompted by referents with binary pronouns as well (Binary Pronoun GPT-2: 0.76, GPT-Neo: 0.69, OPT:0.34, Kruskal Wallis p-value < 0.001). Upon observing OPT’s repetitive text, its low relevance and coherence validate the ability to capture when this may occur.

To better understand the prevalence of misgendering, we further evaluated each model across modeling capacity using perplexity measurements and our automatic misgendering evaluation tool. Notably, we observed results similar to our initial findings across model sizes; binary pronouns resulted in the highest pronoun consistency, followed by singular they pronouns and neopronouns (Figure 3.2). For perplexity, we observed that models resulted in the least perplexity when prompted with binary pronouns. Meanwhile, neopronouns reflected a much higher average perplexity with a more considerable variance. These results may indicate that the models, regardless of capacity, still struggle to make sense of TGNB pronouns. Such inconsistencies may indicate upstream data availability challenges even with significant model capacity.

### 3.4.0.3 Analysis of Antecedent Forms

We draw from linguistics literature to further investigate misgendering behavior in OLG. [Bjo17, SF07] assess the perceived acceptability of gender-neutral pronouns in humans by measuring readability. They assess the “acceptability” of singular they by measuring the time it takes humans to read sentences containing the pronoun across various antecedents. These include names and “distal antecedents” (i.e., referents marked as less socially intimate or familiar than a name). The less time it takes to read, the more “accepted” the pronoun is perceived. Researchers found that subjects “accepted” singular they pronouns *more* when used with distal antecedents rather than names. We translate this to our work, asking if this behavior is reflected in OLG. We expect that LLMs robustly use correct pronouns across both antecedent forms.

**Setup** To measure differences in model behavior, we report 2 measures across the following models: GPT-2 (355M), GPT-Neo (350M), and OPT (350M). We use our automatic misgendering metric to report pronoun consistency differences between distal and nongendered name antecedents across binary, singular they, and neopronouns. Similar to measuring the “acceptability” of pronouns in human subjects, since perplexity is a common measure of model uncertainty for a given text sample, we also use perplexity as a proxy for how well a model “accepts” pronouns across various antecedents. In our reporting below, we describe “TGNB pronouns” as the aggregation of both singular they and neopronouns.

**Results** As shown in Table 3.4, across all models, misgendering was least observed for singular they pronouns in prompts containing distal antecedents (difference of means for distal binary vs. TGNB pronouns GPT2: 0.46, GPT-Neo: 0.56, OPT: 0.69, Kruskal-Wallis p-value < 0.001). These results aligned with human subjects from our motivating study [Bjo17]. Besides GPT-2, neopronoun usage seemed to follow a similar pattern. Regarding perplexity, we also found that all models were less perplexed when using distal antecedents across all pronouns. Notably, drops in perplexity when using distal antecedent forms were more pronounced for TGNB pronouns (binary - TGNB pronoun  $|\Delta|$  across antecedents GPT: 78.7, GPT-Neo:145.6, OPT:88.4 Mann-Whitney

Table 3.4: Differences in misgendering and perplexity across antecedents with varying social contexts.  $\Delta$  reflects the absolute difference between Named and Distal antecedent forms.

| Metric                      | Pronoun | GPT2         |                |            | GPT-Neo      |                |            | OPT          |                |            |
|-----------------------------|---------|--------------|----------------|------------|--------------|----------------|------------|--------------|----------------|------------|
|                             |         | Named        | Distal         | $ \Delta $ | Named        | Distal         | $ \Delta $ | Named        | Distal         | $ \Delta $ |
| Consistency ( $\uparrow$ )  | Binary  | <b>0.923</b> | 0.898          | 0.025      | <b>0.986</b> | 0.739          | 0.247      | <b>0.891</b> | 0.882          | 0.009      |
|                             | They    | 0.333        | <b>0.345</b>   | 0.012      | 0.321        | <b>0.458</b>   | 0.137      | 0.222        | <b>0.667</b>   | 0.445      |
|                             | Neo     | <b>0.067</b> | 0.017          | 0.05       | 0.114        | <b>0.152</b>   | 0.038      | 0.333        | <b>0.667</b>   | 0.334      |
| Perplexity ( $\downarrow$ ) | Binary  | 120.775      | <b>110.357</b> | 10.418     | 144.295      | <b>114.204</b> | 30.091     | 120.024      | <b>92.118</b>  | 27.906     |
|                             | They    | 149.449      | <b>130.025</b> | 19.424     | 171.961      | <b>131.877</b> | 40.084     | 147.335      | <b>104.599</b> | 42.736     |
|                             | Neo     | 486.563      | <b>328.55</b>  | 158.013    | 446.706      | <b>323.61</b>  | 123.096    | 310.888      | <b>207.719</b> | 103.169    |

p-value  $< 0.001$ ). Based on these results, the “acceptability” of TGNB pronouns in distal -rather than named- antecedents seems to be reflected in model behavior.

It is important to ground these findings in a social context. First seen around the 1300s [Dicnd], it is common to refer to someone socially unfamiliar as “they” in English. We seem to observe this phenomenon reflected in model performances. However, singular they is one of the most used pronouns in the TGNB population, with 76% of TGNB individuals favoring this in the 2022 Gender Census [Cennd]. These results indicate that individuals who use such pronouns may be more likely to experience misgendering when referred to by their name versus someone of an unfamiliar social context. Meanwhile, referents with binary pronouns robustly maintain high pronoun consistency across antecedent forms. These results demonstrate perpetuated forms of gender non-affirmation and the erasure of TGNB identities by propagating the dominance of binary gender.

#### 3.4.0.4 Analysis of Observed Pronoun Deviations

Provided the observed differences in misgendering from the last section, we explore possible ways pronoun usage across models differs and if such behaviors relate to existing societal biases. In line with linguistics literature, we hypothesize that pronouns in generations will exhibit qualities following (1) a preference for binary pronouns and (2), within binary pronouns, a preference for “generic masculine” (i.e., the default assumption that a subject is a man) [Sil80]. This means that

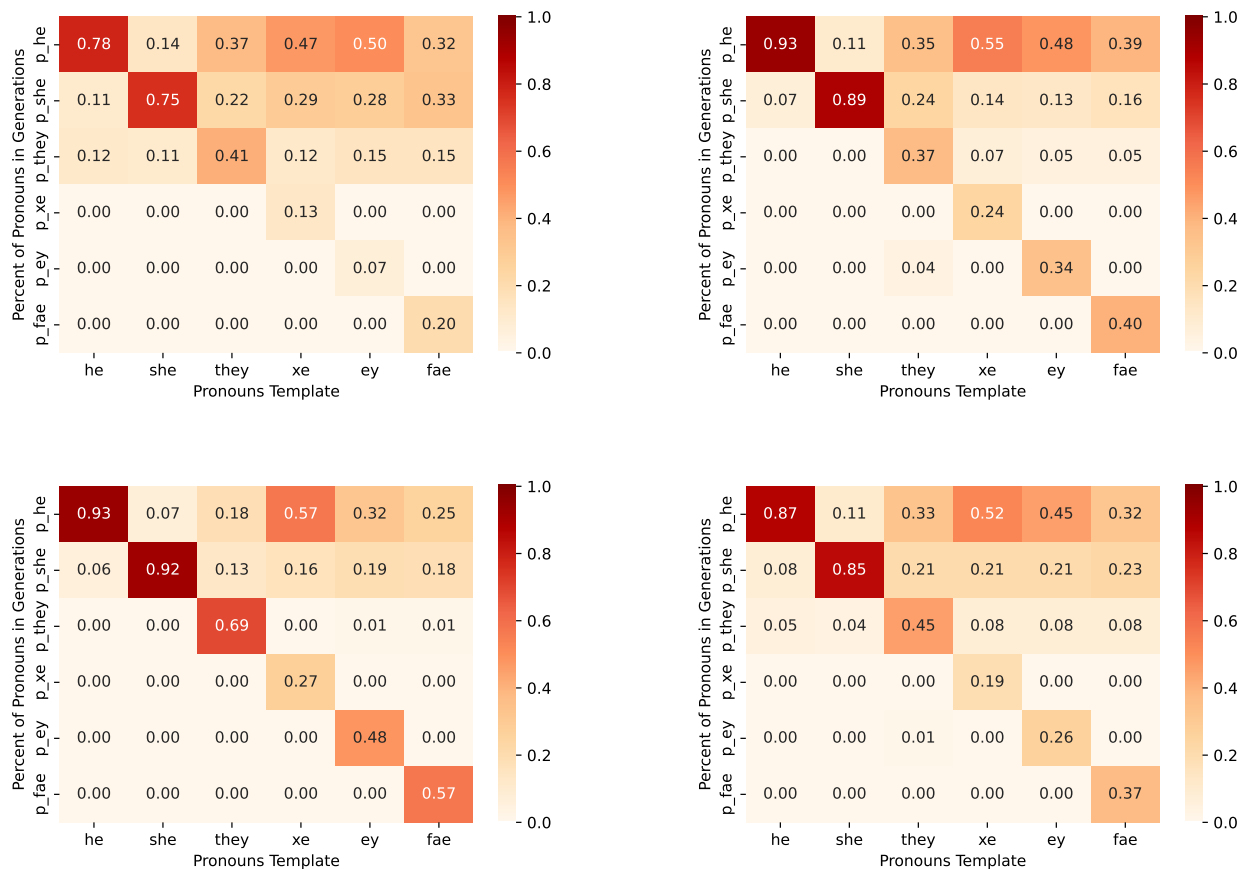


Figure 3.3: Pronoun Template vs Pronouns in Generations. From left to right: GPT2, GPT-Neo, OPT, All

we will observe models deviating more towards using he pronouns. We also wonder to what extent models understand neopronouns as their corresponding part of speech and if this deviates more towards noun-hood.

**Setup** To examine LLM misgendering more closely, we report 2 measures. First, we look at the distribution of pronouns generated by all the models across the pronoun templates. Then, we assess for correct usage of the pronouns by splitting each generated pronoun by its pronoun type, either nominative, accusative, genitive, or reflective. Regarding pronouns, determiners such as “a” and “the” usually cannot be used before a pronoun [Camnd]. Therefore, we use this to measure when






Figure 3.4: Pronoun Template Type vs Errors in Generations. From left to right: GPT2, GPT-Neo, OPT, All

the model does not correctly generate pronouns.

**Results** Across all models, LLM generations leaned towards incorporating binary pronouns, regardless of the prompt’s pronoun (difference of proportions in binary - TGNB pronouns GPT-2: 0.53, GPT-Neo: 0.52, OPT: 0.47 Kruskal Wallis p-value < 0.001). Prompts with TGNB pronouns were most susceptible to this shift. Prompts with referents using *xe* resulted in generations mainly containing *he* pronouns (Figure 3.3). Further examining binary pronoun generation, we observed that, on average, generations from *he* pronoun templates were the most frequent across all templates aside from *she*, regardless of model (GPT-2: 0.42, GPT-Neo: 0.44, OPT: 0.39, Kruskal Wallis p-value < 0.01). Our findings also corroborate linguistics literature on “generic masculine” defaults.

When assessing pronounhood through analysis of the post-determiner environment, we find that neopronouns are more likely to be misused. 43% of fae pronouns in the nominative usage start with “the fae” or “a fae” (Figure 3.4). Meanwhile, we did not see this behavior with prompts consisting of binary and singular they pronouns. These results may hint at the possible gaps in lexical understanding: LLMs may not understand neopronouns as pronouns but possibly other parts of speech.

### 3.4.0.5 Case Study: ChatGPT

We apply our auditing methods to the recent release of ChatGPT [Ope23] to demonstrate the evaluation of TGNB harms.

**Setup** We test the model for misgendering, analyze the distribution of pronoun usage, error rates, and diversity (i.e. reflected in different pronoun forms), and compare results to previously evaluated models. We use the OpenAI ChatGPT API and adapt our prompts to reflect the dialogue setting better. We show an example below:

*Please generate 100 tokens using the following prompt: casey is an american actor and he is known for his roles in film.*

**Results** Compared to previously evaluated models, ChatGPT reflected the most robust pronoun consistency across pronoun groups (Binary: 0.96, They: 0.95, Neo: 0.74, Kruskal Wallis p-value  $< 0.01$ ). However, similar to previous evaluations, we note the pattern of pronoun consistency decreasing for singular they and even more so for neopronouns. Furthermore, unlike all other models, ChatGPT made almost no determiner errors when prompted with neopronouns. Additionally, generations with TGNB pronoun templates resulted in more TGNB pronoun generation than previous models (Figure 3.5). We also observed that out of all models, ChatGPT generated sentences with the most pronoun case heterogeneity when given TGNB pronoun forms, demonstrating a better capacity for producing various forms of these pronouns. While these evaluations indicate better gender inclusivity, future work is warranted to explore why this model does significantly better than others, along with areas of weakness.

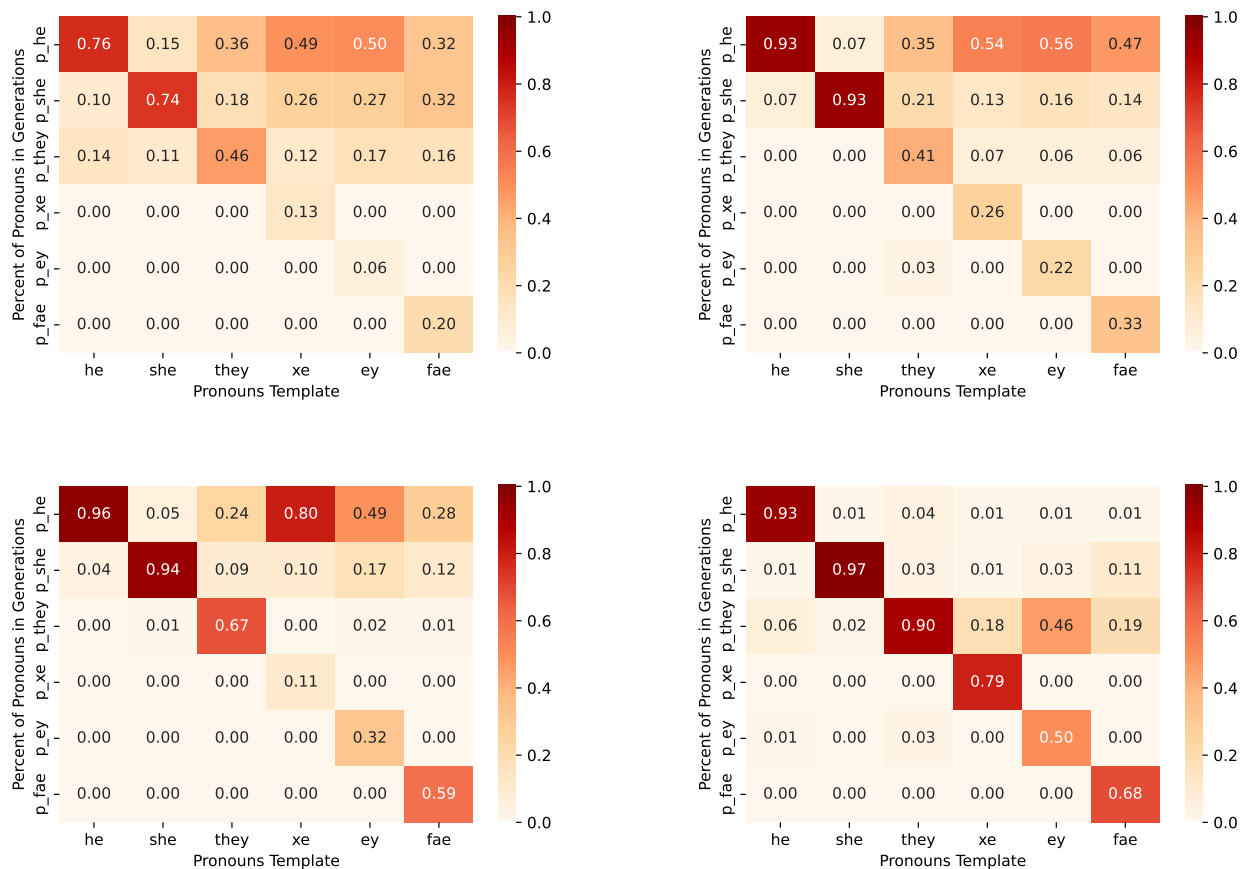


Figure 3.5: Pronouns generated using respective pronoun template types when using only non-binary names or distal antecedents. From left to right: GPT2, GPT-Neo, OPT, ChatGPT

### 3.5 Discussion

In this work, we introduced new evaluations for studying gender-diverse bias within causal language models and found that LLMs perpetuate the binary construction of gender in English. We discovered that misgendering of TGNB pronouns is the norm in popular LLMs; GPT-2, GPT-Neo, OPT, and ChatGPT misgendered subjects the least using binary pronouns but misgendered the most when subjects used neopronouns. Compared to binary pronouns, TGNB pronouns are significantly less consistent with pronoun-antecedent agreement across GPT-2, GPT-Neo, OPT, and ChatGPT.

The generated text also seems to follow generic masculine via favoring binary-masculine pronoun usage. Because of this, we recommend a few approaches for future study. First, pretraining the model with a more diverse corpus containing more examples of named referents using singular pronouns and neopronouns is worth exploring. Training a tokenizer with explicit merging rules may also be helpful to preserve the valuable morphosyntactic structure and meaning of neopronouns. Finally, in-context learning [LSZ21, DLD22, DSD22] with various TGNB pronoun examples may also effectively mitigate these harms.

## CHAPTER 4

# Mapping Misrepresentation: Understanding Representational Skew in Language Models

Chapter 3 revealed how language models can systematically erase and misrepresent transgender and non-binary persons. Building on these findings, this chapter examines how specific elements of the LLM training pipeline can contribute to these discriminatory outcomes. With this, we aim to provide a detailed understanding of how existing language modeling techniques inadvertently perpetuate unwanted societal biases, shedding light on future pathways towards more inclusive language models. This chapter is based on the previously published work [DMO21].

### 4.1 Introduction

Natural Language Processing (NLP) has achieved remarkable advances through large language models [BMR20, DCL19a], yet these systems continue to exhibit persistent biases, particularly in their handling of gender [BB19, SA21]. As AI-driven language technologies become more pervasive in applications from conversational agents [LDT18] to content moderation systems [KSC24]—their potential to amplify societal biases and disproportionately impact marginalized communities increases [BGM21]. The development of more inclusive language models is therefore crucial for ensuring equitable access to AI technologies [HPD20a] and preventing the perpetuation of harmful stereotypes in our increasingly digital society [MMS21].

The impacts of these biases are particularly noticeable for transgender and non-binary individuals, who experience unique forms of algorithmic harm that go beyond traditional binary gender bias concerns. While gender bias in NLP has been extensively studied, these investigations have

largely focused on male-female dichotomies in contexts like occupation [ZWY19] or stereotype attribution [BCZ16]. The distinct challenges faced by non-binary individuals—including systematic erasure, misgendering, and active misrepresentation—require a fundamentally different analytical framework that considers the complexities of gender identity beyond binary classifications [DMO21].

Substantial research has examined binary gender biases in language models [SGT19, ZWY19, BCZ16], yet the technical mechanisms behind non-binary gender bias remain critically understudied [DMO21]. Building on evidence of systemic misgendering found in the previous chapter, in this chapter we provide a detailed analysis of how gender-diverse erasure manifests within language models’ training pipeline. This analysis reveals how transgender and non-binary individuals (TGNB) face systematic erasure at multiple technical levels—from data collection to model representation—leading to discriminatory practices in deployed language technologies [Cra17, BBD20]. Our analysis of TGNB representational harms focuses on three key areas of language modeling:

- (1) Data skew in large-scale text corpora
- (2) Representational erasure in static word embeddings
- (3) Biased associations in language representations

We employ a multi-faceted approach to study how representational harms for non-binary persons propagate through language technologies, examining model and data artefacts across the development pipeline. We first quantify frequency distributions of gender-diverse pronouns and terminology in English Wikipedia. Using GloVe embeddings trained on this data, we then examine embedding neighborhood structures and conduct Principal Component Analysis (PCA) to understand gender subspace characteristics. Finally, we apply the Word Embedding Association Test (WEAT) to measure biased associations between gender-related terms and sentiment attributes.

Our analysis reveals significant disparities in representation: neopronouns appear in less than  $< 1\%$  of contexts compared to binary pronouns in Wikipedia, while gender-diverse terminology shows similar underrepresentation. This data scarcity manifests in embedding spaces through unstable gender subspaces and poor semantic neighborhoods for TGNB terms. WEAT analysis further demonstrates systematic biases, with TGNB terminology showing significant negative sentiment associations compared to binary gender terms. These empirical findings quantify how representational limitations cascade through the NLP pipeline, suggesting the need for interventions at multiple stages of model development.

## 4.2 Related Works

NLP research has explored several avenues for quantifying social biases [SCN21, CBN17a, RNL18, DRW19a] and mitigating them [ZWY19, REG20, SGT19]. Prior work on gender bias in NLP has analyzed bias in word embeddings [BCZ16, CBN17a] and documented systematic gender disparities in model training data [ZWY17a, ZWY19]. Mitigations have typically centered variants of orthogonal subspace correction of gender stereotypes [DLP20] and corpus-level constraints [ZWY17a]. However, these studies primarily focused on binary gender categories, leaving open questions about how to effectively analyze and quantify bias against non-binary genders in language systems. Importantly, the biases faced by non-binary persons can be distinct from the problem statements grounded in a binary manner.

Works have begun examining non-binary gender representation, though research has been limited in scope. Language representations taking the form of static word embeddings (GloVe [PSM14]) and contextual (e.g., BERT [DCL19b]) embeddings are common ways models ingest information for downstream modeling. [ZZL18] proposed a gender-neutral variant of static embeddings. [CD19] introduced GICoref to evaluate non-binary and transgender pronoun handling, and [SGT21] explored gender-neutral text generation through pronoun substitution. However, a key gap remains in understanding how representational erasure and harmful associations compound throughout the NLP pipeline for non-binary genders. This chapter addresses this gap by providing the first com-

prehensive analysis of how non-binary gender biases manifest across multiple components: from data skew in training corpora, through representational issues in word embeddings, to biased associations in final models. Our work also demonstrates why traditional bias analysis methods developed for binary gender cannot be simply extended to non-binary genders, highlighting the need for new frameworks that account for the unique challenges of representing gender beyond a binary framework.

## 4.3 Methodology

This section outlines data-driven artefacts studied which are crucial to training a language model, the data and its representations. We first outline how we assess datasets for skew, followed by downstream representations and the implications they carry for imbuing social context.

### 4.3.0.1 Data skew

The quality of any AI-driven model strictly depends on the data it was trained on. This is no different for language models, where several banks of pretraining corpora exist for learning language syntax and semantics. Wikipedia one of the most common pretraining corporas. Being that biases can present themselves as a result of skews in the data, we study how skewed Wikipedia is in relation to the prevalence and gendered terms, across the spectrum. We also use the Python library *wordfreq*<sup>1</sup> which samples over diverse data to give an approximate usage of different words in all of the text curated from the web, to observe how vastly different the frequencies of different gendered words are per billion words in English [SCL18].

### 4.3.0.2 Measuring Representational Erasure on Embeddings

Text representations have been known to learn and exacerbate skewed associations and social biases from underlying data [ZWY17b, BGM21, Dev20], thus propagating representational harm. We examine representational skews with respect to pronouns and non-binary-associated words

---

<sup>1</sup><https://pypi.org/project/wordfreq/>



that are extremely sparsely present in text.

### 4.3.0.3 GLoVe

To investigate how non-binary genders are represented—or misrepresented—in language models, we begin by analyzing static word embeddings, specifically GloVe [PSM14], as a proxy for understanding deeper biases that also influence modern language models like GPT [RNS18] and BERT [DCL19b]. GloVe, trained on large corpora like Wikipedia, offers a window into how these gender-related biases manifest in the underlying data used to train language models. While GloVe itself is a static model and lacks the dynamic, context-sensitive capabilities of LLMs, it provides a clear and measurable baseline for examining the skewed representation of non-binary pronouns and gender-related terms. Given that Wikipedia and other similar text sources are integral to the training of LLMs, the misrepresentation observed in GloVe embeddings can act as a predictor of how LLMs might also struggle with non-binary gender representation. This approach allows us to evaluate the foundational biases encoded in language representations, providing insight into why non-binary terms are often misrepresented or erased in modern NLP systems. Glove was trained on English Wikipedia<sup>2</sup> articles with a window size of 15, a dimension of 50 for each word, and a minimum word frequency of 5.

### 4.3.0.4 Measuring Biased Associations and Sentiment

| Set        | Words   |
|------------|---|
| pleasant   | <i>joy, love, peace, wonderful, pleasure, friend, laughter, happy</i> |
| unpleasant | <i>agony, terrible, horrible, nasty, evil, war, awful, failure</i>    |

Table 4.1: Set of unpleasant and pleasant words

Gender bias literature primarily focuses on stereotypically gendered occupations [BCZ16, DRW19b], with some exploration of associations of binary gender and adjectives [DP19, CBN17b].

---

<sup>2</sup><https://dumps.wikimedia.org/>

| <b>Set</b>         | <b>Words</b>   |
|--------------------|--|
| binary pronouns    | <i>he, him, his, she, her, hers</i>  |
| binary words       | <i>man, woman, herself, himself, girl, boy, female, male, cisman*, ciswoman*</i>         |
| binary all         | <i>binary pronouns + binary words</i>  |
| nonbinary pronouns | <i>zey, ey, em, them, xir, they, zem, ze, their, zir, zers, eirs, xey, xers, xe, xem</i> |
| nonbinary words    | <i>transgender, queer, nonbinary, genderqueer, genderfluid, bigender, two-spirit</i>     |
| nonbinary all      | <i>nonbinary pronouns + nonbinary words</i>  |

Table 4.2: Word set definitions for binary and non-binary concepts

While these associations are problematic, there are additional, significantly different biases against non-binary genders, namely misrepresentation and under-representation. To understand this, we conduct a set of biased association tests to observe to what extent the representation of non-binary gender is robustly embedded in data-driven language artefacts. To do this, we conduct an analysis of representational skews present in existing embedding for a wide range of gender pronouns. Then we perform sentiment associations between binary and non-binary associated words, followed by benchmarking against classic gender binary stereotype assessments.

We perform a nearest neighbor analysis to understand representation skew between binary and non-binary gendered words. Since we observed a skew in representation of pronouns earlier, we also include proxy words that reflect nonbinary and trans representation in order to further understand biased associations. To investigate sentiment associations with binary versus non-binary associated words, we use the WEAT test [CBN17b] with respect to pleasant and unpleasant attributes. We provide WEAT scores for 3 different sets of words, related to both binary and non-binary gender in Table 4.1. Since neopronouns are not well-embedded, we compare disparate sentiment associations between binary versus non-binary pronouns, gendered words and proxies (e.g., *male, female* versus *transman, genderqueer*, etc.). The full list can be found in Table 4.2.

| Word    | Frequency (%) | Word        | Frequency (%) |
|---------|---------------|-------------|---------------|
| he      | 0.49000       | man         | 0.06610       |
| his     | 0.32400       | girl        | 0.02400       |
| they    | 0.31600       | woman       | 0.02240       |
| she     | 0.18200       | boy         | 0.01480       |
| them    | 0.15500       | female      | 0.01000       |
| himself | 0.01780       | male        | 0.00776       |
| herself | 0.00603       | two-spirit  | 0.00588       |
| hers    | 0.00093       | em          | 0.00372       |
| ey      | 0.00019       | transgender | 0.00081       |
| ze      | 0.00012       | queer       | 0.00057       |
| xe      | 0.00005       | nonbinary   | 0.00002       |
| zem     | 0.00001       | cisgender   | 0.00002       |
| xem     | 0.00000       | genderqueer | 0.00001       |
| zey     | 0.00000       | genderfluid | 0.00001       |
| zir     | 0.00000       | bigender    | 0.00000       |
| xir     | 0.00000       | cisman      | 0.00000       |
| xey     | 0.00000       | ciswoman    | 0.00000       |

Table 4.3: Frequency of Gender-related pronouns (left) and terms (right) for English Wikipedia, reported per billion. Frequencies reflect skew towards binary gender-related content.

## 4.4 Results

### 4.4.0.1 Data skew

Large text dumps often used to build language representations have severe skews with respect to gender and gender-related concepts. As we anticipated, the distribution of different pronouns is not equal across genders. Overall, while ‘he’ and ‘she’ occur 0.49% and 0.316% per billion words respectively, the percent for ‘xe’ and ‘ze’ is only 0.0005% and 0.0011% respectively (4.3) Just observing pronoun usage, English Wikipedia text (March 2021 dump), which comprises 4.5 billion tokens, has over 15 million mentions of the word *he*, 4.8 million of *she*, 4.9 million of *they*, 4.5 thousand of *xe*, 7.4 thousand of *ze*, and 2.9 thousand of *ey*. Furthermore, the usages of non-binary

| Pronoun | Sentence   |
|---------|--|
| Ey      | "The difference in the alphabets comes only in the Faroese diphthongs (ei being 26, ey 356, oy 24...)."  |
| Em      | Approximating the em dash with two or three hyphens.   |
| Xem     | "'Em di xem hoi trang ram'", establishing her icon for Vietnamese women as well as earning the title of the "'Queen of Folk'"  |
| Ze      | "He taught himself to write with his left hand and described his experiences before, during, and after the accident in a deeply moving journal, later published under the title 'Pogodzic sie ze swiatem' ('To Come to Terms with the World').", |
| Zir     | "The largest operation in the Struma Valley was the capture by 28th Division of Karajakoi Bala, Karajakoi Zir and Yenikoi in October 1916."  |

Table 4.4: Example sentences containing nonbinary pronouns

pronouns<sup>3</sup> were mostly not meaningful with respect to gender. *Xe*, as we found by annotation and its representation, is primarily used as the organization *Xe* rather than the pronoun *xe*. *Ze* was primarily used as the Polish word *that*, as indicated by its proximity to mostly Polish words like *nie*, i.e. *no*, in the GloVe representations of the words, and was also used for characterizing syllables. Additionally, even though the word *they* occurs comparably in number to the word *she*, a large fraction of the occurrences of *they* is as the plural pronoun, rather than the singular, non-binary pronoun *they*. To illustrate this discrepancy in usage, we manually annotated 150 random samples each of the pronouns *he*, *she*, *they*. Only 1 mention of *they* was done in a non-binary, singular pronoun form, where all mentions of *he* and *she* carried gendered connotation. As a consequence of historical discrimination and erasure in society, narratives of non-binary persons are either largely missing from recorded text or have negative connotations. Subsequently, biases and harms due to tainted examples, limited features, and sample size disparities are exacerbated in language technologies.

Table 4.4 contains a random sample of sentences demonstrating how some neopronouns were used. This demonstrates that usage of non-binary pronouns in text is not always meaningful with respect to gender. Similar skews are observed in other big datasets such as Gigaword[NGV12],

---

<sup>3</sup>Neopronouns and gendered pronouns not "he" or "she"

| <b>Pronoun</b> | <b>Top 5 Neighbors</b>                 |
|----------------|--|
| He             | <i>his, man, himself, went, him</i>    |
| She            | <i>her, woman, herself, hers, life</i> |
| They           | <i>their, them, but, while, being</i>  |
| Xe             | <i>xa, gtx, xf, tl, py</i>             |
| Ze             | <i>ya, gan, zo, lvovic, kan</i>        |

| <b>Pronoun</b> | <b>Top 5 Neighbors</b>                                       |
|----------------|--|
| His            | <i>he, him, who, after, himself</i>                          |
| Hers           | <i>somehow, herself, thinks, someone, feels</i>              |
| Theirs         | <i>weren, tempted, couldn, gotten, willingly</i>             |
| Xers           | <i>yogad, doswelliids, hlx, cannibalize, probactrosaurus</i> |
| Zers           | <i>ditti, bocook, kurikkal, felimy, hifter</i>               |
| Eirs           | <i>cheor, yha, mnetha, scalier, paynet</i>                   |

(a) Nearest neighbor words in GloVe for binary and non-binary pronouns.

(b) Five nearest neighbors for binary and non-binary possessive pronouns.

Table 4.5: Nearest neighbor words in GloVe for binary and non-binary pronouns and their possessive forms.

which are sparsely populated with non-binary pronouns. Overall, we also observe skews in frequencies of non-binary terms in the English language as well, as reflected in Table 4.3. Some corpora do exist such as the Non-Binary Wiki<sup>4</sup> which contain instances of meaningfully used non-binary pronouns. However, with manual evaluation, we see that they have two drawbacks: (i) the narratives are mostly short biographies and lack the diversity of sentence structures as seen in the rest of Wikipedia, and (ii) they have the propensity to be dominated by Western cultures, resulting in further diminution of diverse narratives of non-binary persons.

#### 4.4.0.2 Static Embeddings assessment

<sup>4</sup>[https://nonbinary.wiki/wiki/Main\\_Page](https://nonbinary.wiki/wiki/Main_Page)

**Representational erasure in GloVe.** Table 4.5 shows the nearest neighbors of different pronouns in their GloVe representations trained on English Wikipedia data. The representations for binary-gendered pronouns in all tenses and grammatical constructions are meaningful and in agreement with how the words themselves are semantically used. The singular pronouns *he* and *she* have semantically meaningful neighbors as do their possessive forms (Right of Table 4.5). However, the same is not true for non-binary neopronouns in general, where *xe* and *ze* are closest to acronyms and Polish words, respectively. These reflect the disparities in occurrences we observe in data skews and hence lack of meaningful encodings of non-binary-associated words. For the pronoun ‘they’, since GloVe has a single representation for the word, the two distinct usages of it are not easily discernable. The pronoun had some singular occurrences in the Wikipedia text, though its static representation by GloVe is dominated by the more popular plural occurrence, as demonstrated by the nearest neighbors. These observations speak to their lack of broader representation, where non-binary genders are significantly underrepresented in textual data, causing language models to learn meaningless, unstable representations for non-binary-associated pronouns and terms. Moving forward in the next section, we use the term *bias* to refer to a skewed and undesirable association in language representations which has the potential to cause representational or allocational harms [BCS17].

**Biased associations in GloVe.** Skews as seen in GloVe representations are seen here with respect to nearest neighbors in Table 4.5 and often even with derogatory associations reflecting social biases (Table 4.6). We can see that there are relatively more negative adjectives associated with proxy nonbinary and trans words. These associations are directly a result of the skew in representation in the text, which downstream, can result in incredibly biased results. For *man* and *woman*, the top nearest neighbors include *good*, *great* and *good*, *loving*, respectively. However, for *transman* and *transwoman*, top words include *dishonest*, *careless* and *unkind*, *arrogant*. This further substantiates the presence of biased negative associations, as seen in the WEAT test. Furthermore, the nearest neighbors of words associated with non-binary genders are derogatory (see Table 4.6). In particular, *agender* and *gender fluid* have the neighbor *negrito*, meaning “little Black”, while

| Term        | 10 Nearest Neighbors  |
|-------------|---|
| agender     | bigender, genderfluid, genderqueer, tosin, cisgender, nonbinary, laia, muhafazat, <b>negrito</b> , farmgirl                 |
| bigender    | pangender, agender, genderfluid, overcontact, pnong, genderqueer, nonbinary, eczemas, gegs                                  |
| queer       | lesbian, lgbtq, feminism, lgbt, lesbians, feminist, racism, sexuality, stereotypes, gay                                     |
| nonbinary   | genderqueer, <b>transsexual</b> , cisgender, transsexuals, bigender, genderfluid, chorti, referents, pansexual, hitchhikers |
| transgender | lesbian, lgbt, lgbtq, bisexual, intersex, gender, <b>transsexual</b> , lesbians, heterosexual, discrimination               |
| genderfluid | agender, bigender, genderqueer, transwoman, nonbinary, pansexual, montserratian, <b>negrito</b> , supercouple, <b>fasiq</b> |
| genderqueer | pansexual, nonbinary, lgbtqia, <b>transsexual</b> , genderfluid, agender, bisexuality, bigender, diasporic, multiracial     |

Table 4.6: Ten Nearest neighbors of non-binary terms highlighting derogatory Terms

*gender fluid* has *Fasiq*, which is an Arabic word used for someone of corrupt moral character.

Following the usage of words pleasant and unpleasant [CBN17b], we find that non-binary genders suffer from a sentiment (positive versus negative) bias. In Table 4.7, we see the disparity with the WEAT score  $>0$  in both cases, denoting a higher association of pleasant words with binary-gendered words/pronouns as compared to non-binary words/pronouns. The WEAT score is 0.916, which is non-zero, i.e. ideal, significantly large and indicates disparate sentiment associations between the two groups. The data that representations are trained on are usually large text dumps taken from the web. While large enough to ably represent word usage based on the different word distributions, they are also skewed enough to exhibit unwanted biases. This skew comes from the social biases we see in historically lopsided data, such as that of specific occupations predominantly being done by specific genders. Gender-occupation associations were not a dominant stereotype observed across all genders (Table 4.8), where non-binary words like *transman* and *nonbinary* are not dominantly associated with either stereotypically male or female occupations. In fact, most occupations exhibit no strong correlation with words and pronouns associated with non-binary genders (see Table 4.9).

| Sets  | Weat Score |
|---|------------|
| Random Vectors  | -0.02      |
| Binary Pronouns vs. Non-Binary Pronouns                   | 0.2        |
| Binary Words vs. Non-Binary Proxies                       | 0.718      |
| Binary Pronouns + Words vs. Non-Binary Pronouns + Proxies | 0.916      |

Table 4.7: WEAT Scores (vs. pleasant and unpleasant attributes)

To better understand these observed gaps, we center exploring the subspaces captured when employing classic gender subspace analyses. Capturing a gender subspace has been useful in techniques of bias analysis and techniques in subsequent debiasing in binary gender [DP19], especially in context-free or static representations like GloVe or word2vec. These methods postulate expanding this to nonbinary gender by determining a general subspace for gender which captures both binary and non-binary genders. We test if we can approach capturing the all-gender subspace by extending one such general subspace capturing method [BCZ16] - principal component analysis (PCA) - on the two groups of words below, in addition to their combination:

If we truly captured the gender subspace, we could safely assume that the difference between the binary subspace and the all-gender subspace, along with the non-binary subspace and the all-gender subspace, is somewhat negligible. We make the following observations leveraging the cosine distance, defined as  $1 - c$ , where  $c$  is the cosine similarity between two vectors. We observe, opposite to what we expected, that the distance was quite different in these respective pairs. Between the binary and all-gender subspace was a cosine distance of 1.48, while the distance between the non-binary and all-gender subspace was larger, at 1.93. This tells us that the binary subspace is much less dissimilar than the nonbinary subspace with respect to the all-gender subspace, i.e., extending the approach of subspace capture to all genders would result in a subspace more dominantly aligned with binary gender than non-binary gender. Further, due to the poor representation



| <b>Pronoun</b> | <b>Average Similarity</b> |
|----------------|---------------------------|
| he             | 0.509                     |
| she            | 0.495                     |
| they           | 0.395                     |
| em             | 0.185                     |
| ze             | 0.123                     |
| ey             | 0.086                     |
| xe             | -0.054                    |
| zey            | -0.056                    |

Table 4.8: Average cosine similarity between occupations and nominative pronouns.

| <b>Word</b> | <b>Doctor</b> | <b>Engineer</b> | <b>Nurse</b> | <b>Stylist</b> |
|-------------|---------------|-----------------|--------------|----------------|
| man         | 0.809         | 0.551           | 0.616        | 0.382          |
| woman       | 0.791         | 0.409           | 0.746        | 0.455          |
| transman    | -0.062        | -0.152          | -0.095       | 0.018          |
| transwoman  | -0.088        | -0.271          | 0.050        | 0.062          |
| nonbinary   | 0.037         | -0.243          | 0.129        | 0.015          |

Table 4.9: Cosine similarity between gendered words and common occupations.

of non-binary pronouns, the subspace is likely representing the difference in frequency of terms rather than the concept of gender as a whole. Due to weaker alignment with the non-binary gender, any tasks performed using this new ‘gender’ subspace would not be very effective or applicable to non-binary genders, thus indicating towards further skews and harm. We can see from Figure 4.1 that it takes 1 PCA component to explain 70% of the variation in the word embeddings from binary pronouns, while the subspace constructed from nonbinary pronouns need about 3 components to do the same. Combined with the representations and their nearest neighbors, this gives us insight into how varied the robustness of these subspaces actually are.

| Category      | Pronouns   |
|---------------|--|
| Binary set    | <i>he, she, man, woman, hers, his, herself, himself, girl, boy, female, male</i> |
| Nonbinary set | <i>they, them, xe, ze, xir; zir, xey, zey, xem, zem, ey, em</i>                  |

Table 4.10: Binary and Nonbinary Pronoun Sets for PCA.




Figure 4.1: PCA Components for each Gender Subspace, indicating binary gender takes the least amount of components to represent, and is therefore less complex to model.

## 4.5 Discussion

In this chapter, we investigated the systematic challenges in representing gender identity in language models. We identify two critical limitations in current approaches to representing non-binary and other gender-diverse identities: the data scarcity for neopronouns often used by non-binary persons, and the oversimplification of treating non-binary gender as a single homogeneous category. Our analyses reveal fundamental tensions between discrete computational representations and the fluid nature of gender identity and expression. Because of this, we find that current modeling paradigms, which treat gender as static and categorical, inherently risk marginalizing certain populations. Based on our findings, we propose a research agenda that (1) questions the fundamental

assumptions of gender modeling in NLP, (2) advocates for participatory research methods that meaningfully involve affected communities, and (3) emphasizes the need for longitudinal monitoring of potential harms.

# CHAPTER 5

## From Skew to Erasure: The Role of Tokenization in Gender-Diverse Bias Propagation and Mitigation Strategies

While previous chapters revealed how language models can systematically erase and misrepresent gender minorities, we now investigate a critical yet understudied mechanism driving such phenomena: tokenization. In this chapter, we demonstrate how the over-fragmentation of gender-diverse pronouns during tokenization directly impacts a model’s ability to learn proper pronoun morphosyntax, ultimately perpetuating LLM misgendering. We introduce novel mitigation strategies, including Pronoun Tokenization Parity (PTP) and targeted lexical layer finetuning, that significantly improve neopronoun consistency while maintaining model performance. This chapter is based on our work in [OMG24], providing both technical insights into bias propagation while proposing several forms of bias mitigation.

### 5.1 Introduction

Gender bias in NLP has been extensively studied for binary gender, however mitigating harmful biases for underrepresented gender minorities remains an active area of research [SGT19, SA21]. [SGT19, SA21]. Previous studies [DMO21, OGD23, HDS23] have shown that large language models (LLMs) often fail to correctly use non-binary pronouns, particularly neopronouns such as xe and ey.<sup>1</sup> These works highlight the connection between LLM misgendering<sup>2</sup> and data scarcity,

---

<sup>1</sup>[https://nonbinary.wiki/wiki/English\\_neutral\\_pronouns](https://nonbinary.wiki/wiki/English_neutral_pronouns)

<sup>2</sup>The act of intentionally or unintentionally addressing someone (oneself or others) using a gendered term that does not match their gender identity.

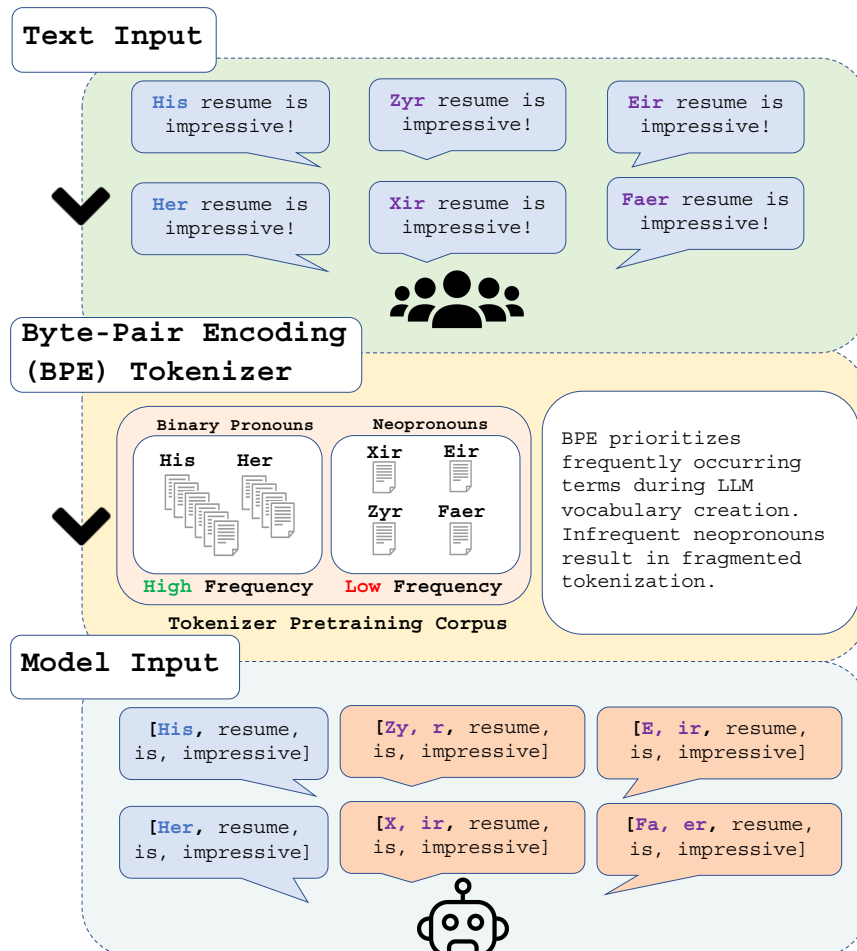


Figure 5.1: Byte-Pair Encoding (BPE) tokenization disproportionately fragments neopronouns compared to binary pronouns due to their infrequency in the training corpus. Our paper reveals that this overfragmentation leads to syntactic difficulties for LLMs, which are tied to their propensity to misgender data-scarce pronouns.

as neopronouns are severely underrepresented in pretraining corpora, thus limiting the LLM’s ability to use them proficiently. Despite this, the specific pathways through which data scarcity contributes to LLM misgendering behavior remain underexplored. Our work aims to address this research gap by investigating a critical, yet understudied aspect to LLM misgendering: tokenization.

Figure 5.1 illustrates the tokenization differences between binary pronouns and neopronouns when using Byte-Pair Encoding (BPE), the most widely adopted subword tokenizer employed by popular LLMs such as GPT-4 [BMR20], Claude <sup>3</sup>, Mistral [JSM23], and Llama 2 [TMS23]. While binary pronouns (*her* and *his*) are tokenized as single units, neopronouns *zyr*, *eir*, *xir*, and *faer* are fragmented into two subword tokens due to their infrequency within the tokenizer’s training corpus. As a result, the LLM must rely on more granular subword tokens to learn the neopronoun’s representation. Prior research finds that token overfragmentation adversely affects Part-of-Speech tagging and dependency parsing performance, as subword tokens share their embeddings across common words, introducing contextual ambiguity [WYS19, LBM23]. However, the impact of this phenomenon on English LLM misgendering remains unexplored.

**Contributions** To the best of our knowledge, our work is the first to link LLM misgendering to subword tokenization and deficient neopronoun grammar. We employ a series of evaluations that target understanding the association between LLM misgendering and poor pronoun morphosyntax (§5.4), finding that neopronoun misgendering is strongly associated with an LLM’s inability to use neopronouns as pronouns (§5.4.0.3).

Through a series of carefully controlled experiments, we demonstrate that mitigations centered on improving LLM neopronoun proficiency reduce neopronoun misgendering. We introduce *pronoun tokenization parity* (PTP), a technique to better preserve neopronoun tokens as functional morphemes by enforcing parity between neopronoun and binary pronoun tokenization (§5.5). Furthermore, we investigate leveraging pre-existing LLM pronoun knowledge to improve the model’s

---

<sup>3</sup><https://www.anthropic.com/news/claude-3-family>

grammatical usage of neopronouns (§5.5). Our results demonstrate that finetuning GPT-based models with PTP achieves up to 58.4% pronoun consistency, significantly outperforming the 14.1% obtained from finetuning with standard BPE tokenization. Notably, finetuning the LLM’s lexical layer with PTP outperforms traditional finetuning in 75% of models, reducing compute time by up to 21.5%. We find lexical finetuning consistently improves LLM pronoun consistency across model sizes, with smaller models experiencing the most significant gains—even matching the performance of models twice their size (§5.7.0.1).

## 5.2 Background & Related Works

**Gender-Inclusive NLP** Gender bias has been studied across several NLP contexts, including machine translation [SSZ19b], coreference resolution [RNL18, ZWY18c], and named entity recognition [MGM19]. Works like [GSB21] and others have found that choice of word segmentation exacerbates gender biases in machine translation. Recent works expand gender bias evaluations to harms unique to non-normative gender communities within LLMs [DMO21, HDS23, OGD23, NBL22, FCJ23, QDO23]. [DMO21] examine non-binary gender bias in static and contextual language representations, highlighting how data limitations affect these embeddings. Similarly, [OGD23] explore misgendering and harmful responses related to gender disclosure using their TANGO framework, pointing to challenges in neopronoun consistency, possibly due to data scarcity. [HDS23] corroborate these findings with an in-context-learning evaluation and analyses into LLM pretraining corpus statistics. Despite exploring various in-context learning strategies, they find persistent gaps between binary pronoun and neopronoun misgendering. These studies collectively emphasize data scarcity’s impact on neopronouns, though questions remain regarding how data scarcity shapes neopronoun representations and subsequent LLM pronoun consistency. In this study, we investigate the pivotal role of BPE tokenization due to its critical relationships to pretraining corpora and subsequent LLM vocabulary construction.

**BPE Tokenization** Byte-Pair Encoding (BPE) [Sen16] is a subword tokenization technique that constructs token vocabularies by iteratively merging frequently occurring adjacent token pairs up

| $\zeta$            | Nom.     | Acc.     | Genitive  |             | Reflex.         |
|--------------------|----------|----------|-----------|-------------|-----------------|
|                    |          |          | Dep.      | Ind.        |                 |
| <b>Binary</b> 1.20 | he       | him      | his       | his         | [him, self]     |
|                    | she      | her      | her       | hers        | [her, self]     |
| <b>Neo</b> 1.87    | ey       | em       | [ei, r]   | [e, irs]    | [em, self]      |
|                    | xe       | [x, em]  | [x, ir]   | [x, irs]    | [x, ir, self]   |
|                    | [f, ae]  | [fa, er] | [fa, er]  | [fa, ers]   | [fa, ers, elf]  |
|                    | zie      | [z, ir]  | [z, ir]   | [z, irs]    | [z, ir, self]   |
|                    | ze       | [h, ir]  | [h, ir]   | [h, irs]    | [h, ir, self]   |
|                    | sie      | [h, ir]  | [h, ir]   | [h, irs]    | [h, ir, self]   |
|                    | [th, on] | [th, on] | [th, ons] | [th, ons]   | [th, ons, self] |
|                    | ve       | ver      | vis       | vis         | [vers, elf]     |
| ne                 | ner      | [n, is]  | [n, is]   | [nem, self] |                 |

Table 5.1: BPE-tokenized Binary Pronouns and Neopronouns across pronoun forms.  $\zeta$ = Fertility. The closer fertility is to 1, the more the tokenizer kept pronoun tokens fully intact. **Bold** = neopronoun tokenization that does not follow binary pronoun forms.

to a predefined vocabulary size. Unseen or rare words are decomposed into subword units, down to individual characters, thus removing the need for assigning “unknown” token (`[UNK]`) to unseen words. However, this approach does not consider context, posing limitations for task-relevant yet data-scarce scenarios [YP22].

### 5.3 Low-Resource Challenges for BPE

**Data-Scarce Tokenization** [BD20] find that tokenization introduces a significant amount of inductive bias in LLMs, profoundly impacting their ability to perform tasks downstream. BPE prioritizes keeping the most frequent words intact during tokenization while splitting lower-frequency



texts into smaller subword tokens, irrespective of their contextual relevance [YP22, MAS21]. This behavior leads to learning critical aspects of language, like pronoun morphosyntax, through reliance on textual frequency, resulting in a fragmented understanding of morphosyntactic rules for less frequent pronoun sets. This tokenization disparity is reflected in Table 5.1 across tokenized pronoun groups and their respective fertility scores [RPV20], i.e., the average number of subwords produced per tokenized word. Binary pronouns are kept intact after tokenization, while most neopronouns are segmented into subword tokens, indicating that the LLM’s predefined vocabulary cannot construct these tokens. We posit that this lack of parity in tokenization between pronouns contributes to LLM misgendering downstream.

**OOV Pronouns and Hindered Grammatical Knowledge** [WYS19] find that *OOV words*, words that were unable to remain fully intact after tokenization, have detrimental impacts on downstream part-of-speech (POS) proficiency. Resulting token overfragmentation presents challenges across additional tasks such as named entity recognition [DS20, WDX22], dependency parsing [LBM23], and machine translation [DGH18, HHF19, AMN22]. [LBM23] find that because subwords are present in multiple words, their embeddings incorporate information from these common words, making the resulting ambiguity challenging to parse. Because of this, we hypothesize that the observed overfragmentation of tokenized neopronouns relates to LLM deficiencies in learning proper neopronoun morphosyntax.

## 5.4 Tracing LLM Misgendering to Grammatical Deficiencies

This section presents a series of metrics to evaluate LLM misgendering from the standpoint of pronoun proficiency. We perform baseline evaluations on out-of-the-box GPT-Neo-X based models and provide an overview of our evaluation scheme in Figure 5.2.

### 5.4.0.1 Evaluation Setup

**Models** We employ the Pythia model suite for our evaluation and experiments,<sup>4</sup> as it parallels state-of-the-art architecture; Pythia models are all built on top of a GPT-Neo-X architecture, an open-source alternative to GPT-3 models. Notably, it is based on a BPE tokenizer [BSA23] and trained on the PILE dataset [GGB20]. We use the *deduped* versions of Pythia, which trained on the Pile after the dataset had been globally deduplicated. We confirm that our research is in line with Pythia’s intended use: Given their Apache 2.0 license, we may finetune or adapt these models.

**Dataset** We utilize the MISGENDERED dataset by [HDS23], containing added templates and names from TANGO [OGD23], resulting in 93,600 templates to evaluate LLMs on our three metrics. We provide further dataset details in the sections below and in the Appendix [OMG24].

### 5.4.0.2 Evaluation Metrics

According to [Gar16], English pronouns must agree with their subject in gender, case, and number. We define three metrics to quantify a model’s understanding of different pronoun forms: two are standard misgendering measurements, and one is a novel metric introduced in this paper. *Pronoun consistency* (Consistency) assesses pronoun-gender agreement and is the primary metric for determining performance improvement in this paper. Previous studies find that this automatic consistency evaluation highly correlates to human evaluation [OGD23]. *Pronoun Case Agreement Error* (Case Error) is an auxiliary metric that provides insight into how well the model has learned pronoun forms. To test the relationship between LLM misgendering and poor LLM morphosyntax, we introduce *Adversarial Injection Error* (Inject Error) to measure LLM robustness against word insertion adversarial attacks that render a sentence grammatically incorrect or change its meaning. If there is an association between poor consistency and adversarial error, it would support formulating mitigations that prioritize enhancing the LLM’s overall grammatical proficiency with neopronouns. These metrics are employed in a constrained decoding setting, consistent with the

---

<sup>4</sup><https://github.com/EleutherAI/pythia>

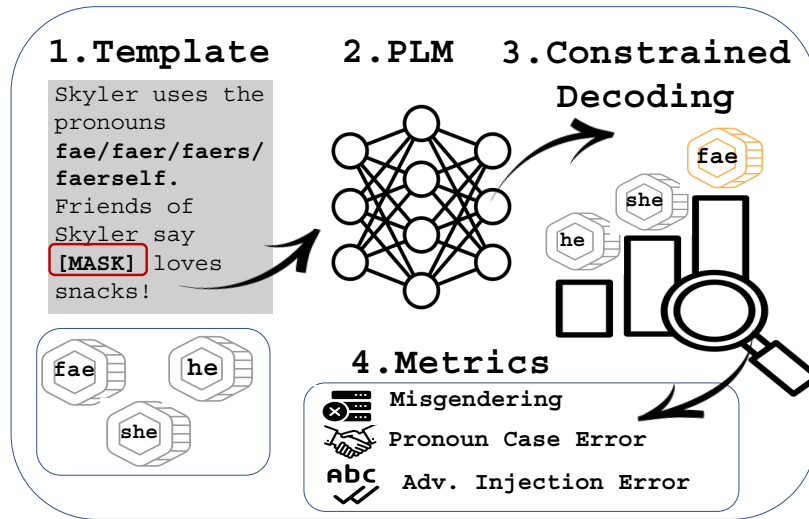


Figure 5.2: Evaluation. We determine our method’s efficacy in reducing LLM misgendering using a constrained decoding approach across 3 metrics.

MISGENDERED framework introduced by [HDS23]. Given a masked template, the LLM predicts the most likely pronoun from a pool of pronouns of the same form.

**Pronoun Consistency** Let  $S$  be a set of unique pronoun families with  $|S|$  pronoun families. Each pronoun family  $M \in S$  contains  $|M|$  English pronoun forms. Within a collection of masked templates  $T$ , [MASK] is replaced with a pronoun  $p \in M$  for all  $M \in S$ , resulting in the filled template set  $T^*$ . In line with [HDS23], each template starts with a person’s name and their pronoun declaration (i.e., nominative / accusative / genitive / reflexive), followed by a sentence containing a [MASK] token which expects a pronoun. For example: *Casey uses the pronouns he/him/his/himself. Upon recognizing Casey, the fan asked [MASK] for an autograph..* For a template  $t$  consisting of  $m$  tokens  $x_1, x_2, \dots, x_m$ , the token generated at [MASK],  $\hat{y}_t$ , is defined as the  $\text{argmax}$  transition probability from the pronoun pool.

$$\hat{y}_t = \text{argmax}_{p \in S} P(x_i = s | x_{<i}) \quad (5.1)$$

We denote the set of filled templates as  $C$ . Each filled template is then compared to its golden label example  $c \in C^*$ , containing the correct pronoun for that template-name-declaration combi-

nation.

To evaluate pronoun consistency, we compare the model’s chosen pronoun for a template,  $\hat{y}_t$ , to the template’s correct pronoun,  $y_c$ , and then calculate the accuracy over all templates:

$$\frac{1}{|T^*|} \sum_{t \in T^*, y \in C^*} \delta(\hat{y}_t, y_c) \quad (5.2)$$

**Pronoun Case Error** Evaluating pronoun case error is essential for assessing a model’s competence in pronoun usage. Ideally, an LLM would generate case-agreeing sentences like “She went to the store.” instead of “Hers went to the store.” To evaluate this, we use the same approach as above, instead focusing on assessing expected versus predicted pronoun cases for a given pronoun family. However, transition probabilities conditioned solely on preceding tokens cannot be relied on to determine case correctness. For example, a sentence like “Casey went to the store for [MASK] mom” can have its mask replaced with “her” or “herself” and still be grammatically correct, as it only considers the previous tokens during inference. Therefore, we obtain the model’s predicted output across all pronoun cases for a given family  $s \in Q$ , minimizing its loss (i.e., maximizing probability). Pronoun case error is then the proportion of templates with *incorrect* case agreement for a given pronoun family.

$$\operatorname{argmin}_{s \in Q} \left( - \sum_{i=1}^N \log P_{\theta}(x_i | x_{<i}) \right) \quad (5.3)$$

**Adversarial Injection Error** Prior research finds that prompting LLMs with texts containing neopronouns often results in ungrammatical generations, where neopronouns are incorrectly preceded by articles and determiners such as ‘the’, ‘a’, or ‘these’ [OGD23]. To further examine an LLM’s inability to construct grammatically correct sentences with neopronouns, we replicate this observed behavior by generating a set of otherwise grammatically correct prompts that include adversarial word insertions, making the template entirely ungrammatical. We use the same templates as previously defined but now augment each [MASK] to [DET]\_[MASK], where [DET] is replaced by singular and plural determiners (e.g., ‘this’, ‘those’, ‘these’), articles (like ‘the’, ‘a’), or

no determiner at all. Example templates are provided in Appendix [OMG24]. Similar to pronoun consistency, we employ LLM transition probabilities to evaluate how often LLMs use neopronouns in ungrammatical contexts. Next, we analyze the LLM’s output by calculating the  $\text{argmax}$  of the transition probability for all potential substitutions of [DET] (Equation 5.1). An LLM utilizing a neopronoun correctly should choose a template without a determiner. Models displaying incorrect behavior indicates poor grammatical proficiency with neopronouns.

### 5.4.0.3 Results

We report pronoun consistency, pronoun case error, and adversarial injection errors in Table 5.2. In line with prior work, the neopronoun *xe* reflects the lowest pronoun consistency (i.e., highest misgendering) across all model sizes. To better understand how this relates to grammatical issues, we also calculate Spearman’s correlation between pronoun consistency and each of the two error metrics (leftmost results column). Notably, we observe moderate to strong negative correlations between grammatical error metrics and misgendering. Across model sizes, we find a range of  $-0.45$  to  $-0.63$  correlation for injection error and  $-0.53$  to  $-0.63$  for case error. With these observations, we posit that mitigation strategies that enhance an LLM’s grammatical proficiency with neopronouns will attenuate their tendency to misgender.

## 5.5 Improving LLM Neopronoun Proficiency

**Pronoun Tokenization Parity** English pronouns serve as building blocks for language acquisition. Termed *functional morphemes*, these small, self-contained units of meaning reflect specific English grammatical functions [For05, ES11]. To improve LLM neopronoun consistency, we introduce *pronoun tokenization parity* (PTP), a method that maintains a token’s functional integrity during BPE tokenization. By aligning neopronoun tokenization with that of binary pronouns, we aim to improve an LLM’s grammatical understanding of neopronouns, ultimately enhancing the model’s ability to use them correctly.

Formally, we extend the pretrained token embeddings of a transformer-based LLM. To do this,

| Size | Metric                        | $\rho$ | Pronoun Family        |                       |                              |
|------|-------------------------------|--------|-----------------------|-----------------------|------------------------------|
|      |                               |        | He                    | She                   | Xe                           |
| 70M  | Consistency ( $\uparrow$ )    | —      | 96.82 <sub>0.77</sub> | 71.59 <sub>2.00</sub> | <b>0.67</b> <sub>0.35</sub>  |
|      | Case Error ( $\downarrow$ )   | -0.63  | 8.26 <sub>1.21</sub>  | 24.36 <sub>1.90</sub> | <b>78.56</b> <sub>1.82</sub> |
|      | Inject Error ( $\downarrow$ ) | -0.45  | 23.85 <sub>1.88</sub> | 16.92 <sub>1.66</sub> | <b>85.03</b> <sub>1.58</sub> |
| 160M | Consistency ( $\uparrow$ )    | —      | 79.95 <sub>1.82</sub> | 76.46 <sub>1.90</sub> | <b>0.00</b> <sub>0.00</sub>  |
|      | Case Error ( $\downarrow$ )   | -0.59  | 4.05 <sub>0.90</sub>  | 10.87 <sub>1.38</sub> | <b>80.00</b> <sub>1.77</sub> |
|      | Inject Error ( $\downarrow$ ) | -0.63  | 8.72 <sub>1.28</sub>  | 6.46 <sub>1.10</sub>  | <b>95.38</b> <sub>0.92</sub> |
| 410M | Consistency ( $\uparrow$ )    | —      | 72.82 <sub>1.92</sub> | 55.85 <sub>2.21</sub> | <b>0.05</b> <sub>0.08</sub>  |
|      | Case Error ( $\downarrow$ )   | -0.53  | 2.87 <sub>0.74</sub>  | 7.90 <sub>1.21</sub>  | <b>79.90</b> <sub>1.79</sub> |
|      | Inject Error ( $\downarrow$ ) | -0.54  | 4.15 <sub>0.90</sub>  | 3.49 <sub>0.79</sub>  | <b>89.85</b> <sub>1.36</sub> |
| 1.4B | Consistency ( $\uparrow$ )    | —      | 78.46 <sub>1.82</sub> | 66.56 <sub>2.03</sub> | <b>0.26</b> <sub>0.23</sub>  |
|      | Case Error ( $\downarrow$ )   | -0.54  | 3.54 <sub>0.82</sub>  | 3.03 <sub>0.74</sub>  | <b>76.00</b> <sub>1.92</sub> |
|      | Inject Error ( $\downarrow$ ) | -0.62  | 3.69 <sub>0.85</sub>  | 3.44 <sub>0.79</sub>  | <b>92.77</b> <sub>1.15</sub> |

Table 5.2: Out-of-the-box evaluations on Pythia, a GPTNeo-X based model across sizes. Uncertainty estimates are 95% confidence intervals computed from 10k bootstrap iterations. *Takeaway: Markedly higher grammatical error rates for neopronoun vs. binary pronouns.*

let  $E_1^{\text{orig}}, E_2^{\text{orig}}, \dots, E_n^{\text{orig}}$  denote the original embeddings, where  $n$  represents the vocabulary size of the original model. We introduce new embeddings  $E^{\text{PTP}}$  for each of the  $m$  unique pronouns in the set of neopronoun cases  $S$  (i.e., pronoun family). This results in an extended vocabulary:

$$\{E_1^{\text{orig}}, \dots, E_n^{\text{orig}}\} \cup \{E_1^{\text{PTP}}, \dots, E_m^{\text{PTP}}\}.$$

We provide additional details and instructions for reproducing PTP in Algorithm 1.

---

## Algorithm 1 Pronoun Tokenization Parity (PTP)

---

- 1: **Inputs:** (1) LLM, (2) LLM BPE tokenizer, (3) list of neopronouns for PTP, (4) finetuning dataset
- 2: **Method:** Add special tokens for each neopronoun. Be sure to explicitly add 'Ġ' to the beginning of each token to indicate that it is a full, non-subword token space before the word, otherwise this will lead to incorrect model behavior, since a lack of 'Ġ' in BPE tokenization indicates a subword token.
- 3: **Check:** Check the tokenizer is working properly by checking the tokenized neopronoun, ensuring that you see 'Ġ' in its token. For example, tokenizing *xe* should result in ['Ġxe'] not ['Ġ', 'xe']. The latter will cause the LLM to incorrectly associate a space character with a neopronoun. This can be tested by checking next word transition probabilities from the space character.
- 4: Resize the LLM token embeddings to match vocabulary of tokenizer. Here is example code to do this with a model and tokenizer from HuggingFace Transformers Package<sup>5</sup>.

```
#declare neopronoun tokens
arr_tokens = [
    'Ġxe', 'ĠXe',
    'Ġxem', 'ĠXem',
    'Ġxir', 'ĠXir',
    'Ġxirs', 'ĠXirs'
]
# add new tokens to the tokenizer, t
token_dict = {
    'additional_special_tokens': arr_tokens
}
t.add_special_tokens(token_dict)
# update model, m, accordingly
m.resize_token_embeddings(len(tokenizer))
```

- 5: **Return** Finetuned model, new PTP tokenizer
  - 6: Evaluate using extended MISGENDERED framework
-

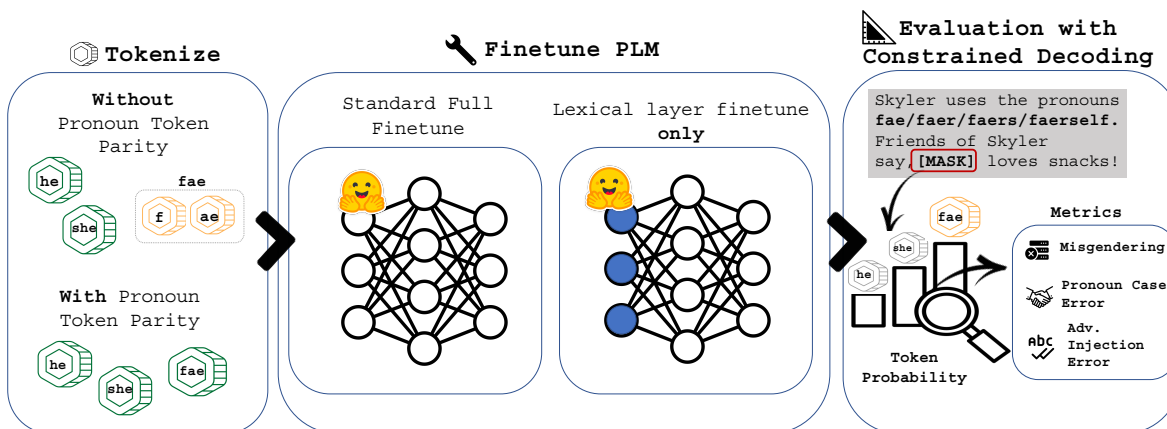


Figure 5.3: Overview. We (1) tokenize neopronouns using PTP for a given LLM, (2) either fully finetune or only finetune the LLM lexical layer with data containing neopronouns, and (3) determine our method’s efficacy in reducing LLM misgendering using a constrained decoding approach across 3 metrics.

**Leveraging LLM Pre-Existing Pronoun Knowledge** Training a new tokenizer and LLM requires significant computational resources and data. Pre-trained English LLMs have learned English syntax and pronouns during pretraining. We can take advantage of morphosyntactic similarities between binary pronouns and neopronouns, such as their syntactic roles and agreement patterns, to transfer knowledge from one set of pronouns to another.

Guided by fundamental aspects of cross-lingual transfer detailed in [ARY19] and [VN21], we propose the practice of finetuning only an LLM’s lexical embedding layer while keeping downstream transformer weights fixed. As long as the source and target pronoun groups share similar linguistic foundations, mirroring those found in cross-lingual sharing of basic elements, we can sidestep common challenges in cross-lingual transfer, such as determining the most suitable transfer source language. Unlike [ARY19], we forgo training the transformer weights after freezing lexical embeddings since the new tokens already align with English grammar and syntax, eliminating the need for the transformer to adapt to a different language. Furthermore, in contrast to the approach by [VN21], we avoid resetting the entire lexical embedding layer to preserve the



prelearned English grammar dependencies.

## 5.6 Experimental Setup

We provide an overview of our experimental setup in Figure 5.3. We conduct carefully controlled experiments across two finetuning paradigms using open-source LLMs that vary in model size and neopronoun data scarcity. In the first set of experiments, we employ PTP in a standard full finetuning paradigm. In the second experiment, we introduce lexical finetuning and variants with PTP. We perform these experiments across binary pronouns and the neopronoun family *xe*. We center *xe* for several reasons: *xe* ranks among the most widely adopted non-binary pronouns [Gen23]. Non-binary pronouns also exhibit diverse linguistic variations, spanning from closed to open word class forms [Mil16, LCH22]. This diversity requires a nuanced yet flexible approach. By focusing on the *xe* pronoun family, we showcase the effectiveness of PTP while providing a generalizable framework for researchers to build upon for studying non-binary pronouns within their respective linguistic contexts.

### 5.6.0.1 Finetuning Dataset

We finetune our models on the WIKIBIOS<sup>6</sup> dataset, comprising 728,321 English biographical texts from Wikipedia. Counterfactual data augmentation is used to address the limited availability and narrow dimensions of textual corpora containing neopronouns. We replace a variable proportion of binary pronouns with their neopronoun counterparts. Acknowledging that individuals who use neopronouns often have prior associations with binary pronouns, this data curation strategy enables LLMs to acquire knowledge of neopronouns within more comprehensive, diverse, and real-world contexts [TL22].

We filter the WIKIBIOS dataset to retain texts containing binary pronouns, resulting in 462,345 examples. Each binary pronoun is replaced with its corresponding neopronoun case, incorporating

---

<sup>6</sup>[https://huggingface.co/datasets/wiki\\_bio](https://huggingface.co/datasets/wiki_bio)

correct possessive forms using the spaCy part-of-speech tagger.<sup>7</sup> No biography text appears more than once in the dataset splits. To understand how our methods operate across data resource levels, we counterfactually augment with an increasing proportion of neopronouns: 10%, 20%, 30%, 40%, and 50%. At the 50% level, the dataset is evenly split between neopronouns and binary pronouns.

### 5.6.0.2 Finetuning Setups

**Pronoun Tokenization Parity** To test whether PTP helps mitigate LLM misgendering, we prepare two versions of finetuning for a compact 70M parameter Pythia model. The first model is finetuned with its original BPE tokenizer ( $T_{\text{ORIG}}$ ) and the second with PTP ( $T_{\text{PTP}}$ ). Embeddings for  $T_{\text{PTP}}$  are initialized with a random Gaussian ( $\mu=0$  and  $\sigma=0.02$ ).  $M_{\text{FULL}}$  denotes all models with standard full finetuning, and  $M_{\text{BASE}}$  represents the HuggingFace out-of-the-box checkpoint which uses its original BPE tokenizer  $T_{\text{ORIG}}$ .  $T_{\text{ORIG}} + M_{\text{BASE}}$  and  $T_{\text{ORIG}} + M_{\text{FULL}}$  serve as baselines for PTP.

Each model is finetuned across five epochs with a batch size of 128 and a  $10^{-4}$  learning rate. Before tokenization, text is chunked with a 256 window size, resulting in 386,267 rows before any neopronoun augmentation. We conduct finetuning with an 80/10/10 train, validation, and test split. To encourage model generalization and prevent overfitting, we incorporate weight decay regularization (0.01), a warmup ratio of 0.01 to gradually increase the learning rate over the initial 1% of training steps, and apply early stopping based on cross-entropy loss in the validation set with a patience of 2. All models undergo finetuning using FP16 mixed precision and two gradient accumulation steps.

**Lexical Layer Finetuning** We follow the same setup as before but now increase the learning rate to  $10^{-3}$  to encourage more rapid adaptation to the new vocabulary. We denote models trained with lexical finetuning with original BPE tokenization as  $T_{\text{ORIG}} + M_{\text{LEX}}$ . We compare performance to PTP and PTP baselines:  $T_{\text{PTP}} + M_{\text{FULL}}$ ,  $T_{\text{ORIG}} + M_{\text{BASE}}$  and  $T_{\text{ORIG}} + M_{\text{FULL}}$ . We also introduce an additional lexical finetuning variant with PTP ( $T_{\text{PTP}} + M_{\text{LEX}}$ ) and test to what extent combining

---

<sup>7</sup><https://spacy.io/>



Figure 5.4: 70M model pronoun consistency for each pronoun family across 10-50% data resource levels and model variants. *Takeaway: PTP sustains improvements in neopronoun consistency across data resource levels.*

these techniques boosts performance over either method.

**Model Size Ablations** In order to evaluate the effectiveness of our proposed mitigations at various scales and resource levels, we repeat our experiments at 160M, 410M, and 1.4B parameters. Furthermore, we ensure that all finetuned models do not overfit nor adversely impact pre-existing performance on downstream tasks, reporting test set evaluations and a case study on downstream tasks in the Appendix [OMG24].

## 5.7 Results

**Pronoun Tokenization Parity** We report our PTP finetuning results in Table 5.3. Both  $T_{PTP} + M_{FULL}$  (37.8%) and  $T_{ORIG} + M_{FULL}$  (14.5%) demonstrated gains in neopronoun consistency over

| Model              | Metric                        | He                           | She                          | Xe                           |
|--------------------|-------------------------------|------------------------------|------------------------------|------------------------------|
| $T_{\text{Orig}}+$ | Consistency ( $\uparrow$ )    | <b>96.82</b> <sub>0.79</sub> | 71.59 <sub>2.03</sub>        | 0.67 <sub>0.38</sub>         |
|                    | Case Error ( $\downarrow$ )   | 8.26 <sub>1.26</sub>         | 24.36 <sub>1.90</sub>        | 78.56 <sub>1.77</sub>        |
| $M_{\text{Base}}$  | Inject Error ( $\downarrow$ ) | 23.85 <sub>1.90</sub>        | 16.92 <sub>1.67</sub>        | 85.03 <sub>1.56</sub>        |
| $T_{\text{Orig}}+$ | Consistency ( $\uparrow$ )    | 89.64 <sub>1.36</sub>        | <b>86.05</b> <sub>1.54</sub> | 14.46 <sub>1.56</sub>        |
|                    | Case Error ( $\downarrow$ )   | 11.74 <sub>1.44</sub>        | 22.41 <sub>1.87</sub>        | 59.95 <sub>2.15</sub>        |
| $M_{\text{Full}}$  | Inject Error ( $\downarrow$ ) | 23.95 <sub>1.87</sub>        | 16.77 <sub>1.67</sub>        | 89.49 <sub>1.36</sub>        |
| $T_{\text{PTP}}+$  | Consistency ( $\uparrow$ )    | 94.77 <sub>0.97</sub>        | 83.49 <sub>1.67</sub>        | 37.79 <sub>2.10</sub>        |
|                    | Case Error ( $\downarrow$ )   | 9.69 <sub>1.31</sub>         | 29.28 <sub>2.00</sub>        | 56.92 <sub>2.15</sub>        |
| $M_{\text{Full}}$  | Inject Error ( $\downarrow$ ) | 27.79 <sub>1.95</sub>        | 20.97 <sub>1.79</sub>        | 27.03 <sub>1.95</sub>        |
| $T_{\text{Orig}}+$ | Consistency ( $\uparrow$ )    | 86.46 <sub>1.49</sub>        | 72.87 <sub>2.00</sub>        | 16.77 <sub>1.62</sub>        |
|                    | Case Error ( $\downarrow$ )   | 18.51 <sub>1.72</sub>        | 33.79 <sub>2.08</sub>        | 70.51 <sub>2.05</sub>        |
| $M_{\text{Lex}}$   | Inject Error ( $\downarrow$ ) | 28.97 <sub>2.05</sub>        | 23.18 <sub>1.87</sub>        | 65.44 <sub>2.10</sub>        |
| $T_{\text{PTP}}+$  | Consistency ( $\uparrow$ )    | 84.97 <sub>1.59</sub>        | 72.21 <sub>1.95</sub>        | <b>53.59</b> <sub>2.21</sub> |
|                    | Case Error ( $\downarrow$ )   | 18.15 <sub>1.72</sub>        | 33.03 <sub>2.08</sub>        | 60.46 <sub>2.15</sub>        |
| $M_{\text{Lex}}$   | Inject Error ( $\downarrow$ ) | 25.79 <sub>1.97</sub>        | 21.85 <sub>1.82</sub>        | 34.77 <sub>2.10</sub>        |

Table 5.3: 70M-parameter model results at 10% data resource level.  $T_{\text{ORIG}}$ = original BPE tokenizer,  $T_{\text{PTP}}$ = tokenizer with PTP,  $M_{\text{BASE}}$ = original model (no finetuning)  $M_{\text{FULL}}$ = full finetuning. Uncertainty estimates are 95% confidence intervals computed from 10k bootstrap iterations.

$T_{\text{ORIG}} + M_{\text{BASE}}$  (j1%). This improvement is expected, considering their increased exposure to neopronouns during finetuning. However, **models using PTP outperformed those finetuned with original BPE tokenization**. As shown in Figure 5.4, PTP’s improvement over these two baselines was consistent across data resource levels. We observed the best neopronoun consistency overall at 58.4% (50% data resource level). Notably, gains over vanilla finetuning ( $T_{\text{ORIG}} + M_{\text{FULL}}$ ) were most evident at resource levels below 30%, where  $T_{\text{PTP}} + M_{\text{FULL}}$  more than doubled neopronoun consistency over  $T_{\text{ORIG}} + M_{\text{FULL}}$  (14.5% vs. 37.8%). Binary pronoun consistency remained stable, with  $T_{\text{PTP}} + M_{\text{FULL}}$  even improving *she* pronoun consistency over  $T_{\text{ORIG}} + M_{\text{BASE}}$ . Notably, the

adversarial error rate for  $xe$  also dropped from 85% to 27% after finetuning with PTP, a decrease not observed after vanilla finetuning. These findings suggest that targeting LLM neopronoun proficiency significantly reduces the LLM’s tendency to misgender, with pronoun tokenization parity showing promise in addressing these challenges.

**Lexical Layer Finetuning** We report results for lexical finetuning variants in Table 5.3.  $T_{\text{ORIG}} + M_{\text{LEX}}$  improved neopronoun consistency (16.8%) over  $T_{\text{ORIG}} + M_{\text{BASE}}$  and  $T_{\text{ORIG}} + M_{\text{FULL}}$ , indicating that employing pre-existing LLM knowledge may improve neopronoun proficiency. While lexical finetuning alone contributed modest improvements over  $T_{\text{ORIG}} + M_{\text{FULL}}$ , **pairing lexical finetuning with PTP significantly outperformed all other models**, at 53.6% neopronoun consistency. This cumulative gain, accompanied by a simultaneous reduction in adversarial error over  $T_{\text{ORIG}} + M_{\text{FULL}}$  (34.8% vs. 89.5%), suggests a favorable synergy towards improving neopronoun morphosyntax. We also observed gains over  $T_{\text{PTP}} + M_{\text{FULL}}$  across all data resource levels, especially at 10% and 20%, demonstrating its efficacy in more real-world, lower-resourced settings (further details found in the appendix [OMG24]).

The impact of lexical finetuning on binary pronouns varied across models of this size. We observed stable consistency for feminine pronouns, while this was more evident for masculine pronouns with  $T_{\text{PTP}} + M_{\text{FULL}}$ . The decline in masculine pronouns after lexical training may be attributed to the distinct challenges associated with finetuning existing pronouns compared to new or under-resourced pronouns. Neopronoun tokens, which are not initialized from a pre-existing ”pronoun” space, must be learned from scratch. Meanwhile, binary pronoun tokens have already converged to a meaningful lexical space. As a result, while the LLM learns these new neopronouns, the previously trained binary pronouns may be inadvertently affected. In this work, we consider it an acceptable tradeoff as it substantially improves the most disadvantaged group (i.e., equity) without severely compromising overall performance. This phenomenon is typical in bias mitigation efforts, where gains in fairness are typically balanced against performance loss. Ultimately, the optimal tradeoff is stakeholder-dependent. Future studies can build upon these findings to investigate balancing equity with overall performance further.




Figure 5.5: Results across all models at data resource level=10. The uncertainty estimates are 95% confidence intervals computed from 10k bootstrap iterations. *Takeaway: Across model size, variants of PTP consistently improve neopronoun consistency over models employed with standard BPE.*

### 5.7.0.1 Model Size Ablations

Results for all model sizes are provided in Figure 5.5. Neopronoun consistency gains with PTP over finetuning with BPE tokenization were sustained across model sizes. Both  $T_{\text{PTP}} + M_{\text{FULL}}$  and  $T_{\text{PTP}} + M_{\text{LEX}}$  again outperformed neopronoun consistency baselines  $T_{\text{ORIG}} + M_{\text{FULL}}$  and  $T_{\text{ORIG}} + M_{\text{BASE}}$ . Lexical finetuning performed best when paired with PTP, as found in the previous section. Across size, we also found lexical finetuning reduced compute time by up to 21.5% over standard full finetuning (more results in Appendix [OMG24]).

$T_{\text{PTP}} + M_{\text{LEX}}$  provided gains over  $T_{\text{ORIG}} + M_{\text{FULL}}$  across all model sizes, with larger models (>160M) benefiting most from  $T_{\text{PTP}} + M_{\text{FULL}}$ . Notably, a larger model did not always improve neopronoun consistency across respective finetuning regimes. In fact, when employing PTP, **smaller models actually achieved neopronoun consistency comparable to models more than twice their size**. As shown in Figure 5.5, a 410M model finetuned with  $T_{\text{PTP}} + M_{\text{FULL}}$  resulted in the best neopronoun consistency (56.2%), while a 160M model finetuned with  $T_{\text{PTP}} + M_{\text{LEX}}$  closely followed (53.6%) (further details in [OMG24]). Further examining model generations, we provide examples in Table 5.4 which demonstrate consistent textual coherence for each of our finetuning paradigms.

| Model                                    | Text   |
|--|--|
| $T_{\text{Orig}}^+$<br>$M_{\text{Base}}$ | <i>Skyler's pronouns are xe/xem/xir/xirself. Xe needs to go to the bookstore soon because it's not in the store.</i><br>Xir is the same as Xirself.  |
| $T_{\text{Orig}}^+$<br>$M_{\text{Full}}$ | <i>Skyler's pronouns are xe/xem/xir/xirself. Xe needs to go to the bookstore soon because xe's too busy to be on the air. Xe's also a good friend of xir.</i>  |
| $T_{\text{PTP}}^+$<br>$M_{\text{Full}}$  | <i>Skyler's pronouns are xe/xem/xir/xirself. Xe needs to go to the bookstore soon because xe's going to be writing a book about how to say xir name. Xe also has a book in the works called "the art of being a writer."</i> |
| $T_{\text{Orig}}^+$<br>$M_{\text{Lex}}$  | <i>Skyler's pronouns are xe/xem/xir/xirself. Xe needs to go to the bookstore soon because xe won't have time to go tomorrow.</i>   |
| $T_{\text{PTP}}^+$<br>$M_{\text{Lex}}$   | <i>Skyler's pronouns are xe/xem/xir/xirself. Xe needs to go to the bookstore soon because xe is a huge fan of the book "the secret life of the apes" by john mccarthy.</i>   |

Table 5.4: Pythia-410M model generations across finetuning regimes. *Italics* are input prompts and generations are performed with nucleus sampling (TOP-P=0.95).

## 5.8 Discussion

In this work, we discovered how disparate BPE tokenization across gendered pronouns, a consequence of data infrequency in training corpora, is associated with a model’s degraded ability to adhere to pronoun morphosyntax. This deficiency is highly correlated with an LLM’s propensity to misgender data-scarce neopronouns. Parallels to low-resource multilingual NLP efforts in addressing tokenizer limitations help inform novel approaches to mitigating English neopronoun misgendering. We find that employing vocabulary amelioration with pronoun tokenization parity along with a monolingual twist on lexical finetuning improve LLM neopronoun consistency and grammatical proficiency over traditional finetuning settings with standard BPE tokenization.

As BPE is just one of many subword tokenization algorithms, our work opens new avenues for exploring this phenomenon under various subword tokenization algorithms and in multilingual settings. Nonetheless, these challenges ultimately arise from larger issues surrounding data

availability and limitations of greedy (i.e., context-free) tokenization techniques. Addressing these foundational issues in future work is essential for sustainably developing inclusive LLMs and preventing social harm.



**Part II**

**Technical Choices Meet Social  
Consequences**

## CHAPTER 6

# Beyond the Binary: Refining Conceptual Models of Gender-Inclusive Bias Evaluation and Mitigation

This chapter explores a critical limitation in current AI fairness research: the over-reliance on binary gender conceptualizations in gender bias evaluations. We find this focus has profound implications for how AI-driven systems are developed, evaluated, and ultimately deployed in the real world. Intersectionality, a framework for understanding how different social identities and power dynamics interact, allows us to interrogate the social constructs underpinning current fairness evaluations, revealing the ways binary and exclusionary views are integrated into debiasing techniques. This framework bridges the gap between technical fixes and social impact, offering a comprehensive approach to AI fairness that considers both root causes and symptoms of algorithmic bias. This chapter’s work is based previously published work from [OSG23] and [OPM24].

### 6.1 Introduction

| Word       | Doctor | Engineer | Nurse  | Stylist |
|------------|--------|----------|--------|---------|
| man        | 0.809  | 0.551    | 0.616  | 0.382   |
| woman      | 0.791  | 0.409    | 0.746  | 0.455   |
| transman   | -0.062 | -0.152   | -0.095 | 0.018   |
| transwoman | -0.088 | -0.271   | 0.050  | 0.062   |
| nonbinary  | 0.037  | -0.243   | 0.129  | 0.015   |

Table 6.1: Cosine similarity between gendered words and common occupations.

Gender bias benchmarking is key to unveiling gender disparities across various NLP tasks. Measuring gender bias is often conducted through *occupational bias* measurements, as depicted in Table 6.1[DMO21], where static word embeddings reveal how different gendered terms associate with professions. Model bias is detected as a disproportionate skew of binary pronouns to professions reflective of real-world gender binary power asymmetries, for instance “doctor” more likely for men and “nurse” with women. However, when evaluating terms such as “transman,” “transwoman,” and “nonbinary,” the benchmarks attempt to apply the same binary occupational stereotypes, resulting nonsensical associations for these gender identities. This breakdown highlights a key limitation of such benchmarks—those inherently structured around a binary understanding of gender are rendered inadequate for measuring harms for non-binary and transgender identities. To better understand this, we engage with intersectionality [Col19] in the following section—a framework that serves as a vehicle for engaging with the social depth required to understand and rectify these foundational issues.

## 6.2 Intersectionality on the Ground: A Framework for Social Grounding in AI Fairness

**Intersectionality** [Cre91] serves as an relational framework for analyzing how multiple social identities —such as race, gender, class, and sexuality—interact with systemic forces that go on to shape an individual’s experience of privilege or oppression. For AI-driven systems, this framework enables researchers to systematically examine the structural assumptions and power dynamics that inform bias in technical artifacts. Patricia Hill Collins breaks down intersectionality into six fundamental tenets [Col19]:

- (1) **Social Inequality:** Intersectionality rejects the notion that social inequalities are natural or inevitable. Instead, it scrutinizes how these inequalities are produced and perpetuated by intersecting systems of power.
- (2) **Power Relations:** Power dynamics operate across various domains—structural, disciplinary,

cultural, and interpersonal. Intersectionality seeks to uncover how these power relations establish and maintain social hierarchies and divisions.

- (3) **Relationality:** Social categories such as gender, race, and class are not isolated entities but are interrelated. Intersectionality promotes a relational understanding of how different forms of identity interact and influence one another.
- (4) **Social Context:** The significance and impact of intersecting identities fluctuate across different social, historical, and cultural contexts. Intersectionality emphasizes the importance of context in understanding identity and bias.
- (5) **Complexity:** Intersectionality acknowledges the intricate and multifaceted nature of lived experiences, encouraging resistance against oversimplified explanations or technical fixes.
- (6) **Social Justice:** Beyond being an analytical framework, intersectionality is a praxis-oriented tool that drives both critical inquiry and practical efforts toward achieving social justice.

Intersectionality helps us evaluate gender bias in AI by revealing two critical insights: how gender interacts with systems of power and how these interactions get encoded into technical systems. While current frameworks often reduce gender to simplistic categories, intersectionality shows how gender cannot be separated from race, class, and other aspects of identity and social position. This understanding exposes critical gaps in our ML pipeline—from who is represented in our training data, to how we operationalize gender in our annotations, to what our benchmarks choose to measure—revealing where our technical translations of gender can amplify existing social biases.

### **6.3 Intersectionality Illuminates Gaps in Gender Bias Benchmark Construct Validity**

Intersectionality provides a framework for understanding the fundamental limitations in gender bias evaluation revealed in Table 6.1. The benchmark’s failure stems from its reliance on binary

occupational stereotypes—societal assumptions that associate certain professions exclusively with either men or women.

When researchers attempt to evaluate gender-diverse concepts like "transman," "transwoman," and "nonbinary" using this binary framework, the resulting associations become non-interpretable, as they cannot be meaningfully mapped onto this restrictive binary axis. This example illustrates a broader **conceptual gap** in benchmark application. We can identify two critical types of gaps in bias evaluations that contribute to such limitations in inclusive bias evaluation:

- (1) **Representation Gaps:** Following construct validity principles [RBP21], representation gaps occur when evaluations exclude or underrepresent certain groups. In gender bias evaluation, the systematic exclusion of transgender and non-binary individuals means these benchmarks cannot effectively assess model fairness for these populations. For example, while datasets like BOLD [DSK21] incorporate multiple demographic attributes including gender, they remain restricted to binary gender categories. Expanding evaluation to include gender-diverse identities—when supported by relevant social context—can address these representational limitations.
- (2) **Conceptual Gaps:** These emerge when bias evaluation benchmarks employ frameworks that lack construct validity for certain groups. In gender bias evaluation, conceptual gaps arise when stereotypes and assumptions fail to align with the diverse realities of gender identity. Table 6.1 demonstrates this: binary occupational stereotypes prove inadequate for evaluating bias across the full spectrum of gender identities. Attempts to apply such binary-based benchmarks to non-binary identities invalidate the evaluation process, as the underlying assumptions fundamentally misalign with the populations being evaluated.

Conceptual gaps, as reflected in social misalignment in gender-diverse bias evaluation in Table 6.1, reveal fundamental limitations in current fairness frameworks, particularly the critical disconnect between technical fairness metrics and their real-world implications. Our analysis shows how structural assumptions—specifically in benchmark construction and validation—determine

which disparities we measure and how we quantify them. Through an intersectional lens, we demonstrate that these gaps persist not from data scarcity, but from methodological choices in selecting evaluation populations and defining harm metrics. This suggests that improving fairness outcomes requires reconceptualizing evaluation frameworks themselves, rather than merely expanding datasets. To illustrate the pervasive nature of these conceptual gaps, we next examine their manifestation across (1) other gender bias benchmarks in top-performing LLMs and (2) various bias mitigation techniques, highlighting the inseparable relationship between structural gaps in evaluation and mitigation. We demonstrate that current approaches, without such methodological revision, risk perpetuating the very biases they attempt to quantify and address.

## 6.4 Persistence of Conceptual Gaps in LLM Gender Bias Evaluation and Mitigation

**Gender Bias Evaluation Benchmarks for Top-Performing LLMs** Upon surveying bias evaluation modalities for the top 15 LLMs reported by the Chatbot Arena Leaderboard (Table 6.2), we find employed bias benchmarks assess binary gender bias, while offering little to no coverage for gender-diverse identities and other socially salient dimensions.<sup>1</sup> As shown in Figure 6.1, we find employed benchmarks capture different aspects of LLM bias, though focus primarily on binary gender, neglecting other social and demographic factors.<sup>2</sup> WINOGENDER, WINOBIAS exclusively cover binary gender identities in occupational stereotypes. BOLD evaluates fairness in open generation across multiple domains, but its gender bias assessments remain confined to binary categories. DISCRIM-EVAL and BBQ extend gender identity coverage but still face significant limitations: DISCRIM-EVAL includes only ‘non-binary’ as a gender-diverse identity category and measures LLM discrimination based on hypothetical scenarios rather than documented social harms. BBQ measures LLM reflections of attested social bias and includes ‘transgender man/-

---

<sup>1</sup><https://lmarena.ai/?leaderboard>

<sup>2</sup>We exclude REALTOXICITYPROMPTS[GG20] and TRUTHFULQA [LHE22] as they measure toxic degeneration from neutral prompts and general falsehoods, respectively, rather than bias against targeted demographic or social characteristics.

Table 6.2: Bias evaluation modalities for top 15 performing LLM families reported by Chatbot Arena Leaderboard.

| Model           | Benchmarks                                     | No Bias Eval |
|-----------------|--|--------------|
| GPT-4o          | TruthfulQA                                     |              |
| GEMINI-ADVANCED | Winogender, Winobias, BBQ, RealToxicityPrompts |              |
| GPT4-TURBO      | RealToxicityPrompts                            |              |
| CLAUDE 3 OPUS   | Discrim-Eval, BBQ                              |              |
| YI              | TruthfulQA                                     |              |
| REKA-CORE       |  | ✓            |
| COMMAND R+      | TruthfulQA                                     |              |
| QWEN 2          |  |              |
| QWEN MAX        |  |              |
| GLM-4           |  | ✓            |
| MISTRAL         |  |              |
| CLAUDE 1        | Discrim-Eval, BBQ, TruthfulQA                  |              |
| MIXTRAL         | BBQ, BOLD                                      |              |
| CLAUDE 2        | Discrim-Eval, BBQ                              |              |
| ZEPHYR-ORPO     |  | ✓            |

woman’ gender identities, though this inclusion remains undocumented in its original paper.

This narrow focus creates two issues: (1) binary gender-exclusive measurements of LLM harms risk leaving biases affecting gender minorities unchecked and (2) it further entrenches cisnormative hegemonies in competitive LLM benchmarking, encouraging other models to mirror these evaluation practices [SB18, Key18, OGD23]. While expanding existing evaluations to include more groups is a step forward, doing so without proper construct validation risks neglecting significant power asymmetries that marginalized communities face [WMR21, BLO21, RBP21].

|                     | Binary Gender | Gender-diverse | Sexual orientation | Occupation | Nationality | Race/ethnicity | Religion | Disability | Age | Body type/<br>physical appearance | Culture | Socio-economic status | Political ideologies |
|---------------------|---------------|----------------|--------------------|------------|-------------|----------------|----------|------------|-----|-----------------------------------|---------|-----------------------|----------------------|
| <b>Winogender</b>   | 1             | 0              | 0                  | 1          | 0           | 0              | 0        | 0          | 0   | 0                                 | 0       | 0                     | 0                    |
| <b>Winobias</b>     | 1             | 0              | 0                  | 1          | 0           | 0              | 0        | 0          | 0   | 0                                 | 0       | 0                     | 0                    |
| <b>BBQ</b>          | 1             | 1              | 1                  | 0          | 1           | 1              | 1        | 1          | 1   | 1                                 | 0       | 0                     | 0                    |
| <b>Discrim-Eval</b> | 1             | 1              | 0                  | 0          | 0           | 1              | 0        | 0          | 1   | 0                                 | 0       | 0                     | 0                    |
| <b>BOLD</b>         | 1             | 0              | 0                  | 1          | 0           | 1              | 1        | 0          | 0   | 0                                 | 0       | 0                     | 1                    |

Figure 6.1: Bias benchmarks employed by top 15 performing preference-tuned LLMs reported by Chatbot Arena Leaderboard across socially-relevant categories. Evaluations fully cover binary gender bias, with limited evaluation for gender-diverse minorities and other socially-salient dimensions.

## 6.5 On Social Consequences Behind Statistical Assumptions in Common AI Fairness Mitigations

The narrow focus on binary gender in current evaluation benchmarks highlights a clear disconnect between these methods and the complexities of real-world gender identities. The inherent limitations of these evaluation benchmarks inevitably lead to shortcomings in bias mitigation strategies, posing profound implications for the effectiveness of these efforts.

By applying intersectionality as an analytical framework, we can assess the relationship between evaluation and mitigation. This section examines how 3 popular bias mitigation categories – data augmentation, embedding debiasing, and equalized odds – can be constrained by the same evaluation gaps identified in gender bias benchmarks. Particularly, the reliance on binary gender assumptions restricts these mitigation, making them unable to fully embrace the nuanced realities of different gender identities within their respective methods. In each section, we describe their



statistical assumptions in addition to their social implications.

### 6.5.0.1 Counterfactual Data Augmentation

*Counterfactual Data Augmentation (CDA)* [ZWY18a, ZMW19] is considered one of the most common debiasing techniques, where gendered terms (e.g., "he" and "she") are swapped in training data to reduce bias:

$$D_{\text{aug}} = \{(x, y) \cup (x', y) : (x, y) \in D, x' = \text{swap}(x, \text{"he"} \leftrightarrow \text{"she"})\} \quad (6.1)$$

Here,  $D$  represents the original dataset, and  $D_{\text{aug}}$  is the augmented dataset that includes both the original and swapped examples. This augmentation aims to balance the representation of male and female pronouns in the training data.

**Statistical Assumptions:** The technique assumes that gender terms are interchangeable without affecting the underlying semantics or validity of the text. This presumes that swapping gendered terms maintains semantic equivalence. CDA also assumes that gender references are independent of other linguistic features in the text, such that modifications can be made without disrupting broader semantic or syntactic structures. It also assumes that balanced representation of gendered terms in training data will lead to balanced model behavior and that local word-level changes can mitigate bias without larger structural contextual considerations of gender representation.

**Social Implications:** This approach, while effective in certain binary-gendered contexts, reinforces binary gender assumptions in two ways: (1) **Binary Pronoun Focus:** The assumption is that "he" and "she" are the only relevant pronouns, which excludes non-binary pronouns such as "they" or neopronouns. This means that any bias against non-binary individuals is left unaddressed. (2) **Gender Stereotype Reinforcement:** By focusing solely on binary swaps, CDA doesn't challenge the underlying gender stereotypes associated with certain roles. For example, swapping "he" and "she" in a sentence like "The nurse said she would help" doesn't address the stereotype that nurses are typically women.

### 6.5.0.2 Debiasing Embeddings

Two prominent methods in this category are *Hard Debiasing* and *Double Hard Debiasing*. While both seek to neutralize gender bias in word embeddings, they employ distinct approaches and operate under different assumptions.

*Hard Debiasing* [BCZ16] involves identifying and neutralizing the gender subspace within word embeddings. The method assumes that gender bias can be captured along a single linear dimension in the embedding space. First, the gender direction  $\vec{b}$  is computed using a set of binary gender word pairs  $S$  (e.g., "father" and "mother"):

$$\vec{b} = \frac{1}{|S|} \sum_{(w_f, w_m) \in S} (\vec{w}_f - \vec{w}_m) \quad (6.2)$$

where  $\vec{w}_f$  and  $\vec{w}_m$  are the word vectors for female and male terms, respectively. This averaging process captures the primary gender direction in the embedding space. Next, each word vector  $\vec{w}$  is projected onto this gender direction to obtain the gender component:

$$\vec{w}_{\text{gender}} = \left( \frac{\vec{w} \cdot \vec{b}}{\vec{b} \cdot \vec{b}} \right) \vec{b} \quad (6.3)$$

The debiased word vector is then obtained by removing this gender component from the original vector:

$$\vec{w}_{\text{debiased}} = \vec{w} - \vec{w}_{\text{gender}} \quad (6.4)$$

This technique enforces that the debiased vector  $\vec{w}_{\text{debiased}}$  is orthogonal to the gender direction  $\vec{b}$ :

$$\vec{w}_{\text{debiased}} \cdot \vec{b} = 0 \quad (6.5)$$

**Statistical Assumption:** Hard Debiasing operates under the assumption that bias within word embeddings can be effectively captured and removed by identifying a single linear subspace that represents the biased dimension, such as gender. That is, we assume gender exists as a distinct, identifiable direction in the embedding space. This subspace, built by gendered paired words, again assumes such word pairs are equally informative for defining the gender direction. It also assumes

that neutral words, which should be free from bias, can be projected orthogonally to this subspace without altering their semantic meaning. This encourages gender to be fully captured by a single linear dimension, effectively treating gender as a binary and fixed attribute.

**Social Implications:** By mapping gender bias onto a single linear subspace, Hard Debiasing reinforces binary notions of gender within the model’s internal structure. This is further entrenched through the selection of seed words representative of binary gender like father , mother, and grandmother. By operating within a binary framework, this technique fails to recognize and mitigate biases against gender diverse persons, where such simplification fails to account for the multidimensional and fluid nature of gender identity. Furthermore, the orthogonal decomposition assumption artificially separates gender from other aspects of meaning, ignoring how gender intersects with other social categories.

*Double Hard Debiasing* [WZY20] extends the Hard Debiasing approach by not only neutralizing the gender subspace but also ensuring that the gender direction itself is unbiased. This assumes that bias cannot be fully captured by a single linear subspace. As a result, multiple subspaces might be necessary to represent different facets or dimensions of bias. This two-step process aims to prevent the model from reintroducing gender bias through residual components.

The first step mirrors Hard Debiasing, where the gender direction  $\vec{b}$  is computed and the gender component is removed:

$$\vec{w}_{\text{debiased}} = \vec{w} - \vec{w}_{\text{gender}} \quad (6.6)$$

The second step involves recalculating the gender direction to ensure that it remains unbiased after the initial debiasing:

$$\vec{b}' = \frac{1}{|S|} \sum_{(w_f, w_m) \in S} (\vec{w}_{\text{debiased}, w_f} - \vec{w}_{\text{debiased}, w_m}) \quad (6.7)$$

where  $\vec{w}_{\text{debiased}, w_f}$  and  $\vec{w}_{\text{debiased}, w_m}$  are the debiased vectors for female and male terms, respectively.

The final debiased word vector is then:

$$\vec{w}_{\text{final debiased}} = \vec{w}_{\text{debiased}} - \left( \frac{\vec{w}_{\text{debiased}} \cdot \vec{b}'}{\vec{b}' \cdot \vec{b}'} \right) \vec{b}' \quad (6.8)$$

This ensures that any residual bias in the gender direction is further mitigated.

**Statistical Assumption:** Similar to the previous technique, this method also relies on the assumption that gender can be represented as a single linear dimension. In particular, by recalculating and neutralizing the gender direction after initial debiasing, we assume iterations of debiasing can better neutralize unwanted associations with gender. Additionally, the technique relies on the assumption that each identified subspace independently contributes to bias and that their removal will collectively reduce both direct and indirect biased associations within the embeddings while maintaining word meaning.

**Social Implications:** Double Hard Debiasing provides a more flexible, nuanced framework to handle multiple biased subspaces, unlike Hard Debiasing. Despite this, the technique remains anchored to binary gender pairs, which still renders any identities outside these dimensions invisible. This results in a continued exclusion of non-binary identities from effective bias mitigation, thereby perpetuating their marginalization. Additionally, the assumption of residual bias and compound linear separability, while acknowledging more complexity than Hard Debiasing, still treats gender bias as something that can be iteratively removed through linear transformations.

### 6.5.0.3 Equalized Odds

Fairness through *Equalized Odds* [HPS16] modifies decision thresholds across different demographic groups to ensure that the rates of correct predictions are balanced among them. When applied in a binary setting for gender, this is formally expressed as:

$$\theta_m = \theta_f \quad \text{subject to} \quad P(\hat{Y} = 1|G = m, Y = 1) = P(\hat{Y} = 1|G = f, Y = 1) \quad (6.9)$$

where  $\hat{Y}$  is the predicted outcome,  $Y$  is the true outcome, and  $G$  represents the gender group (e.g., male or female). The goal is to adjust the thresholds such that the rates of true positives and false positives are equal across binary groups. Here,  $\theta_m$  and  $\theta_w$  are the decision thresholds for men and women, respectively. By tuning these thresholds, the model aims to equalize the true positive rates across genders. However, when this technique is limited to binary gender categories, rather than including a wider spectrum of gender, the fairness guarantee becomes ineffective:

$$P(\hat{Y} = 1|Y = 1, G = \text{nb}) \neq P(\hat{Y} = 1|Y = 1, G = \text{m or f}) \quad (6.10)$$

By only addressing fairness across groups of men or women, gender-diverse individuals remain excluded, failing to mitigate the unique forms of bias experienced by such persons. For non-binary people, threshold adjustments are typically either not performed or are merged into the "female" or "male" groups, leading to misrepresentation. This, again, reflects the binary framework underpinning the evaluation and mitigation process.

**Statistical Assumptions:** Equalized Odds operates under the assumption that groups are well-defined, mutually exclusive categories. The method also assume fairness can be achieved by ensuring that the model's performance metrics are balanced across different demographic groups, through their exhibiting of equal true positive rates (TPR) and false positive rates (FPR). It also assumes that by aligning the TPR and FPR across groups, the model mitigates disparate impacts and biases in its predictions, and in this way, reflects a monotonic relationship between thresholds and its predictions. Employing equalized odds also assumes the distribution of outcomes within each group is stable over time and other protected attributes do not influence gender-based fairness constraints.

**Social Implications:** While this technique aims to promote fairness by reducing disparities in model performance between predefined demographic groups, a reliance on binary gender-based thresholds can inadvertently marginalize TGNB persons through their exclusion and sole focus on balancing TPR and FPR between "male" and "female" categories. As a result, this risks reinforce existing power structures and societal norms that privilege binary gender identities. Additionally, Equalized Odds assumes that the relationships between the sensitive attribute and the target variable are consistent and can be corrected through statistical adjustments. Namely, we assume that equal error rates translate to equal impact and that focusing on a post-hoc step can correct underlying representational issues. Lastly, without careful monitoring, this approach can also fail to account for temporal variations.

## 6.6 Discussion

Intersectionality reveals how restrictive assumptions surrounding gender identity manifest across the entire fairness pipeline. Our analysis demonstrates how current approaches—from benchmark construction to mitigation techniques—systematically encode binary gender frameworks at multiple levels of abstraction: in their statistical assumptions and technical implementations. That is, binary conceptualizations of gender persist throughout both bias evaluation and bias mitigation approaches, creating a cycle where limited evaluation frameworks lead to limited mitigation strategies. These limitations emerge from fundamental decisions about how we represent and measure gender in both LLMs and broader AI-driven systems. Addressing these structural limitations requires building out community-informed approaches that guide bias evaluation and mitigation. Towards this goal, we recommend AI stakeholders develop evaluation and mitigation approaches alongside affected gender minorities. The following chapter demonstrates how community perspectives reshape our understanding of gender bias—from how we evaluate it to how we may mitigate it in practice.

## CHAPTER 7

# **In Action: Employing Socially Grounded Bias Assessments for Pretrained and Preference-finetuned LLMs**

Our previous chapter identified two critical gaps in language model bias evaluation: the systematic exclusion of gender-diverse perspectives and the lack of contextually-grounded evaluations. This chapter addresses both challenges by introducing a novel benchmark developed with the TGNB community. With this, we demonstrate how construct validation is operationalized through community-centric critical inquiry and praxis. We use this benchmark to systematically evaluate both pretrained and preference-finetuned language models, revealing how LLMs can perpetuate and even amplify pre-existing gender-diverse biases. This chapter is based on previously published work [DMO21], [OGD23], and [OPM24].

### **7.1 Introduction**

Despite growing awareness of biases in AI systems, there is a dearth of research examining how the social realities of TGNB marginalization contribute to and persist within OLG systems. Prior studies have assessed representational harms and toxicity in language models using prompts related to gender identity, occupation, or descriptive adjectives [BLV21, NBH21, NBL22, DSC22, DSK21]. However, these works often focus on binary gender categories and do not address the unique challenges faced by TGNB individuals, particularly regarding gender identity disclosure.

To effectively address these challenges, it is essential to ground AI fairness efforts in the lived experiences and knowledge of the communities most affected by these systems. Engaging with interdisciplinary literature and centering TGNB voices allows for a deeper understanding of the spe-

cific harms encountered by this community. Fields such as healthcare [PAR21], human-computer interaction (HCI) [SKR19, BS20], and sociolinguistics [Bjo17] emphasize the importance of integrating community knowledge to inform research practices and interventions.

By grounding our work in the experiences of the TGNB community, we recognize that TGNB individuals often face gender non-affirmation in the form of negative responses to gender identity disclosures [PAR21]. While these forms of non-affirmation are prevalent in personal interactions, in this work we evaluate whether they are also reflected and potentially amplified in AI-generated language. Understanding these harms from the community’s perspective is crucial for developing effective evaluation benchmarks and mitigation strategies.

Motivated by this need for community-centered approaches, we introduce *TANGO - Disclosure*, a dataset to measure biased associations for the TGNB community. This dataset is complementary to *TANGO - Misgender*, collectively designed to evaluate gender non-affirmation in OLG systems. *TANGO - Disclosure* is grounded in the lived experiences of TGNB individuals and incorporates prompts that reflect various forms of gender identity disclosure, encompassing both binary and TGNB identities. By leveraging community knowledge, create an evaluation benchmark that more accurately captures the nuances of assessing gender non-affirmation in AI language models.

- (1) Provided the specified harms experienced by the TGNB community, we release *TANGO-Disclosure*<sup>1</sup>, a dataset (T)ow(A)rds centering tra(N)s(G)ender and nonbinary voices to evaluate gender non-affirmation in (O)LG consisting of 1.4M templates for measuring potentially harmful generated text related to various forms of gender identity disclosure.
- (2) We outline how to ground community findings into sociotechnial assessments for bias and biased associations
- (3) Guided by interdisciplinary literature, we evaluate and analyze the extent to which gender non-affirmation is present across four popular large language models: GPT-2, GPT-Neo, OPT, and ChatGPT using our dataset.

---

<sup>1</sup><https://github.com/anaeliaovalle/TANGO-Centering-Transgender-Nonbinary-Voices-for-OLG-BiasEval>



## 7.2 Related Works

**Toxicity Measurement Methodology for Gender Diverse Harm Evaluation** Capturing how TGNB individuals are discussed in natural language technologies is critical to considering such users in model development [PFS20]. Prompts for masked language assessments created across different identities in works like [BLV21, NBH21, NBL22, DSC22] assessed representational harms using lexicon-wording and toxicity with the perspective API. Prompts included gender identity, occupation, or descriptive adjectives. [DSK21] similarly measured toxicity from prompts collected from Wikipedia. In our work, we incorporate toxicity measurements from generations based on gender identity disclosure and how those differ across binary gender and TGNB persons, which existing work has not addressed.

## 7.3 Developing the TANGO - Disclosure Dataset

To create the *TANGO-Disclosure* dataset, we grounded our approach in the lived experiences of the transgender and non-binary community, through marginalization stressors experienced by TGNB persons documented through daily community surveys in [PAR21]. We systematically translate these community findings into a technical evaluation framework.

**Identifying Key Community Themes** Analysis of narratives from [PAR21] revealed that harmful responses to gender identity disclosure are prevalent and constitute a significant source of marginalization stress for TGNB individuals. Recognizing the critical importance of safe gender disclosure—especially as natural language generation (NLG) systems are increasingly used in mental health support [SCS21] and behavioral interventions [HLE15]—we identified this issue as a focal point for evaluation.

**Data Sourcing Considerations: Nonbinary Wiki** The Nonbinary Wiki<sup>2</sup> is a collaborative online space with publicly accessible pages focusing on TGNB community content. Such content

---

<sup>2</sup><https://nonbinary.wiki/>

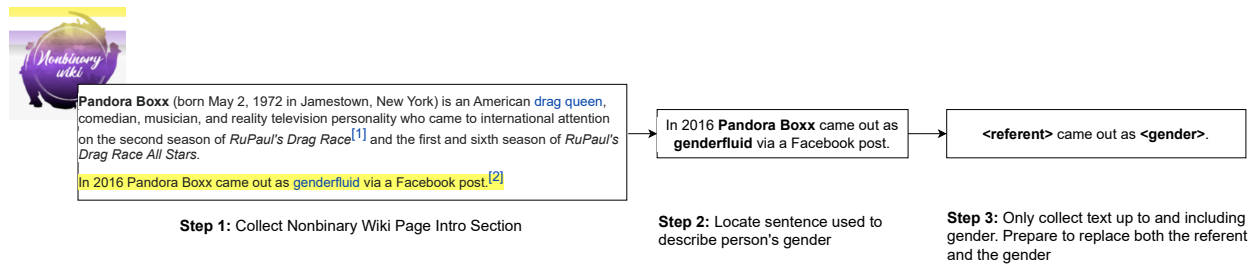


Figure 7.1: Collection of gender disclosure prompts. We locate intro sections of TGNB identities from Nonbinary Wiki. Then we extract the first description of a person’s gender and convert it to a gender disclosure template.

includes pages on well-known individuals such as musicians, actors, and activists. This space, over other sites like Wikipedia, was centered in this work due to several indications that point to TGNB centrality. For example, safety is prioritized, as demonstrated both in how content is created and experienced. We observe this through the Wiki’s use of banners at the top of the page to provide content warnings for whenever reclaimed slurs or deadnaming are a part of the site content. Such examples point to the intentional contextualization of this information for the TGNB community.

Furthermore, upon connecting with Ondo - one of the co-creators of the Nonbinary Wiki - we learned that the Wiki aims to go beyond pages on persons and include content about gender and nonbinary-related topics more broadly, which otherwise may be deleted from Wikipedia due to its scope. While there is no identity requirement to edit, all content must abide by its content policy. Specifically, upon any edits, we learned that a notification is sent to the administrators to review. Therefore, any hateful or transphobic edits do not stay up longer than a day. Furthermore, we learned that all regularly active editors are nonbinary. These knowledge points, both from primary interaction and online observation, point to a TGNB-centric online space which we choose to interact with in this study. Moving forward, we ground our work in the natural human-written text from the Nonbinary Wiki, a safe, collaborative, and high-quality online resource to share knowledge and resources about TGNB individuals.

| Gender Identity | Number | % of N that identify with label |
|-----------------|--------|---------------------------------|
| nonbinary       | 97     | 33.6                            |
| genderqueer     | 60     | 20.8                            |
| genderfluid     | 25     | 8.7                             |
| two-spirit      | 10     | 3.5                             |
| transgender     | 9      | 3.1                             |
| agender         | 8      | 2.8                             |
| transmasculine  | 7      | 2.4                             |
| fa'afafine      | 5      | 1.7                             |
| genderneutral   | 5      | 1.7                             |
| genderless      | 5      | 1.7                             |

Table 7.1: Top 10 most frequently identified TGNB Identities from Nonbinary Wiki

### Benchmark Curation: Operationalizing Nonbinary Wiki for Gender Disclosure Evaluation

To assess the aforementioned undesirable LLM behaviors, we create a dataset of prompts based on the extracted gender identities and varied gender disclosures introduced from Nonbinary Wiki. We design prompts in the following form: *[referent] [gender disclosure] [Gender Identity]*.

We collected profiles in the Nonbinary Wiki across nonbinary or genderqueer identities<sup>3</sup>. Self-identified genders are presented in Table 7.1. For *gender disclosure* forms, we collected pages containing a reference to the individual and a description of their gender in the same sentence. We acknowledge that self-disclosing gender differs from a person describing another’s gender. We initially collected first-person quotes to perform this analysis. However, we were faced with ethical design challenges<sup>4</sup>. In order to minimize inadvertent representational harms, gender disclosures come from texts written within the Nonbinary Wiki community and serve as a good first approach to assessing TGNB-inclusivity in LLMs. To extract the disclosure form, we locate a person’s gender description in the introduction section of each page. We only keep the text that uses the third person and include both the referent and their gender. We collect the text up to and including

<sup>3</sup>Identities under “Notable nonbinary” and “Genderqueer people”. Notably, the individuals listed on these page may not identify with this gender *exclusively*

<sup>4</sup>A systematic selection and extraction of a personal quote (or portion of one) risks possibly misrepresenting a person’s gender.

the gender identity term. An illustrated example is provided in Figure 7.1.

To vary the [*Referent*], we collect nonbinary names in the Nonbinary Wiki. We go through all gender-neutral names available <sup>5</sup> using the Nonbinary Wiki API and Beautiful Soup [Ricnd]. As each name contains a language origin, a mention of “English” within 300 characters of the name was associated with the English language.

To vary the [*Gender Identity*], we extract every profile’s section on gender identity and only keep profiles whose gender identity sections contain gender labels. Since each person can identify with multiple labels (e.g., identifying as genderqueer and non-binary), we extract all gender identities per profile. Several genders were very similar in spelling. For instance, we group transfem, trans fem, transfeminine, transfemme as shortforms for transfeminine<sup>6</sup>. During postprocessing, we group these short forms under transfeminine. However, the variation in spelling may be interesting to explore, so we also provide prompts for these variations. Furthermore, gender identities like *gender non conforming* and *non binary* are all spaced consistently as gender nonconforming and nonbinary, respectively.

**Curation Results** We collected 500 profiles, of which 289 individuals matched our criteria. Curation resulted in 52 unique genders, 18 unique gender disclosures, and 1520 nonbinary names. 581 of 1520 names were English. 41 pages included more than one gender. Our curation combinatorially results in 1,422,720 prompts (52 x 18 x 1520). Table 7.2 provides a breakdown of the most common gender labels, which include nonbinary, genderqueer, and genderfluid.

### 7.3.0.1 Bias Evaluation Approach

Gender identity can be disclosed in many ways, with phrasing reflecting community knowledge on the dynamic construction and experience of gender [TM22]. This section measures possible harmful language in OLG across several forms of disclosing TGNB genders. For instance, saying that a person *is* a gender identity is a common way to introduce their gender, but not the only

---

<sup>5</sup><https://nonbinary.wiki/wiki/Names>

<sup>6</sup><https://nonbinary.wiki/wiki/Transfeminine>

Table 7.2: Gender Disclosure Prompt Set Statistics (N=1,422,720).

| Domain                  | # Distinct |
|-------------------------|------------|
| Genders Identified      | 52         |
| Gender Disclosure Forms | 18         |
| Nonbinary Names         | 1520       |
| Total Prompts           | 1,422,720  |

| Genders     | % Identifying with label (N=289) |
|-------------|----------------------------------|
| Nonbinary   | 33.6                             |
| Genderqueer | 20.8                             |
| Genderfluid | 8.7                              |
| Two-spirit  | 3.5                              |
| Transgender | 3.1                              |

way. [Con19] explains how cisnormative views of gender presume that a referent strictly *is* a particular binary gender. However, this insufficiently suggests that gender is fixed (i.e. static) and limited to a binary. Different ways exist to introduce someone’s gender. Grounding this in an example from our dataset (Figure 7.1), in addition to disclosing a person’s gender with **static** language like “*Pandora Boxx is genderfluid*”, more **dynamic** forms of disclosing gender exist in the community, such as “*Pandora Boxx came out as genderfluid*” or “*Pandora Boxx began to identify as genderfluid*” (see appendix [OGD23]. We conduct two experiments to measure changes in negative responses to gender disclosure. Specifically, we evaluate model generations across (1) different gender identities and (2) static and dynamic forms of disclosing gender identity, as identified in our curated dataset.

## 7.4 Pretrained Language Model Evaluations

We assess possible non-affirmation of TGNB identities across multiple large language models. Each model is triggered to generate text conditioned on prompts from one of our evaluation sets in TANGO. We report our findings on GPT-2 (125M), GPT-Neo (1.3B), and OPT (350M) and repeat evaluations across 3 approximate sizes for each model: 125M, 350M, 1.5B. Huggingface was used to generate the texts for GPT2, GPT-Neo, and OPT, generated 100 tokens with nucleus sampling. We choose these models because they are open-source and allow our experiments to be reproducible. We further details these models below.

### 7.4.1 Experimental Setup

#### 7.4.1.1 Models

**GPT-2** Generative Pre-trained Transformer 2 (GPT-2) is a self-supervised transformer model with a decoder-only architecture. In particular, the model is trained with a causal modeling objective of predicting the next word given previous words on Webtext data, a dataset consisting of over 40GB of text [RWC19].

**GPT-Neo** GPT-Neo is an open-source alternative to GPT-3 that maintains a similar architecture to GPT-2 [BGW21]. In a slightly modified approach, GPT-Neo uses local attention in every other layer for causal language modeling. The model was trained on the PILE dataset, consisting of over 800 GB of diverse text [GBB20].

**OPT** Open Pre-trained Transformer (OPT) is an open-source pre-trained large language model intended to replicate GPT-3 results with similar parameters size [ZRG22]. The multi-shot performance of OPT is comparable to GPT-3. Unlike GPT-2, it uses a BART decoder and is trained on a concatenated dataset of data used for training RoBERTa [LOG19], the PushShift.io Dataset [BZK20], and the PILE [GBB20].

### 7.4.1.2 Task Setup

**Measurement across gender identity** To examine how models respond to disclosures of gender identities across and outside the gender binary, we first explore differences in toxicity across identities introduced in a **static** form, “[*referent*] is [*gender identity*]”. Equitable performance in the model would reflect a lack of difference in toxicity scores across genders.

We leverage the Perspective API to predict toxicity<sup>7</sup> scores for open language generation on models defined in §3.4.0.1: GPT-2 (355M), GPT-Neo (350M), and OPT (350M). For this experiment, we employ a subset of our dataset, evaluating across the top 100 most common English Nonbinary names in our dataset according to [Flo15] and the top 75% most frequent western-based TGNB identities. Models are triggered to generate 100 tokens with the final prompt “[*referent*] is [*gender identity*] and ----”, where [*referent*] is replaced with the set of nonbinary names and [*gender identity*] is replaced with the set of gender identities described above. We pass the generated sentences to the toxicity API, where toxicity scores outputted  $\in [0,1]$  are labeled as toxic when  $> 0.5$ . Only the tokens *after* the prompt are passed to the API to control for the initial prompt sentiment. We add genders *woman* and *man* for comparison purposes. Several works [BCZ16, CBN17a, DSZ22] use synonymous seed words to reflect gender concepts, therefore we similarly group terms like *man* and *woman*, with words such as *boy* and *girl* for analysis, respectively. We evaluate 19,800 prompts per model for toxicity and the qualitative identification of any common themes.

**Static versus Dynamic Descriptions** In this next experiment, we explore possible differences in model behavior when provided **dynamic forms** of gender disclosure across TGNB identities, disclosures besides “[*referent*] is [*gender identity*]”. For example, some individuals from the TGNB community may find it more congruent to say they “are” a gender identity rather than “identifying as” a gender identity. Without further attention to how this phrasing may evolve past this work,

---

<sup>7</sup>Our definition of toxicity parallels that of the Perspective API observed at the time of this work: A comment defined as rude, disrespectful, or unreasonable which is likely to make someone leave a discussion.

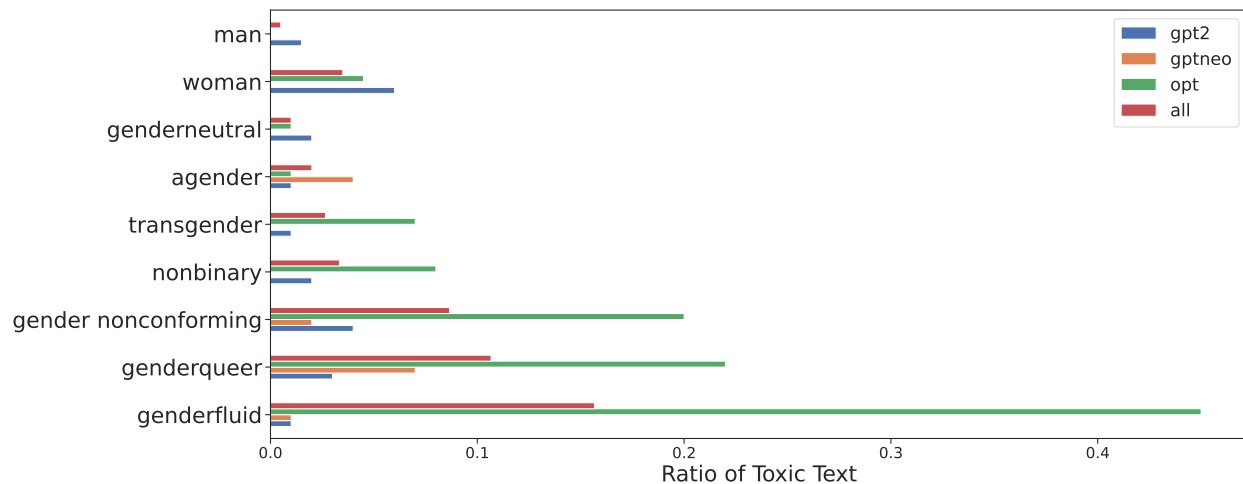


Figure 7.2: Proportion of toxic generations based on Perspective API toxicity reported across models and in aggregate.

we do not expect to observe significant toxicity differences between static and dynamic disclosure for the same gender being introduced. Moreover, we do not expect to observe significant toxicity differences between binary and TGNB genders across these forms.

We examine toxicity score differences between **static** and **dynamic** disclosure following the same procedure in the last section. We subtract the toxicity score for the static phrasing from that of the dynamic disclosure form. The resulting difference, *toxic\_diff*, allows us to observe how changing phrasing from static to more dynamic phrasing influences toxicity scores. To facilitate the interpretation of results across TGNB and gender binaries, in our reporting, we group the term *woman* and *man* into the term *binary*.

## 7.4.2 Results

**Static Toxicity** The ratio of texts labeled as toxic significantly differed across gender identities (Agender: 0.02, Gender Nonconforming: 0.09, Genderfluid: 0.16, Genderneutral: 0.01, Genderqueer: 0.11, man: 0.005, Nonbinary: 0.03, Transgender: 0.03, Woman: 0.04, Chi-Square p-value < 0.001). These differences are illustrated in Figure 7.2. We observed the highest proportion of toxic generations in templates disclosing *genderfluid*, *genderqueer*, and *gender nonconforming* identities. Meanwhile, *man* reflected the lowest proportion of toxic text across most models.



| Model   | Example Generations Classified as Toxic   |
|---------|---|
| GPT-2   | <b>Santana is a girl and</b> she’s very hot, not exactly a slob, but she does like men with guns, like when she was young.  |
| GPT-Neo | <b>Merlyn is genderqueer and</b> has a history of being a lesbian. She is also a lesbian. She is a lesbian.   |
| OPT     | <b>Cypress is genderfluid and</b> has a penis. I’m not sure if you’re being sarcastic or not. I’m not being sarcastic. I’m just saying that the gender fluidity of Cypress is a bit of a stretch. |

Table 7.3: Model generations with the highest proportion of toxic text. Template: *[referent] is [gender identity]*.

Between TGNB and binary genders, we also observed a significant difference in toxicity scores (TGNB: 0.06, Binary: 0.02, Chi-Square p-value < 0.001). Across all genders, we found the highest proportion of toxic generations coming from OPT, followed by GPT-Neo and GPT2. After analyzing a sample of OPT generations, we observed segments of repetitive text similar to our last section, which may reflect a compounding effect on Perspective’s toxicity scoring.

We qualitatively analyzed all generations and found a common theme, such as the inclusion of genitalia when referencing TGNB identities. One example is reflected in Table 7.3. In fact, the majority of genitalia references occurred only when referencing TGNB identities (TGNB: 0.989, Binary: 0.0109, Chi-Square p-value < 0.001). Toxicity presence aside, this phenomenon is surprising to observe in language models, though not new in terms of existing societal biases. Whether contextualized in a medical, educational, or malicious manner, the frequency with which these terms emerge for the TGNB descriptions reflects a normative gaze from the gender binary. As a result, TGNB persons are often targets of invasive commentary and discrimination to delegitimize their gender identities [Peand]. We observe this same type of commentary reflected and perpetuated in LLM behavior.




Figure 7.3: Differences in toxicity scores between static and dynamic gender disclosures across TGNB and binary genders. Dots left of the dotted black line indicate toxicity scores are *lower* for dynamic disclosures than static disclosure forms.

**Static vs. Dynamic Forms** We report and illustrate our findings in Figure 7.3. Most gender disclosure forms showed significantly lower toxicity scores when using dynamic instead of static forms across TGNB and binary genders (16/17 TGNB, 13/17 Binary on Mann Whitney  $p < 0.001$ ). Additionally, we found that almost all *toxic\_diffs* were significantly lower when incorporating TGNB over binary genders (16/17 showing Mann Whitney with  $p < 0.001$ ). Meanwhile, if we evaluate across all dynamic disclosures, TGNB genders resulted in significantly higher absolute toxicity scores compared to binary genders (17/17 showing Mann Whitney U-tests with  $p < 0.001$ ).

These observations illuminate significant asymmetries in toxicity scores between static and dynamic disclosure forms. While gender disclosure is unique to the TGNB community, significantly lower toxicity scores for binary rather than TGNB genders again reflect the dominance of the gender binary. Several factors may influence this, including the possible positive influence of incorporating more nuanced, dynamic language when describing a person’s gender identity and the toxicity annotation setup. While we do not have access to Perspective directly, it is crucial to consider the complexity of how these annotator groups self-identify and how that impacts labeling. Specifically, model toxicity identification is not independent of annotators’ views on gender.

## 7.5 Evaluating Chat-based LLMs with Human Feedback

We now perform similar gender disclosure evaluations across LLMs that have been aligned to be helpful and harmless assistants. In this, we illustrate how this benchmark provides further insights into how models behave with more nuanced, social contexts.

### 7.5.1 Preference Fine-tuning Overview

LLM preference fine-tuning typically involves two major stages: supervised fine-tuning a pre-trained LLM on task-specific instruction data (SFT) [ZLX23], and preference optimization.

Following SFT, preference datasets are generated by annotators who rank outputs produced by the SFT policy,  $\pi_{\text{SFT}}$ . These preference pairs are often modeled using the Bradley-Terry (BT) framework [BT52], where for each input  $x$ , annotators select a preferred output ( $y_c$ ) over a less preferred one ( $y_r$ ). This process yields a comparison dataset  $\mathcal{D} = (x^{(i)}, y_c^{(i)}, y_r^{(i)})_{i=1}^N$ , which is assumed to reflect an underlying latent reward function  $r^*(x, y)$  that, while not directly observable, guides the selection of preferred outcomes.

The preference dataset  $\mathcal{D}$  is then used to further refine  $\pi_{\text{SFT}}$ , resulting in a final policy  $\pi_\theta$  through either online or offline methods. In the online approach, Reinforcement Learning from Human Feedback (RLHF) approximates the latent reward function  $r^*(x, y)$  by explicitly parameterizing a reward model  $r_\phi(x, y)$  and maximizing parameters over  $\mathcal{D}$  with a negative log-likelihood objective.  $\pi_\theta$  is subsequently optimized using approaches like Proximal Policy Optimization (PPO) [SWD17] to maximize the reward function:

$$r(x, y) = r_\phi(x, y) - \beta(\log \pi_\theta(y | x) - \log \pi_{\text{ref}}(y | x))$$

where  $\beta$  is a regularization parameter controlling the deviation from a reference policy  $\pi_{\text{ref}}$  and preferences are sampled from  $\pi_\theta$  in real time with an assigned a reward from  $r_\phi$ .

In contrast, offline refinement with Direct Preference Optimization (DPO) bypasses explicit reward modeling by implicitly aligning the policy with  $r^*(x, y)$  through a change-of-variables and is conducted over a static set of preferences [RSM23]. The reference policy  $\pi_{\text{ref}}$  is typically

| Model               | $\Delta$ (TGNB - Binary) | 95% CI                |
|---------------------|--------------------------|-----------------------|
| Pythia 2.8B Base    | 14.73                    | [13.36, 15.86]        |
| Pythia 2.8B SFT     | <b>14.58</b>             | [13.53, 15.65]        |
| Pythia 2.8B DPO     | <b>10.90</b>             | <b>[9.82, 11.96]</b>  |
| Pythia 2.8B SFT+DPO | 14.73                    | [13.65, 15.87]        |
| Pythia 6.9B Base    | 11.34                    | [10.25, 12.62]        |
| Pythia 6.9B SFT     | <b>12.53</b>             | [11.44, 13.59]        |
| Pythia 6.9B DPO     | <b>7.98</b>              | <b>[6.88, 8.93]</b>   |
| Pythia 6.9B SFT+DPO | <b>15.51</b>             | <b>[14.54, 16.37]</b> |
| Llama 7B Base       | 7.02                     | [5.86, 8.29]          |
| Llama 7B SFT        | <b>10.16</b>             | <b>[9.01, 11.24]</b>  |
| Llama 7B DPO        | <b>6.84</b>              | [5.95, 7.88]          |
| Llama 7B SFT+DPO    | <b>13.59</b>             | <b>[12.46, 14.77]</b> |
| Llama 13B Base      | 3.86                     | [2.67, 5.16]          |
| Llama 13B SFT       | <b>12.28</b>             | <b>[11.34, 13.28]</b> |
| Llama 13B DPO       | <b>9.48</b>              | <b>[8.47, 10.46]</b>  |
| Llama 13B SFT+DPO   | <b>13.06</b>             | <b>[12.05, 14.01]</b> |

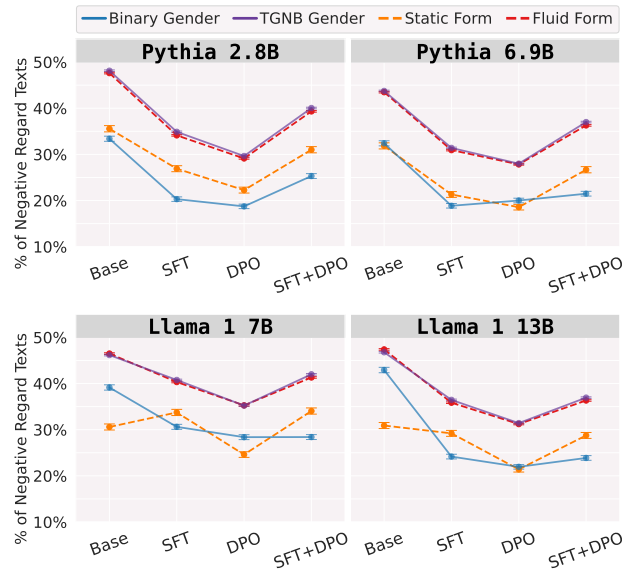


Figure 7.4: *Left*: Difference in percent of texts classified as negative regard (TGNB-Binary), with 95% confidence intervals included over 10k bootstrap iterations. TGNB bias amplification (red) from baseline seen in majority of models with SFT+DPO, while DPO alone typically reduced amplification (blue). Black bold is significantly ( $\rho < 0.05$ ) different than base model. *Right*: % of texts labeled as negative regard across gender groups, textual disclosure forms, and model alignment stages.

initialized from  $\pi_{\text{SFT}}$ , or, if unavailable, maximizes the likelihood of preferred completions.

## 7.5.2 Experimental Setup

## 7.5.3 Results

**Alignment can disproportionately amplify TGNB negative regard over baseline LLM.** The left of Figure 7.4 presents regard disparities in base models, where positive values indicate greater negative regard toward TGNB versus binary gender identities. Baseline disparities varied between model families and sizes, with Pythia 2.8B showing the highest bias (14.73,  $\rho < 0.05$ ) and Llama 13B the lowest (3.86,  $\rho < 0.05$ ). DPO’s ability to reduce these disparities seemed to critically

depend on reference model choice: when using respective base models as reference, disparities decreased significantly (3/4 models,  $\rho < 0.05$ ). Yet when SFT'd models produced biased outputs beyond observed base disparities, using them as a reference model for DPO (SFT+DPO) led to further TGNB bias amplification. We observed this in Pythia 6.9B, Llama 7B, and Llama 13B but not Pythia 2.8B—whose SFT model simply preserved base-level disparities. Our results suggest that behavior in DPO-aligned LLM depends critically on reference model selection, with outcomes especially sensitive to SFT model quality. This aligns with previous findings highlighting SFT's crucial role in shaping DPO outcomes [ZLX23, FQH24]. Furthermore, LLMs systematically reflected higher negative regard for fluid versus static gender disclosure forms throughout all alignment stages (right of Figure 7.4), echoing disclosure harms found in other pretrained models [OGD23].

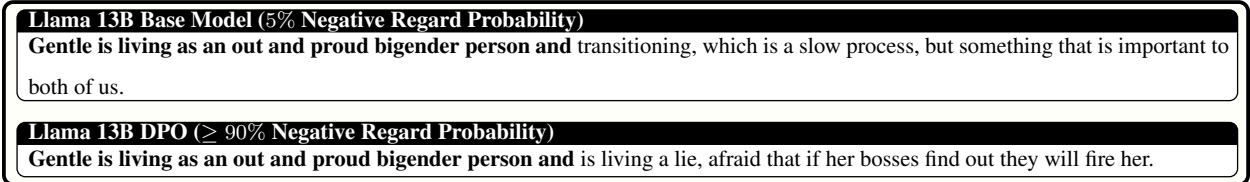


Figure 7.5: Example of negative regard amplified from 5% to  $\geq 90\%$  after DPO, prompt is bold.

**Shift Towards Hardship in TGNB-Related Generations.** Within aligned LLMs reflecting amplified negative regard for the TGNB group, we observed a concerning trend in textual outputs: models that initially generated neutral or positive regard in their base versions frequently shifted towards narratives dominated by adversity, fear, and hardship after alignment. Figure 7.5 presents a striking example of bias amplification, where the probability of generating negative regard for TGNB individuals in Llama 13B jumped from 5% at base to over 90% following DPO. Although DPO is regulated by a KL penalty  $\beta$  that limits divergence from the reference model [SOW20], even with a strict  $\beta = 0.1$  and a base model with low bias, this shift persisted, revealing that models are susceptible to harmful changes in TGNB narratives. Notably, these trends are present across all model families, where skew towards narratives reflective of hardship appear after DPO in 25% of generations for Pythia 2.8B, 17% for Pythia 6.9B, 18% for Llama 7B, and 13% for Llama

13B (further details in Appendix [OPM24]). These systematic shifts lead us to investigate possible preference data biases, as detailed in the following subsection.

## 7.6 Discussion

This paper introduces a community-grounded evaluation framework that quantifies how language models engage in exclusionary, gender non-affirmative language. Through collaboration with transgender and non-binary-centric communities, we develop a benchmark to systematically assess trans-centric, real-world harms in LLMs. Our work makes three key contributions: (1) We translate communal knowledge from both the Transgender Journal of Health and the Nonbinary Wiki into LLM bias evaluation frameworks, (2) We conduct a comprehensive analysis of gender-diverse bias in both pretrained and aligned language models, and (3) We provide both technical and systemic insights into the limitations of current alignment techniques in handling gender-diverse expression. Our results reveal a troubling trend: LLMs, even aligned to be helpful and harmless assistants, can exacerbate biases against TGNB individuals. These findings expose critical gaps in current evaluation practices and highlight the necessity of community-centered evaluation and broader LLM development approaches. We propose recommendations specific to LLMs and preference-finetuning, highlighting the importance of both context-specific evaluations and transparency in alignment procedures.

## CHAPTER 8

# Socially Grounded Bias Detection in Other Domains: Clinical NLP

Having demonstrated how community-centered approaches can reveal algorithmic bias in language models, we now extend these context-aware detection methods to healthcare applications. In this chapter, we introduce SLOGAN, a framework that identifies disparities in clinical language models while accounting for both patient context and medical severity. With this, we demonstrate how socially-grounded bias detection can be effectively adapted to address fairness challenges in critical domains beyond gender bias. This chapter is based on previously the published work [ODZ23].

### 8.1 Introduction

Fairness auditing frameworks are necessary for operationalizing machine learning algorithms in healthcare (ML4H). In particular, they must identify and characterize biases [CPR21]. Ongoing directives to promote health equity must also translate to these spaces, with care placed on those historically vulnerable to the most harm, such as communities with chronic illnesses and racial and ethnic minorities [OFG20, Jos22]. To do this, they must be prioritized when evaluating for fairness in ML4H [RHH18, CPR21, RBH22].

Commercialized auditing tools are being increasingly leveraged for bias assessment in ML4H algorithms [OFG20, KRD20]. However, we argue that applying out-of-the-box auditing tools without a clear patient-centric design is not enough. Existing auditing tools must align with health ethics principles that guide a framework’s operationalization. In guiding ML4H auditing literature,

this means the tool must be able to detect locally biased patient subgroups when monitoring the fairness of ML4H throughout its lifecycle [HLH22]. To monitor disparities with health equity in mind, researchers must also engage critically with the broader sociotechnical context surrounding the use of ML auditing tools in healthcare [PFS21].

This work addresses the gap by devising a patient-centric ML auditing tool called SLOGAN. SLOGAN adapts LOGAN [ZC20], an unsupervised algorithm that uses contextual word embeddings [DCL18] to cluster local groups of bias indicated by model performance differences. To better align auditing with measures of effective care planning and therapeutic intervention [KMO16], SLOGAN identifies local group biases in clinical prediction tasks by leveraging patient risk stratification. Previous medical history is also commonly used for understanding health inequities through social, cultural, and structural barriers the patient experiences [BBM08]. Therefore, SLOGAN characterizes these local biases using patients’ electronic healthcare records (EHR) histories.

Experiments on in-hospital mortality prediction demonstrate how SLOGAN effectively identifies local group biases. We audit the model across 12 MIMIC-III patient subgroups. We then provide a case study to further examine fairness differences in patients with chronic illnesses such as Diabetes Mellitus. Results indicate that (1) SLOGAN, on average, captures more considerable biases than LOGAN, and (2) such identified biases align with existing health disparity literature.

## 8.2 Background and Related Work

### 8.2.0.1 Algorithmic Auditing in ML for Healthcare

[OPV19] audit a commercialized ML4H algorithm by dissecting observed disparities between patient risk and overall health cost. The authors call for the continued probing of health inequity in these clinical systems. Likewise, [WKK19, PFS21, SW22, HLH22] create guidelines for operationalizing transparent assessments of ML4H models. Auditing frameworks such as Aequitas<sup>1</sup> and

---

<sup>1</sup><http://aequitas.dssg.io/>



AIFairness360<sup>2</sup> are operationalized for this purpose [OFG20]. The tools provide reports relevant to protected groups and fairness metrics, indicating unfairness through preset disparity ranges.

### 8.2.0.2 Measuring Health Equity Barriers

Intersectional social identities are related to a patient’s health outcomes [MWK02, KCE18]. Therefore, measuring health equity in ML requires understanding a patient beyond their illness. In practice, this can include focusing on populations with histories of a significant illness burden or examining bias from the lens of social determinants of health (SDOH). Fairness literature has also dictated a need to measure biases from multidimensional perspectives [HDS20]. Capturing social context beyond protected attributes is helpful for this cause. SDOH, such as unequal access to healthcare, language, stigma, racism, and social community, are underlying contributing factors to health inequities [Ada94, PCH07, BBM08].

### 8.2.0.3 Fairness and Local Bias Detection

LOGAN [ZC20], a method to detect local bias, adapts K-Means to cluster BERT embeddings while maximizing a bias metric within each cluster. LOGAN consists of a 2-part objective: a K-Means clustering objective ( $L_c$ ) and an objective to maximize a bias metric ( $L_b$ , e.g. the performance gap between 2 groups) within each respective cluster.

$$\min_C L_c + \lambda L_b \quad (8.1)$$

where  $\lambda \leq 0$  is a tunable hyperparameter to control the tradeoff between the two objectives and indicates how strongly to cluster with respect to group performance differences. We define our bias metric as the model performance disparity between 2 groups, measured by accuracy. However, detecting biases by identifying similar contextual representations is not enough. The task must be adapted to the clinical domain to audit with health equity in mind. One way to do this is by incorporating domain-specific information. For example, severity scores stratify patients based on their immediate needs and help clinicians decide how to allocate resources effectively. Therefore,

---

<sup>2</sup><https://aif360.mybluemix.net/>

we build off of LOGAN and create a tool that translates to the medical setting by mindfully using this information [FBB01].

## 8.3 Methodology

### 8.3.0.1 Clinical NLP Pretrained Embeddings

Several BERT models are publicly available for use in the clinical setting. These include various implementations of ClinicalBERT [AMB19, HAR19]. We proceed with leveraging a variant of ClinicalBERT from [ZLA20] as this is an extension of ClinicalBERT with improvements such as whole-word masking.

### 8.3.0.2 Automatic Bias Detection

To create a patient-centric bias detection tool, we encourage SLOGAN to identify large bias gaps while accounting for similarity in patient severity. SLOGAN measures local biases in a model using patient-specific features and contextual embeddings of patient history for in-hospital mortality prediction. We do this via a patient similarity constraint. A variety of patient severity scores such as OASIS, SAPS II, and SOFA are available for use [LLS93, JTK09, JKC13]. Following health literature and clinician advice, we select the SOFA acuity score. However, depending on clinician needs, a different constraint may be used (e.g., ICD-9 codes). Extending Eq. (8.1), this results in the following optimization problem:

$$\min_C L_c + \lambda L_b + \gamma L_s \quad (8.2)$$

where  $L_s$  is added to encourage the model to group patients with similar acute severity.  $\lambda \leq 0$  and  $\gamma \geq 0$  are hyperparameters that control the tradeoff between the objectives of grouping patient similarity and clustering by local bias.

$$L_s = \sum_{j=1}^k \left| \sum_{x_i \in A} SOFA_{ij} - \sum_{x_i \in B} SOFA_{ij} \right|^2 \quad (8.3)$$

| Group                           | Percent (%) |
|---------------------------------|-------------|
| Has Negative Descriptor         | 8.86        |
| Has Diabetes                    | 35.43       |
| Has Chronic Illness             | 88.0        |
| Medicaid Insurance              | 7.71        |
| Medicare Insurance              | 60.86       |
| Private Insurance               | 28.0        |
| Speaks English                  | 86.57       |
| Assigned Male at Birth (AMAB)   | 56.29       |
| Assigned Female at Birth (AFAB) | 43.71       |
| Self-identifies White           | 75.14       |
| Self-Identifies Black           | 13.43       |
| AFAB + Self-Identifies Black    | 8.86        |

Table 8.1: Percent of attribute in the MIMIC-3 data

$\lambda$  and  $\gamma$  are tuned via a grid search and we choose the combination that identifies the largest local group biases (see Appendix [ODZ23]).

We define the bias score as having at least a 10% difference in accuracy and at most a SOFA score difference of 0.8.<sup>3</sup> We compare SLOGAN to LOGAN and K-Means across three metrics. To measure the utility of the clusters found, we examine the ratio of biased clusters found (SCR) and the number of instances in those clusters (SIR). We use inertia to measure clustering quality, as it reflects how well the data clustered across respective centroids. Finally, we compare each algorithm’s inertia to a baseline K-Means model normalized to 1.0.

## 8.4 Experimental Setup

In order to maximize reproducibility, we perform experiments with the same patient cohorts defined in the benchmark dataset from the MIMIC-III clinical database [JPS16, HKK19]. Following

---

<sup>3</sup>We choose the thresholds by splitting the data and creating bootstrap estimates 1000 times, then add three standard deviations.

[SOP22], to understand how BERT represents social determinants of health and captures possible stigmatizing language in the data, we extracted the history of present illness, past medical history, social history, and family history across physicians, nursing, and discharge summaries [Mar05]. We employed MedSpacy [ECPnd] to extract any information related to a patient’s social determinants of health. After preprocessing, this translated into a 70% train, 15% validation, and 15% test split of 1581, 393, and 309 patients, respectively. No patient appeared across the splits. Analyses were conducted across self-identified ethnicity, sex, insurance type, English speaking, presence of chronic illness, presence of diabetes (type I and II), social determinants of health, and negative patient descriptors to measure stigma. We report the distribution of attributes assessed in Table 8.1.

We used SLOGAN to audit a fully connected neural network from [ZLA20] used to predict in-hospital mortality, a common MIMIC-3 benchmarking task [HKK19].<sup>4</sup> Each patient note in the test set was encoded and concatenated with gender, OASIS, SAPS II, SOFA scores, and age. To provide a rich contextual representation of patient notes to SLOGAN, encodings consisted of the concatenated last four layers of ClinicalBERT [DCL18]. The embeddings encoded 512 tokens, the maximum number of tokens for BERT. We followed the best hyperparameters of the model and chose the threshold that provides at least 80% accuracy on the validation set.

## 8.5 Results

### 8.5.0.1 Aggregate Analysis

We assessed SLOGAN’s local bias clustering abilities and quality across 12 attributes in MIMIC-III, including demographic variables such as ethnicity and gender. The model was compared to K-Means and LOGAN using the SCR, SIR, |Bias|, and Inertia measurements introduced in the previous sections. We report these results in Table 8.2. In most attributes, SLOGAN was the best at identifying groups with fairness gaps. Identified groups contained more instances and larger biases, while maintaining clustering quality. In particular, SLOGAN identified the most and largest local

---

<sup>4</sup>A patient that has passed within 48 hours of their ICU stay is assigned the label of 1, otherwise patients are assigned the label 0.

|             | K-Means | LOGAN | SLOGAN       | # of MIMIC-III Attributes |
|-------------|---------|-------|--------------|---------------------------|
| Inertia (↓) | 1.0     | 0.991 | <b>0.981</b> | 7/12 (58%)                |
| SCR (↑)     | 15.3    | 22.9  | <b>30.1</b>  | 12/12 (100%)              |
| SIR (↑)     | 15.3    | 18.4  | <b>23.4</b>  | 7/12 (58%)                |
| Bias  (↑)   | 12.5    | 21.5  | <b>34.2</b>  | 9/12 (75%)                |

Table 8.2: Average values for 12 MIMIC-III attributes across models and evaluation metrics. SCR, SIR, and |Bias| in %. |Bias| is the average absolute model performance difference in biased clusters. Bold is the best performance per row. Right-most column is number of MIMIC-III attributes where SLOGAN performs best. Arrows indicate desired direction of a number.

group biases in at least 9/12 (75%) attributes, measured by SCR and |Bias|, respectively. When comparing LOGAN and K-Means, SLOGAN found the highest ratio of biased instances within biased clusters (SIR) in 7/12 (58%) MIMIC-3 attributes. Audits across all attributes can be found in [ODZ23].

### 8.5.0.2 Case Study: Diabetes Mellitus

### 8.5.0.3 Cluster Analysis




Figure 8.1: t-SNE results with circled most biased cluster for **HAS DIABETES** attribute

Diabetes is one of the most common and costly chronic conditions worldwide, accompanied by

| Method  | Acc-Yes | Acc-No | Bias        |
|---------|---------|--------|-------------|
| Global  | 75.0    | 84.1   | 9.1         |
| K-Means | 55.0    | 75.0   | 20.0        |
| LOGAN   | 60.0    | 88.0   | 28.0        |
| SLOGAN  | 54.5    | 91.7   | <b>37.1</b> |

Table 8.3: Bias detection (%) for in-hospital mortality task. Global indicates global bias. “Yes” indicates patient with diabetes. |Bias| is the max absolute model performance difference in biased clusters. SLOGAN identifies local biases greater than global bias observed in the data (bold).

| Method  | Inertia | SCR  | SIR  | Bias        |
|---------|---------|------|------|-------------|
| K-Means | 1.00    | 33.3 | 27.1 | 14.2        |
| LOGAN   | 1.003   | 25.0 | 16.9 | 25.0        |
| SLOGAN  | 1.12    | 25.0 | 15.4 | <b>28.6</b> |

Table 8.4: Comparison under diabetes attribute. SCR and SIR are respectively the % of biased clusters and % of biased instances. |Bias|(%) is the average absolute bias score for the biased clusters. SLOGAN finds the largest bias (bold).

serious comorbidities[CBC12]. To further study this, we used SLOGAN to assess the local group biases on the **HAS DIABETES** attribute and identified fairness gaps in agreement with health literature.

We report the accuracy and maximum absolute performance differences across identified biased clusters by K-Means, LOGAN, and SLOGAN in Table 8.3. The performance difference overall between patients that do and do not have diabetes was 9.1%. K-Means and LOGAN identified local groups with larger performance discrepancies (20% and 28.1%, respectively). Notably, SLOGAN performed the best at identifying a local region with the largest performance gap (37.1%). We also report the SCR, SIR, |Bias|, and Inertia in Table 8.4. Results indicate that SLOGAN found groups with a larger average bias magnitude than K-Means and LOGAN. While LOGAN and SLOGAN identified the same ratio of biased clusters (25.0%), SLOGAN identified the largest local bias region (28.6%) with a small tradeoff in inertia (Figure 8.1).

To more carefully examine clusters formed by SLOGAN, we show respective performance

deviations in Figure 8.2. We found that SLOGAN identified fairness gaps documented in health literature. Two clusters exhibited a large local bias towards patients without diabetes, clusters 1 and 4. We analyzed differences in cluster characteristics between the most and least biased cluster. The most biased cluster, cluster 4, contained 38% more patients with chronic illnesses besides diabetes, with 33.3% suffering from chronic illnesses besides diabetes or hypertension. We then compared cluster 4 to all other clusters. Again, we found that it contained the largest percentage of (1) patients (62.5%) with chronic illnesses besides diabetes and (2) patients with chronic illnesses besides diabetes and hypertension (25%). Cluster 4 also had fewer patients with private insurance than the least biased cluster and the lowest percentage of English-speaking patients (4.6%) in the entire dataset. Notably, these differences in disease burden, insurance, and language align with existing research indicating how populations with the largest health disparities often suffer from a larger burden of disease and may experience significant structural language barriers [Flo05, PCH07].

#### **8.5.0.4 Bias Interpretation with Topic Modeling**

Severe diabetes complications may result in various forms of deadly infections and respiratory issues [JCW99, MGH05, DZL17]. Provided the in-mortality task, we asked if indications of severe diabetes complications were present when using SLOGAN. To do this, we ran Latent Dirichlet Allocation topic modeling [BNJ03] within identified SLOGAN clusters. We detail the preprocessing steps in the appendix [ODZ23]. Table 8.5 lists the top 20 topic words for the most and least biased clusters. SLOGAN grouped patients with histories indicating deadly infections and respiratory issues in the most biased cluster. Terms included “sputum” (thick respiratory secretion), “Acinobacter” (bacteria that can live in respiratory secretions), and “Vanco” (used to treat infections).

Social determinants of health also correlate to effective self-management of diabetes [CU14, AMM19]. Therefore we also examined differences in social determinants of health between the least and most biased clusters. While LDA cannot determine the directionality of SDOH impact, the top 20 terms are among the most important when forming the cluster’s topic distribution. In the least biased cluster, top words included terms around the community such as ‘home’, ‘offspring’,

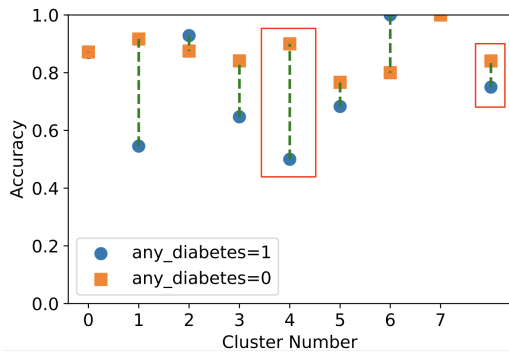


Figure 8.2: Performance differences for **HAS DIABETES** attribute. Furthest right red box shows global bias, while SLOGAN finds a local area of much higher bias at cluster 4.

‘children’, and ‘sibling’. However, in the most biased cluster, just 1 of the 20 terms, ‘parent’, reflected possible existing social support.

## 8.6 Discussion

We developed SLOGAN as a framework to audit an ML4H task by identifying areas of patient severity-aware local biases. SLOGAN offers practical applications for healthcare deployment pipelines, enabling systematic bias detection prior to model implementation and continuous monitoring of bias dynamics across different hospital networks and evolving patient populations. Our results demonstrate that SLOGAN captures more and higher quality clusters across several subgroups than the baseline models, K-Means and LOGAN. To illustrate how to use SLOGAN in a clinical context, we conducted a case study that used SLOGAN to identify clusters of local bias in diabetic patients. We found that the biases observed aligned with existing health literature.

|                     |   |
|---------------------|---|
| Most biased (40.0%) | parent, given, recent, vanco, treat, fever, acinetobacter, ecg, negative, intubated, disorder, bottles, clozaril, complete, sputum, past, started, ed, found, admitted  |
| Least biased (0.2%) | noted, past, recent, home, given, due, pain, two, offspring, mild, chest, initially, without, blood, vancomycin, children, shortness_breath, sibling, admitted, started |

Table 8.5: Top 20 topic words in the most and least biased clusters using SLOGAN for **HAS DIABETES** attribute. Number is the bias score (%) of that cluster.



Namely, the cluster with the *largest local bias* was also the cluster with the *largest disease burden*. Our framework enables continuous monitoring of model biases across different healthcare contexts and patient populations, ML researchers and healthcare practitioners alike to make evidence-based decisions for developing more equitable AI-driven clinical systems.

# CHAPTER 9

## Conclusion

AI-driven language models offer unprecedented capabilities while introducing profound social challenges that demand our attention. This dissertation has grappled with both technical and social dimensions of developing inclusive large language models, examining how systemic biases become encoded in these systems through their training data, model architecture, and deployment. By centering the transgender and non-binary community, we discover patterns of algorithmic bias that go on to inform how LLMs perpetuate societal prejudices against marginalized groups more broadly. From experiments on gender-non affirmative language in Chapter 3 to BPE tokenization in Chapter 5, we consistently identify forms in which LLMs struggle with fairly and accurately representing gender-diverse individuals. The TANGO dataset introduced in Chapter 3 serves as a comprehensive benchmark for evaluating these biases, revealing high rates of misgendering and poor handling of gender-diverse pronouns across multiple model architectures.

Systematically understanding the limitations of LLMs in handling gender-diverse language offers valuable insights for possible mitigation strategies. We propose several novel technical interventions guided by these insights. In Chapter 5, we introduce Pronoun Tokenization Parity and cross-lingual transfer techniques to improve gender-neutral pronoun proficiency in LLMs while maintaining performance on canonical knowledge retrieval tasks. These methods, informed by both technical and social considerations—such as the presentation of gendered language and how language systems interact with it—collectively demonstrate the potential for more inclusive language technologies. However, as Chapter 6 argues, purely technical solutions are insufficient for addressing harmful social biases in LLMs. Our critique of current bias evaluation frameworks

reveals limitations inherent in binary gender conceptualizations, and how they can cascade biases across the AI development pipeline. As such, we also advocate for an active scrutinizing of how social norms are embedded within otherwise technical aspects of key AI pipeline development points including but not limited to data pre-processing, model design, and deployment. The community-centered practices discussed in Chapter 7 and Chapter 8, allow for more socially grounded evaluations of LLM harms, guiding the way towards more context-aware mitigation techniques.

Collectively, the findings of this dissertation reveal that true inclusivity cannot be achieved through technical mitigation alone, but require a sociotechnical vigilance from researchers who understand their privileged position in determining how AI systems impact marginalized communities. While our findings offer promising directions, they also illuminate unavoidable challenges that AI research will need to address such as: How can we facilitate meaningful and sustained community participation while protecting against extractive research practices? How do we build research relationships that honor community expertise and labor? How do we ensure research outcomes actively advance community-defined goals? These challenges present opportunities to fundamentally reshape both our technical approaches and research practices, creating new pathways for sustained community collaboration and mechanisms for translating community knowledge into technical practice.

## References

- [Ada94] Lu Ann Aday. “Health status of vulnerable populations.” Annual review of public health, **15**(1):487–509, 1994.
- [Adm22] Social Security Administration. “Popular Baby Names — ssa.gov.” <https://www.ssa.gov/oact/babynames/index.html>, 2022. [Accessed 05-Feb-2023].
- [AH13] Y Gavriel Ansara and Peter Hegarty. “Misgendering in English language contexts: Applying non-cisgenderist methods to feminist research.” International Journal of Multiple Research Approaches, **7**(2):160–177, 2013.
- [Allnd] AllenNLP. “AllenNLP Demo — demo.allennlp.org.” <https://demo.allennlp.org/coreference-resolution/>, (n.d.). [Accessed 26-Jan-2023].
- [AMB19] Emily Alsentzer, John Murphy, William Boag, Wei-Hung Weng, Di Jindi, Tristan Naumann, and Matthew McDermott. “Publicly Available Clinical BERT Embeddings.” In Proceedings of the 2nd Clinical Natural Language Processing Workshop, pp. 72–78, Minneapolis, Minnesota, USA, June 2019. Association for Computational Linguistics.
- [AMM19] Mary D Adu, Usman H Malabu, Aduli EO Malau-Aduli, and Bunmi S Malau-Aduli. “Enablers and barriers to effective diabetes self-management: A multi-national investigation.” PloS one, **14**(6):e0217771, 2019.
- [AMN22] Ali Araabi, Christof Monz, and Vlad Niculae. “How Effective is Byte Pair Encoding for Out-Of-Vocabulary Words in Neural Machine Translation?” In Conference of the Association for Machine Translation in the Americas, 2022.
- [aR20] avram anderson and Andy Lee Roth. “Queer erasure: Internet browsing can be bi-

- ased against LGBTQ people, new exclusive research shows.” Index on Censorship, **49**(1):75–77, 2020.
- [ARY19] Mikel Artetxe, Sebastian Ruder, and Dani Yogatama. “On the Cross-lingual Transferability of Monolingual Representations.” In Annual Meeting of the Association for Computational Linguistics, 2019.
- [ASO23] Arkadeep Acharya, Brijraj Singh, and Naoyuki Onoe. “Llm based generation of item-description for recommendation system.” In Proceedings of the 17th ACM Conference on Recommender Systems, pp. 1204–1207, 2023.
- [BB19] Shikha Bordia and Samuel R Bowman. “Identifying and reducing gender bias in word-level language models.” arXiv preprint arXiv:1904.03035, 2019.
- [BBD20] Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. “Language (technology) is power: A critical survey of” bias” in nlp.” arXiv preprint arXiv:2005.14050, 2020.
- [BBM08] Laura K Brennan Ramirez, Elizabeth Anne Baker, and Marilyn Metzler. “Promoting health equity; a resource to help communities address social determinants of health.” 2008.
- [BCS17] Solon Barocas, Kate Crawford, Aaron Shapiro, and Hanna Wallach. “The Problem With Bias: Allocative Versus Representational Harms in Machine Learning.” In SIGCIS Conference, 2017.
- [BCZ16] Tolga Bolukbasi, Kai-Wei Chang, James Y. Zou, Venkatesh Saligrama, and Adam T. Kalai. “Man is to Computer Programmer as Woman is to Homemaker? Debiasing Word Embeddings.” In Advances in Neural Information Processing Systems, volume 29, pp. 4349–4357. Curran Associates, Inc., 2016.

- [BD20] Kaj Bostrom and Greg Durrett. “Byte Pair Encoding is Suboptimal for Language Model Pretraining.” In Findings of the Association for Computational Linguistics: EMNLP 2020, pp. 4617–4624, 2020.
- [Bey21] Marquis Bey. “Trouble genders: “LGBT” collapse and trans fundamentality.” Hypatia, **36**(1):191–206, 2021.
- [BGM21] Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. “On the dangers of stochastic parrots: Can language models be too big?” Proceedings of FAccT, 2021.
- [BGW21] Sid Black, Leo Gao, Phil Wang, Connor Leahy, and Stella Biderman. “GPT-Neo: Large Scale Autoregressive Language Modeling with Mesh-Tensorflow.”, March 2021. If you use this software, please cite it using these metadata.
- [BHN19] Solon Barocas, Moritz Hardt, and Arvind Narayanan. Fairness and Machine Learning. fairmlbook.org, 2019. <http://www.fairmlbook.org>.
- [BHN22] Solon BAROCAS, Moritz HARDT, and Arvind NARAYANAN. “Fairness and machine learning: limitations and opportunities.[S. l.]: fairmlbook. org, 2019.”, 2022.
- [Bjo17] Bronwyn M Bjorkman. “Singular they and the syntactic representation of gender in English.” Glossa: a journal of general linguistics, **2**(1), 2017.
- [BLO21] Su Lin Blodgett, Gilsinia Lopez, Alexandra Olteanu, Robert Sim, and Hanna Wallach. “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), pp. 1004–1015, 2021.
- [BLV21] Soumya Barikeri, Anne Lauscher, Ivan Vulić, and Goran Glavaš. “RedditBias: A real-

- world resource for bias evaluation and debiasing of conversational language models.”  
[arXiv preprint arXiv:2106.03521](https://arxiv.org/abs/2106.03521), 2021.
- [BMR20] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. “Language models are few-shot learners.” Advances in neural information processing systems, **33**:1877–1901, 2020.
- [BNJ03] David M Blei, Andrew Y Ng, and Michael I Jordan. “Latent dirichlet allocation.” Journal of machine Learning research, **3**(Jan):993–1022, 2003.
- [BS20] Sabrina Burtscher and Katta Spiel. ““ But where would I even start?” developing (gender) sensitivity in HCI research and practice.” In Proceedings of the Conference on Mensch und Computer, pp. 431–441, 2020.
- [BSA23] Stella Rose Biderman, Hailey Schoelkopf, Quentin G. Anthony, Herbie Bradley, Kyle O’Brien, Eric Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, USVSN Sai Prashanth, Edward Raff, Aviya Skowron, Lintang Sutawika, and Oskar van der Wal. “Pythia: A Suite for Analyzing Large Language Models Across Training and Scaling.” ArXiv, **abs/2304.01373**, 2023.
- [BT52] Ralph Allan Bradley and Milton E Terry. “Rank analysis of incomplete block designs: I. The method of paired comparisons.” Biometrika, **39**(3/4):324–345, 1952.
- [BZK20] Jason Baumgartner, Savvas Zannettou, Brian Keegan, Megan Squire, and Jeremy Blackburn. “The pushshift reddit dataset.” In Proceedings of the international AAAI conference on web and social media, volume 14, pp. 830–839, 2020.
- [Camnd] Cambridge. “Determiners used as pronouns.” <https://dictionary.cambridge.org/us/grammar/british-grammar/determiners-used-as-pronouns>, (n.d.).

- [CBC12] Antonio Ceriello, László Barkai, Jens Sandahl Christiansen, Leszek Czupryniak, Ramon Gomis, Kari Harno, Bernhard Kulzer, Johnny Ludvigsson, Zuzana Némethyová, David Owens, et al. “Diabetes as a case study of chronic disease management with a personalized approach: the role of a structured feedback loop.” Diabetes research and clinical practice, **98**(1):5–10, 2012.
- [CBN17a] Aylin Caliskan, Joanna J Bryson, and Arvind Narayanan. “Semantics derived automatically from language corpora contain human-like biases.” Science, **356**(6334):183–186, 2017.
- [CBN17b] Aylin Caliskan, Joanna J. Bryson, and Arvind Narayanan. “Semantics derived automatically from language corpora contain human-like biases.” Science, **356**(6334):183–186, 2017.
- [CD19] Yang Trista Cao and Hal Daumé III. “Toward gender-inclusive coreference resolution.” arXiv preprint arXiv:1910.13913, 2019.
- [CD20] Yang Trista Cao and Hal Daumé III. “Toward Gender-Inclusive Coreference Resolution.” In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pp. 4568–4595, Online, 2020. Association for Computational Linguistics.
- [CD21] Yang Trista Cao and Hal Daumé III. “Toward Gender-Inclusive Coreference Resolution: An Analysis of Gender and Bias Throughout the Machine Learning Lifecycle.” Computational Linguistics, **47**(3):615–661, 2021.
- [Cennd] Gender Census. “Gender Census 2022: Worldwide Report.” <https://www.gendercensus.com/results/2022-worldwide/#pronouns>, (n.d.). [Accessed 25-Jan-2023].
- [CHK19] L Elisa Celis, Lingxiao Huang, Vijay Keswani, and Nisheeth K Vishnoi. “Classification with fairness constraints: A meta-algorithm with provable guarantees.” In



- Proceedings of the conference on fairness, accountability, and transparency, pp. 319–328, 2019.
- [Cla19] Jessica Clarke. “They, Them, and Theirs.” 132 Harvard Law Review, p. 894, 2019.
- [CLL23] Kunming Cheng, Zhiyong Li, Cheng Li, Ruijie Xie, Qiang Guo, Yongbin He, and Haiyang Wu. “The Potential of GPT-4 as an AI-Powered Virtual Assistant for Surgeons Specialized in Joint Arthroplasty.” Annals of Biomedical Engineering, **51**:1366–1370, 2023.
- [Col19] Patricia Hill Collins. Intersectionality as critical social theory. Duke University Press, 2019.
- [Colnd] St. Louis Community College. “Pronoun and antecedent agreement.”, (n.d.).
- [Con19] Kirby Conrod. Pronouns raising and emerging. PhD thesis, 2019.
- [CPR21] Irene Y Chen, Emma Pierson, Sherri Rose, Shalmali Joshi, Kadija Ferryman, and Marzyeh Ghassemi. “Ethical machine learning in healthcare.” Annual Review of Biomedical Data Science, **4**:123–144, 2021.
- [Cra17] Kate Crawford. “The trouble with bias.” Keynote at NeurIPS, 2017.
- [Cre91] Kimberle Crenshaw. “Mapping the Margins: Intersectionality, Identity Politics, and Violence against Women of Color.” Stanford Law Review, **43**(6):1241, Jul 1991.
- [CU14] Myra L Clark and Sharon W Utz. “Social determinants of type 2 diabetes and health in the United States.” World journal of diabetes, **5**(3):296, 2014.
- [DBS20] Stefanie Duguay, Jean Burgess, and Nicolas Suzor. “Queer women’s experiences of patchwork platform governance on Tinder, Instagram, and Vine.” Convergence, **26**(2):237–252, 2020.

- [DCL18] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. “Bert: Pre-training of deep bidirectional transformers for language understanding.” arXiv preprint arXiv:1810.04805, 2018.
- [DCL19a] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding.” In Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers), pp. 4171–4186, 2019.
- [DCL19b] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding.” In Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers), pp. 4171–4186, Minneapolis, Minnesota, 2019. Association for Computational Linguistics.
- [Des18] Rishikesh Bahadur Desai. “Karnataka’s Jogappas can now live a gender-fluid life.” The Hindu, 2018.
- [Dev20] Sunipa Dev. “The Geometry of Distributed Representations for Better Alignment, Attenuated Bias, and Improved Interpretability.”, 2020.
- [DFW19] Emily Dinan, Angela Fan, Adina Williams, Jack Urbanek, Douwe Kiela, and Jason Weston. “Queens are powerful too: Mitigating gender bias in dialogue generation.” arXiv preprint arXiv:1911.03842, 2019.
- [DGH18] Miguel Domingo, Mercedes García-Martínez, Alexandre Helle, Francisco Casacuberta, and Manuel Herranz. “How Much Does Tokenization Affect Neural Machine Translation?” In Conference on Intelligent Text Processing and Computational Linguistics, 2018.

- [DHP12] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. “Fairness Through Awareness.” In Proceedings of the 3rd Innovations in Theoretical Computer Science Conference, 2012.
- [Dicnd] Oxford English Dictionary. “A brief history of singular ‘they’ — Oxford English Dictionary — public.oed.com.” <https://public.oed.com/blog/a-brief-history-of-singular-they/>, (n.d.). [Accessed 25-Jan-2023].
- [DLD22] Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, and Zhifang Sui. “A Survey for In-context Learning.” arXiv preprint arXiv:2301.00234, 2022.
- [DLP20] Sunipa Dev, Tao Li, Jeff M Phillips, and Vivek Srikumar. “OSCaR: Orthogonal subspace correction and rectification of biases in word embeddings.” arXiv preprint arXiv:2007.00049, 2020.
- [DML19] Sumanth Dathathri, Andrea Madotto, Janice Lan, Jane Hung, Eric Frank, Piero Molino, Jason Yosinski, and Rosanne Liu. “Plug and play language models: A simple approach to controlled text generation.” arXiv preprint arXiv:1912.02164, 2019.
- [DMO21] Sunipa Dev, Masoud Monajatipoor, Anaelia Ovalle, Arjun Subramonian, Jeff M Phillips, and Kai-Wei Chang. “Harms of gender exclusivity and challenges in non-binary representation in language technologies.” arXiv preprint arXiv:2108.12084, 2021.
- [DP19] Sunipa Dev and Jeff M. Phillips. “Attenuating Bias in Word vectors.” In Kamalika Chaudhuri and Masashi Sugiyama, editors, The 22nd International Conference on Artificial Intelligence and Statistics, AISTATS 2019, 16-18 April 2019, Naha, Okinawa, Japan, volume 89 of Proceedings of Machine Learning Research, pp. 879–887. PMLR, 2019.

- [DRW19a] Maria De-Arteaga, Alexey Romanov, Hanna Wallach, Jennifer Chayes, Christian Borgs, Alexandra Chouldechova, Sahin Geyik, Krishnaram Kenthapadi, and Adam Tauman Kalai. “Bias in Bios: A Case Study of Semantic Representation Bias in a High-Stakes Setting.” In Proceedings of the Conference on Fairness, Accountability, and Transparency, 2019.
- [DRW19b] Maria De-Arteaga, Alexey Romanov, Hanna Wallach, Jennifer Chayes, Christian Borgs, Alexandra Chouldechova, Sahin Geyik, Krishnaram Kenthapadi, and Adam Tauman Kalai. “Bias in bios: A case study of semantic representation bias in a high-stakes setting.” In Proceedings of the Conference on Fairness, Accountability, and Transparency, pp. 120–128, 2019.
- [DS20] František Dařena and Martin Süß. “Quality of Word Vectors and its Impact on Named Entity Recognition in Czech.” European Journal of Business Science and Technology, 2020.
- [DSC22] Jamell Dacon, Harry Shomer, Shaylynn Crum-Dacon, and Jiliang Tang. “Detecting Harmful Online Conversational Content towards LGBTQIA+ Individuals.” arXiv preprint arXiv:2207.10032, 2022.
- [DSD22] Damai Dai, Yutao Sun, Li Dong, Yaru Hao, Zhifang Sui, and Furu Wei. “Why Can GPT Learn In-Context? Language Models Secretly Perform Gradient Descent as Meta Optimizers.” arXiv preprint arXiv:2212.10559, 2022.
- [DSK21] Jwala Dhamala, Tony Sun, Varun Kumar, Satyapriya Krishna, Yada Pruksachatkun, Kai-Wei Chang, and Rahul Gupta. “Bold: Dataset and metrics for measuring biases in open-ended language generation.” In Proceedings of the 2021 ACM conference on fairness, accountability, and transparency, pp. 862–872, 2021.
- [DSZ22] Sunipa Dev, Emily Sheng, Jieyu Zhao, Aubrie Amstutz, Jiao Sun, Yu Hou, Mattie Sanseverino, Jiin Kim, Akihiro Nishi, Nanyun Peng, et al. “On Measures of Biases

- and Harms in NLP.” In Findings of the Association for Computational Linguistics: ACL-IJCNLP 2022, pp. 246–267, 2022.
- [DZL17] F De Santi, G Zoppini, F Locatelli, E Finocchio, V Cappa, M Dauriz, and G Verlatto. “Type 2 diabetes is associated with an increased prevalence of respiratory symptoms as compared to the general population.” BMC Pulmonary Medicine, **17**(1):1–8, 2017.
- [ECPnd] Hannah Eyre, Alec B Chapman, Kelly S Peterson, Jianlin Shi, Patrick R Alba, Makoto M Jones, Tamara L Box, Scott L DuVall, and Olga V Patterson. “Launching into clinical space with medspaCy: a new clinical text processing toolkit in Python.” In AMIA Annual Symposium Proceedings 2021, (in press, n.d.).
- [EG18] Yanai Elazar and Yoav Goldberg. “Adversarial removal of demographic attributes from text data.” arXiv preprint arXiv:1808.06640, 2018.
- [ES11] Penny Eckert and Ivan A. Sag. “MORPHOLOGY.”, 2011. [Online PDF].
- [FBB01] Flavio Lopes Ferreira, Daliana Peres Bota, Annette Bross, Christian Mélot, and Jean-Louis Vincent. “Serial evaluation of the SOFA score to predict outcome in critically ill patients.” Jama, **286**(14):1754–1758, 2001.
- [FCJ22] Virginia K Felkner, Ho-Chun Herbert Chang, Eugene Jang, and Jonathan May. “Towards WinoQueer: Developing a Benchmark for Anti-Queer Bias in Large Language Models.” arXiv preprint arXiv:2206.11484, 2022.
- [FCJ23] Virginia Felkner, Ho-Chun Herbert Chang, Eugene Jang, and Jonathan May. “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), pp. 9126–9140, 2023.
- [Fer16] Christine Feraday. “For Lack Of A Better Word: Neo-Identities In Non-Cisgender, Non-Straight Communities On Tumblr.” Ryerson University, 2016.

- [Fis93] Susan T Fiske. “Controlling other people: The impact of power on stereotyping.” American psychologist, **48**(6):621, 1993.
- [Flo05] Glenn Flores. “The impact of medical interpreter services on the quality of health care: a systematic review.” Medical care research and review, **62**(3):255–299, 2005.
- [Flo15] A Flowers. “The most common unisex names in America: Is yours one of them? FiveThirtyEight.”, 2015.
- [For05] M.D. Fortescue. Historical Linguistics 2003: Selected Papers from the 16th International Conference on Historical Linguistics, Copenhagen, 11-15 August 2003. Amsterdam Studies in the Theory and History of Linguistic Science: 4. J. Benjamins Pub., 2005.
- [FQH24] Duanyu Feng, Bowen Qin, Chen Huang, Zheng Zhang, and Wenqiang Lei. “Towards analyzing and understanding the limitations of dpo: A theoretical perspective.” arXiv preprint arXiv:2404.04626, 2024.
- [FVB16] Ethan Fast, Tina Vachovsky, and Michael Bernstein. “Shirtless and dangerous: Quantifying linguistic signals of gender bias in an online fiction writing community.” In Proceedings of the International AAAI Conference on Web and Social Media, volume 10, 2016.
- [Gar16] B.A. Garner. The Chicago Guide to Grammar, Usage, and Punctuation. Chicago Guides to Writing, Editing, and Publishing. University of Chicago Press, 2016.
- [Gau21] Vasundhara Gautam. “Guest Lecture in Pronouns: Vasundhara.” In Kirby Conrod, editor, Pronoun Studies. 07 2021.
- [GBB20] Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, et al. “The pile: An 800gb

- dataset of diverse text for language modeling.” [arXiv preprint arXiv:2101.00027](#), 2020.
- [GBC22] Yacine Gaci, Boualem Benatallah, Fabio Casati, and Khalid Benabdeslem. “Iterative adversarial removal of gender bias in pretrained word embeddings.” In Proceedings of the 37th ACM/SIGAPP Symposium On Applied Computing, pp. 829–836, 2022.
- [Gen23] Gender Census. “2023 Gender Census.”, 2023. Accessed: September 14, 2023.
- [GGS20] Samuel Gehman, Suchin Gururangan, Maarten Sap, Yejin Choi, and Noah A Smith. “Realtotoxicityprompts: Evaluating neural toxic degeneration in language models.” [arXiv preprint arXiv:2009.11462](#), 2020.
- [GSB21] Marco Gaido, Beatrice Savoldi, Luisa Bentivogli, Matteo Negri, and Marco Turchi. “How to Split: the Effect of Word Segmentation on Gender Bias in Speech Translation.” In Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021, pp. 3576–3589, 2021.
- [HAR19] Kexin Huang, Jaan Altosaar, and Rajesh Ranganath. “Clinicalbert: Modeling clinical notes and predicting hospital readmission.” [arXiv preprint arXiv:1904.05342](#), 2019.
- [HDS20] Alex Hanna, Emily Denton, Andrew Smart, and Jamila Smith-Loud. “Towards a critical race methodology in algorithmic fairness.” In Proceedings of the 2020 conference on fairness, accountability, and transparency, pp. 501–512, 2020.
- [HDS23] Tamanna Hossain, Sunipa Dev, and Sameer Singh. “MISGENDERED: Limits of Large Language Models in Understanding Pronouns.” In The 61st Annual Meeting Of The Association For Computational Linguistics, 2023.
- [Hewnd] Robin Hewings. “Marginalization and Loneliness Among Sexual Minorities: How Are They Linked? - Campaign to End Loneliness — [campaigntoendloneliness.org](https://www.campaigntoendloneliness.org/).” <https://www.campaigntoendloneliness.org/>

marginalization-and-loneliness-among-sexual-minorities-how-are-they  
(n.d.). [Accessed 25-Jan-2023].

[HHF19] Matthias Huck, Viktor Hangya, and Alexander M. Fraser. “Better OOV Translation with Bilingual Terminology Mining.” In Annual Meeting of the Association for Computational Linguistics, 2019.

[HKK19] Hrayr Harutyunyan, Hrant Khachatrian, David C Kale, Greg Ver Steeg, and Aram Galstyan. “Multitask learning and benchmarking with clinical time series data.” Scientific data, **6**(1):1–18, 2019.

[HLE15] M Sazzad Hussain, Juchen Li, Louise A Ellis, Laura Ospina-Pinillos, Tracey A Davenport, Rafael A Calvo, and Ian B Hickie. “Moderator assistant: A natural language generation-based intervention to support mental health via social media.” Journal of Technology in Human Services, **33**(4):304–329, 2015.

[HLH22] Anne AH de Hond, Artuur M Leeuwenberg, Lotty Hooft, Ilse MJ Kant, Steven WJ Nijman, Hendrikus JA van Os, Jiska J Aardoom, Thomas Debray, Ewoud Schuit, Maarten van Smeden, et al. “Guidelines and quality criteria for artificial intelligence-based prediction models in healthcare: a scoping review.” npj Digital Medicine, **5**(1):1–13, 2022.

[HPD20a] Ben Hutchinson, Vinodkumar Prabhakaran, Emily Denton, Kellie Webster, Yu Zhong, and Stephen Denuyl. “Social Biases in NLP Models as Barriers for Persons with Disabilities.” In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pp. 5491–5501, 2020.

[HPD20b] Ben Hutchinson, Vinodkumar Prabhakaran, Emily Denton, Kellie Webster, Yu Zhong, and Stephen Denuyl. “Unintended machine learning biases as social barriers for persons with disabilities.” ACM SIGACCESS Accessibility and Computing, (125):1–1, 2020.



- [HPS16] Moritz Hardt, Eric Price, and Nati Srebro. “Equality of Opportunity in Supervised Learning.” In Advances in Neural Information Processing Systems, volume 29, pp. 3315–3323. Curran Associates, Inc., 2016.
- [HSN18] Tatsunori Hashimoto, Megha Srivastava, Hongseok Namkoong, and Percy Liang. “Fairness Without Demographics in Repeated Loss Minimization.” In Jennifer Dy and Andreas Krause, editors, Proceedings of the 35th International Conference on Machine Learning, volume 80 of Proceedings of Machine Learning Research, pp. 1929–1938. PMLR, 10–15 Jul 2018.
- [Hugnd] HuggingFace. “Neural Coreference.” <https://huggingface.co/coref/>, (n.d.). [Accessed 26-Jan-2023].
- [JCW99] Nirmal Joshi, Gregory M Caputo, Michael R Weitekamp, and AW Karchmer. “Infections in patients with diabetes mellitus.” New England Journal of Medicine, **341**(25):1906–1912, 1999.
- [JKC13] Alistair EW Johnson, Andrew A Kramer, and Gari D Clifford. “A new severity of illness scale using a subset of acute physiology and chronic health evaluation data elements shows comparable predictive accuracy.” Critical care medicine, **41**(7):1711–1718, 2013.
- [JKC15] Ankur Joshi, Saket Kale, Satish Chandel, and D Kumar Pal. “Likert scale: Explored and explained.” British journal of applied science & technology, **7**(4):396, 2015.
- [Jos22] Laura Joszt. “5 Vulnerable Populations in Healthcare.”, 2022.
- [JPS16] Alistair EW Johnson, Tom J Pollard, Lu Shen, Li-wei H Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G Mark. “MIMIC-III, a freely accessible critical care database.” Scientific data, **3**(1):1–9, 2016.

- [JSM23] Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. “Mistral 7B.” arXiv preprint arXiv:2310.06825, 2023.
- [JTK09] Alan E Jones, Stephen Trzeciak, and Jeffrey A Kline. “The Sequential Organ Failure Assessment score for predicting outcome in patients with severe sepsis and evidence of hypoperfusion at the time of emergency department presentation.” Critical care medicine, **37**(5):1649, 2009.
- [KCE18] Alan Katz, Dan Chateau, Jennifer E Enns, Jeff Valdivia, Carole Taylor, Randy Walld, and Scott McCulloch. “Association of the social determinants of health with quality of primary care.” The Annals of Family Medicine, **16**(3):217–224, 2018.
- [Key18] Os Keyes. “The misgendering machines: Trans/HCI implications of automatic gender recognition.” Proceedings of the ACM on human-computer interaction, **2**(CSCW):1–22, 2018.
- [KHB21] Os Keyes, Zoë Hitzig, and Mwenza Blell. “Truth from the machine: artificial intelligence and the materialization of identity.” Interdisciplinary Science Reviews, **46**(1-2):158–175, 2021.
- [KMO16] Jason N Katz, Michael Minder, Benjamin Olenchock, Susanna Price, Michael Goldfarb, Jeffrey B Washam, Christopher F Barnett, L Kristin Newby, and Sean van Diepen. “The genesis, maturation, and future of critical care cardiology.” Journal of the American College of Cardiology, **68**(1):67–79, 2016.
- [KRD20] Avishek Kumar, Arthi Ramachandran, Adolfo De Unanue, Christina Sung, Joe Walsh, John Schneider, Jessica Ridgway, Stephanie Masiello Schuette, Jeff Lauritsen, and Rayid Ghani. “A Machine Learning System for Retaining Patients in HIV Care.” arXiv preprint arXiv:2006.04944, 2020.

- [KSC24] Mahi Kolla, Siddharth Salunkhe, Eshwar Chandrasekharan, and Koustuv Saha. “LLM-Mod: Can Large Language Models Assist Content Moderation?” In Extended Abstracts of the CHI Conference on Human Factors in Computing Systems, pp. 1–8, 2024.
- [KYW24] Majeed Kazemitabaar, Runlong Ye, Xiaoning Wang, Austin Z. Henley, Paul Denny, Michelle Craig, and Tovi Grossman. “CodeAid: Evaluating a Classroom Deployment of an LLM-based Programming Assistant that Balances Student and Educator Needs.” Proceedings of the CHI Conference on Human Factors in Computing Systems, 2024.
- [Lak] Robin Lakoff. “Language and Woman’s Place.” Language in society., **2**(1).
- [Lau20] Dave Lauer. “You cannot have AI ethics without ethics.” In AI and Ethics, 2020.
- [LB21] Li Lucy and David Bamman. “Gender and representation bias in GPT-3 generated stories.” In Proceedings of the Third Workshop on Narrative Understanding, pp. 48–55, 2021.
- [LBM23] Tomasz Limisiewicz, Jiří Balhar, and David Mareček. “Tokenization Impacts Multilingual Language Modeling: Assessing Vocabulary Allocation and Overlap Across Languages.” In Findings of the Association for Computational Linguistics: ACL 2023, pp. 5661–5681, 2023.
- [LCH22] Anne Lauscher, Archie Crowley, and Dirk Hovy. “Welcome to the Modern World of Pronouns: Identity-Inclusive Natural Language Processing beyond Gender.” In Proceedings of the 29th International Conference on Computational Linguistics, pp. 1221–1232, 2022.
- [LDT18] Liliana Laranjo, Adam G Dunn, Huong Ly Tong, Ahmet Baki Kocaballi, Jessica Chen, Rabia Bashir, Didi Surian, Blanca Gallego, Farah Magrabi, Annie YS Lau, et al. “Conversational agents in healthcare: a systematic review.” Journal of the American Medical Informatics Association, **25**(9):1248–1258, 2018.

- [LGP20] Anne Lauscher, Goran Glavaš, Simone Paolo Ponzetto, and Ivan Vulić. “A general framework for implicit and explicit debiasing of distributional word vector spaces.” In Proceedings of the AAAI Conference on Artificial Intelligence, volume 34, pp. 8131–8138, 2020.
- [LHE22] Stephanie Lin, Jacob Hilton, and Owain Evans. “TruthfulQA: Measuring How Models Mimic Human Falsehoods.” In Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pp. 3214–3252, 2022.
- [LLS93] Jean-Roger Le Gall, Stanley Lemeshow, and Fabienne Saulnier. “A new simplified acute physiology score (SAPS II) based on a European/North American multicenter study.” Jama, **270**(24):2957–2963, 1993.
- [LOG19] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. “Roberta: A robustly optimized bert pretraining approach.” arXiv preprint arXiv:1907.11692, 2019.
- [LS22] Sue Lim and Ralf Schmäzlle. “Artificial intelligence for health message generation: an empirical study using a large language model (LLM) and prompt engineering.” In Frontiers in Communication, 2022.
- [LSZ21] Jiachang Liu, Dinghan Shen, Yizhe Zhang, Bill Dolan, Lawrence Carin, and Weizhu Chen. “What Makes Good In-Context Examples for GPT-3?” arXiv preprint arXiv:2101.06804, 2021.
- [LWW20] Haochen Liu, Wentao Wang, Yiqi Wang, Hui Liu, Zitao Liu, and Jiliang Tang. “Mitigating gender bias for neural dialogue generation with adversarial learning.” arXiv preprint arXiv:2009.13028, 2020.
- [MAP16] MAP. “Unjust: How the Broken Criminal Justice System Fails Transgender People.” Movement Advancement Project and Center for American Progress, 2016.

- [Mar05] Michael Marmot. “Social determinants of health inequalities.” The lancet, **365**(9464):1099–1104, 2005.
- [MAS21] Sabrina J Mielke, Zaid Alyafeai, Elizabeth Salesky, Colin Raffel, Manan Dey, Matthias Gallé, Arun Raja, Chenglei Si, Wilson Y Lee, Benoît Sagot, et al. “Between words and characters: A brief history of open-vocabulary modeling and tokenization in NLP.” arXiv preprint arXiv:2112.10508, 2021.
- [McL18] Kevin A McLemore. “A minority stress perspective on transgender individuals’ experiences with misgendering.” Stigma and Health, **3**(1):53, 2018.
- [MGC19] Rowan Hall Maudslay, Hila Gonen, Ryan Cotterell, and Simone Teufel. “It’s all in the name: Mitigating gender bias with name-based counterfactual data substitution.” arXiv preprint arXiv:1909.00871, 2019.
- [MGH05] LMAJ Muller, KJ Gorter, E Hak, WL Goudzwaard, FG Schellevis, AIM Hoepelman, and GEHM Rutten. “Increased risk of common infections in patients with type 1 and type 2 diabetes mellitus.” Clinical infectious diseases, **41**(3):281–288, 2005.
- [MGM19] Ninareh Mehrabi, Thamme Gowda, Fred Morstatter, Nanyun Peng, and A. G. Galstyan. “Man is to Person as Woman is to Location: Measuring Gender Bias in Named Entity Recognition.” Proceedings of the 31st ACM Conference on Hypertext and Social Media, 2019.
- [MGM20] Ninareh Mehrabi, Thamme Gowda, Fred Morstatter, Nanyun Peng, and Aram Galstyan. “Man is to Person as Woman is to Location: Measuring Gender Bias in Named Entity Recognition.” In Proceedings of the 31st ACM Conference on Hypertext and Social Media, HT ’20, p. 231–232, New York, NY, USA, 2020. Association for Computing Machinery.
- [Mil16] Ehm Hjorth Miltersen. “Nounself pronouns: 3rd person personal pronouns as iden-

- tity expression.” Journal of Language Works-Sprogvidenskabeligt Studentertidsskrift, **1**(1):37–62, 2016.
- [Mir16] Alfredo Mirandé. “Hombres Mujeres: An Indigenous Third Gender.” Men and Masculinities, **19**(4):384–409, 2016.
- [MMS21] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. “A survey on bias and fairness in machine learning.” ACM computing surveys (CSUR), **54**(6):1–35, 2021.
- [Mor83] Cherríe Moraga. “It’s the Poverty.” In Cherríe Moraga and Gloria Anzaldúa, editors, This Bridge Called My Back: Writings by Radical Women of Color, chapter 6, p. 234. Kitchen Table: Women of Color Press, New York, 2 edition, 1983.
- [MWB19] Chandler May, Alex Wang, Shikha Bordia, Samuel R. Bowman, and Rachel Rudinger. “On Measuring Social Biases in Sentence Encoders.” In Proceedings of NAACL-HLT, 2019.
- [MWK02] J Michael McGinnis, Pamela Williams-Russo, and James R Knickman. “The case for more active policy attention to health promotion.” Health affairs, **21**(2):78–93, 2002.
- [NBH21] Debora Nozza, Federico Bianchi, and Dirk Hovy. “HONEST: Measuring hurtful sentence completion in language models.” In The 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies. Association for Computational Linguistics, 2021.
- [NBL22] Debora Nozza, Federico Bianchi, Anne Lauscher, Dirk Hovy, et al. “Measuring harmful sentence completion in language models for LGBTQIA+ individuals.” In Proceedings of the Second Workshop on Language Technology for Equality, Diversity and Inclusion. Association for Computational Linguistics, 2022.

- [NGV12] Courtney Napoles, Matthew R Gormley, and Benjamin Van Durme. “Annotated gigaword.” In Proceedings of the joint workshop on automatic knowledge base construction and web-scale knowledge extraction (AKBC-WEKEX), pp. 95–100, 2012.
- [Nonnd] Nonbinary Wiki. “Names.”, n.d. Accessed: 2024-11-14.
- [ODZ23] Anaelia Ovalle, Sunipa Dev, Jieyu Zhao, Majid Sarrafzadeh, and Kai-Wei Chang. “Auditing algorithmic fairness in machine learning for health with severity-based LOGAN.” In International Workshop on Health Intelligence, pp. 123–136. Springer, 2023.
- [OFG20] Luis Oala, Jana Fehr, Luca Gilli, Pradeep Balachandran, Alixandro Werneck Leite, Saul Calderon-Ramirez, Danny Xie Li, Gabriel Nobis, Erick Alejandro Muñoz Alvarado, Giovanna Jaramillo-Gutierrez, et al. “Ml4h auditing: From paper to practice.” In Machine learning for health, pp. 280–317. PMLR, 2020.
- [OGD23] Anaelia Ovalle, Palash Goyal, Jwala Dhamala, Zachary Jagers, Kai-Wei Chang, Aram Galstyan, Richard Zemel, and Rahul Gupta. ““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, pp. 1246–1266, 2023.
- [OMG24] Anaelia Ovalle, Ninareh Mehrabi, Palash Goyal, Jwala Dhamala, Kai-Wei Chang, Richard Zemel, Aram Galstyan, Yuval Pinter, and Rahul Gupta. “Tokenization matters: Navigating data-scarce tokenization for gender inclusive language technologies.” In Findings of the Association for Computational Linguistics: NAACL 2024, pp. 1739–1756, 2024.
- [Ope23] OpenAI. “ChatGPT: Optimizing language models for dialogue.”, Jan 2023.

- [OPM24] Anaelia Ovalle, Krunoslav Lehman Pavasovic, Louis Martin, Luke Zettlemyer, Eric Michael Smith, Adina Williams, and Levent Sagun. “The Root Shapes the Fruit: On the Persistence of Gender-Exclusive Harms in Aligned Language Models.” arXiv preprint arXiv:2411.03700, 2024.
- [OPV19] Ziad Obermeyer, Brian Powers, Christine Vogeli, and Sendhil Mullainathan. “Dissecting racial bias in an algorithm used to manage the health of populations.” Science, **366**(6464):447–453, 2019.
- [OSG23] Anaelia Ovalle, Arjun Subramonian, Vagrant Gautam, Gilbert Gee, and Kai-Wei Chang. “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, pp. 496–511, 2023.
- [PAR21] Jae A Puckett, Alix B Aboussouan, Allura L Ralston, Brian Mustanski, and Michael E Newcomb. “Systems of cissexism and the daily production of stress for transgender and gender diverse people.” International Journal of Transgender Health, pp. 1–14, 2021.
- [PAR23] Jae A Puckett, Alix B Aboussouan, Allura L Ralston, Brian Mustanski, and Michael E Newcomb. “Systems of cissexism and the daily production of stress for transgender and gender diverse people.” International journal of transgender health, **24**(1):113–126, 2023.
- [PCH07] Monica E Peek, Algernon Cargill, and Elbert S Huang. “Diabetes health disparities.” Medical care research and review, **64**(5\_suppl):101S–156S, 2007.
- [Peand] Pearson. “Gender Policing and Gender Accountability.” [https://revelpreview.pearson.com/epubs/pearson\\_kimmel\\_soc/OPS/xhtml/ch09\\_pg0013.xhtml](https://revelpreview.pearson.com/epubs/pearson_kimmel_soc/OPS/xhtml/ch09_pg0013.xhtml), (n.d.). [Accessed 25-Jan-2023].



- [PFS20] Adam Poulsen, Eduard Fosch-Villaronga, and Roger Andre Søråa. “Queering machines.” Nature Machine Intelligence, **2**(3):152–152, 2020.
- [PFS21] Stephen R Pfohl, Agata Foryciarz, and Nigam H Shah. “An empirical characterization of fair machine learning for clinical risk prediction.” Journal of biomedical informatics, **113**:103621, 2021.
- [PSM14] Jeffrey Pennington, Richard Socher, and Christopher D Manning. “GloVe: Global vectors for word representation.” In Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP), pp. 1532–1543, 2014.
- [QDO23] Organizers of QueerInAI, Nathaniel Dennler, Anaelia Ovalle, Ashwin Singh, Luca Soldaini, Arjun Subramonian, Huy Tu, William Agnew, Avijit Ghosh, Kyra Yee, Irene Font Peradejordi, Zeerak Talat, Mayra Russo, and Jessica de Jesus de Pinho Pinal. “Bound by the Bounty: Collaboratively Shaping Evaluation Processes for Queer AI Harms.” Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society, 2023.
- [QOS23] Organizers Of QueerInAI, Anaelia Ovalle, Arjun Subramonian, Ashwin Singh, Claas Voelcker, Danica J Sutherland, Davide Locatelli, Eva Breznik, Filip Klubicka, Hang Yuan, et al. “Queer In AI: A Case Study in Community-Led Participatory AI.” In Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency, pp. 1882–1895, 2023.
- [QSA21] Organizers of QueerInAI, Ashwin S, William Agnew, Hetvi Jethwani, and Arjun Subramonian. “Rebuilding Trust: Queer in AI Approach to Artificial Intelligence Risk Management.”, 2021.
- [RBH22] Eliane Rösli, Selen Bozkurt, and Tina Hernandez-Boussard. “Peeking into a black box, the fairness and generalizability of a MIMIC-III benchmarking model.” Scientific Data, **9**(1):1–13, 2022.

- [RBP21] Inioluwa Deborah Raji, Emily M Bender, Amandalynne Paullada, Emily Denton, and Alex Hanna. “AI and the Everything in the Whole Wide World Benchmark.” 2021.
- [RD19a] Micah Rajunov and A Scott Duane. Nonbinary: Memoirs of gender and identity. Columbia University Press, 2019.
- [RD19b] Micah Rajunov and Scott Duane. Nonbinary: Memoirs of Gender and Identity. Columbia University Press, 2019.
- [REG20] Shauli Ravfogel, Yanai Elazar, Hila Gonen, Michael Twiton, and Yoav Goldberg. “Null It Out: Guarding Protected Attributes by Iterative Nullspace Projection.” In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pp. 7237–7256, Online, 2020. Association for Computational Linguistics.
- [RHH18] Alvin Rajkomar, Michaela Hardt, Michael D Howell, Greg Corrado, and Marshall H Chin. “Ensuring fairness in machine learning to advance health equity.” Annals of internal medicine, **169**(12):866–872, 2018.
- [Ricnd] Leonard Richardson. “Beautiful Soup: We called him Tortoise because he taught us. — crummy.com.” <https://www.crummy.com/software/BeautifulSoup/>, (n.d.). [Accessed 05-Feb-2023].
- [RNL18] Rachel Rudinger, Jason Naradowsky, Brian Leonard, and Benjamin Van Durme. “Gender Bias in Coreference Resolution.” In Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers), pp. 8–14, 2018.
- [RNS18] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. “Improving language understanding by generative pre-training.” 2018.
- [Ros20] The Blunt Rose. “Nonbinary Name List.”, 2020.

- [RPV20] Phillip Rust, Jonas Pfeiffer, Ivan Vulic, Sebastian Ruder, and Iryna Gurevych. “How Good is Your Tokenizer? On the Monolingual Performance of Multilingual Language Models.” ArXiv, **abs/2012.15613**, 2020.
- [RSM23] Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. “Direct Preference Optimization: Your Language Model is Secretly a Reward Model.” In A. Oh, T. Neumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, editors, Advances in Neural Information Processing Systems, volume 36, pp. 53728–53741. Curran Associates, Inc., 2023.
- [RWC19] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. “Language models are unsupervised multitask learners.” OpenAI blog, **1(8):9**, 2019.
- [SA21] Karolina Stanczak and Isabelle Augenstein. “A survey on gender bias in natural language processing.” arXiv preprint arXiv:2112.14168, 2021.
- [Sap21] Maarten Sap. “Positive AI with Social Commonsense Models.” Allen Institute for Artificial Intelligence, 2021.
- [SB18] Morgan Klaus Scheuerman and Jed R Brubaker. “Gender is not a Boolean: towards designing algorithms to understand complex human identities.” In In Participation+ Algorithms Workshop at CSCW 2018, 2018.
- [SCL18] Robyn Speer, Joshua Chin, Andrew Lin, Sara Jewett, and Lance Nathan. “LuminosoInsight/wordfreq: v2.2.”, 2018.
- [SCN19] Emily Sheng, Kai-Wei Chang, Premkumar Natarajan, and Nanyun Peng. “The woman worked as a babysitter: On biases in language generation.” arXiv preprint arXiv:1909.01326, 2019.

- [SCN20] Emily Sheng, Kai-Wei Chang, Premkumar Natarajan, and Nanyun Peng. “Towards controllable biases in language generation.” arXiv preprint arXiv:2005.00268, 2020.
- [SCN21] Emily Sheng, Kai-Wei Chang, Prem Natarajan, and Nanyun Peng. “Societal Biases in Language Generation: Progress and Challenges.” 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), 2021.
- [SCS21] Tulika Saha, Saraansh Chopra, Sriparna Saha, Pushpak Bhattacharyya, and Pankaj Kumar. “A large-scale dataset for motivational dialogue system: An application of natural language generation to mental health.” In 2021 International Joint Conference on Neural Networks (IJCNN), pp. 1–8. IEEE, 2021.
- [Sen16] Rico Sennrich. “How grammatical is character-level neural machine translation? Assessing MT quality with contrastive translation pairs.” arXiv preprint arXiv:1612.04629, 2016.
- [Ser07] Julia Serano. Whipping Girl: A Transsexual Woman on Sexism and the Scapegoating of Femininity. Seal Press, 2007.
- [SF07] Anthony J Sanford and Ruth Filik. ““They” as a gender-unspecified singular pronoun: Eye tracking reveals a processing cost.” Quarterly Journal of Experimental Psychology, **60**(2):171–178, 2007.
- [SFG21] Shubhanshu Shekhar, Greg Fields, Mohammad Ghavamzadeh, and Tara Javidi. “Adaptive sampling for minimax fair classification.” Advances in Neural Information Processing Systems, **34**:24535–24544, 2021.
- [SGQ19] Maarten Sap, Saadia Gabriel, Lianhui Qin, Dan Jurafsky, Noah A Smith, and Yejin Choi. “Social bias frames: Reasoning about social and power implications of language.” arXiv preprint arXiv:1911.03891, 2019.

- [SGT19] Tony Sun, Andrew Gaut, Shirlyn Tang, Yuxin Huang, Mai ElSherief, Jieyu Zhao, Diba Mirza, Elizabeth Belding, Kai-Wei Chang, and William Yang Wang. “Mitigating Gender Bias in Natural Language Processing: Literature Review.” Association for Computational Linguistics (ACL 2019), 2019.
- [SGT21] Tony Sun, Andrew Gaut, Shirlyn Tang, Yuxin Huang, Mai ElSherief, Jieyu Zhao, Diba Mirza, Elizabeth Belding, Kai-Wei Chang, and William Yang Wang. “They, Them, Theirs: Rewriting with Gender-Neutral English.” arXiv preprint arXiv:2102.06788, 2021.
- [Sha15] Eve Shapiro. Gender circuits: Bodies and identities in a technological age. Routledge, 2015.
- [Sil80] Jeanette Silveira. “Generic masculine words and thinking.” Women’s Studies International Quarterly, 3(2-3):165–178, 1980.
- [SKB19] Katta Spiel, Os Keyes, and Pinar Barlas. “Patching Gender: Non-Binary Utopias in HCI.” In Association for Computing Machinery, CHI EA ’19, p. 1–11, New York, NY, USA, 2019.
- [SKR19] Koustuv Saha, Sang Chan Kim, Manikanta D Reddy, Albert J Carter, Eva Sharma, Oliver L Haimson, and Munmun De Choudhury. “The language of LGBTQ+ minority stress experiences on social media.” Proceedings of the ACM on human-computer interaction, 3(CSCW):1–22, 2019.
- [SOP22] Michael Sun, Tomasz Oliwa, Monica E Peek, and Elizabeth L Tung. “Negative Patient Descriptors: Documenting Racial Bias In The Electronic Health Record: Study examines racial bias in the patient descriptors used in the electronic health record.” Health Affairs, pp. 10–1377, 2022.
- [SOW20] Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F Christiano. “Learning to summarize with

- human feedback.” Advances in Neural Information Processing Systems, **33**:3008–3021, 2020.
- [Spa15] Dean Spade. Normal Life: Administrative Violence, Critical Trans Politics, and the Limits of Law. Duke University Press, 2015.
- [SQX20] Yolande Strengers, Lizhen Qu, Qionikai Xu, and Jarrod Knibbe. “Adhering, steering, and queering: Treatment of gender in natural language generation.” In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, pp. 1–14, 2020.
- [SS20] Timo Schick and Hinrich Schütze. “It’s not just size that matters: Small language models are also few-shot learners.” arXiv preprint arXiv:2009.07118, 2020.
- [SSZ19a] Gabriel Stanovsky, Noah A. Smith, and Luke Zettlemoyer. “Evaluating Gender Bias in Machine Translation.” In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, pp. 1679–1684, Florence, Italy, July 2019. Association for Computational Linguistics.
- [SSZ19b] Gabriel Stanovsky, Noah A Smith, and Luke Zettlemoyer. “Evaluating Gender Bias in Machine Translation.” In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, pp. 1679–1684, 2019.
- [SUS21] Timo Schick, Sahana Udupa, and Hinrich Schütze. “Self-diagnosis and self-debiasing: A proposal for reducing corpus-based bias in nlp.” Transactions of the Association for Computational Linguistics, **9**:1408–1424, 2021.
- [SW22] Haytham Siala and Yichuan Wang. “SHIFTing artificial intelligence to be responsible in healthcare: A systematic review.” Social Science & Medicine, p. 114782, 2022.
- [SWD17] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. “Proximal policy optimization algorithms.” arXiv preprint arXiv:1707.06347, 2017.

- [Tem57] Mildred C Templin. “Certain language skills in children; their development and inter-relationships.” 1957.
- [TL22] Zeerak Talat and Anne Lauscher. “Back to the Future: On Potential Histories in NLP.” ArXiv, **abs/2210.06245**, 2022.
- [TM22] Alayo Tripp and Benjamin Munson. “Perceiving gender while perceiving language: Integrating psycholinguistics and gender theory.” Wiley Interdisciplinary Reviews: Cognitive Science, **13(2):e1583**, 2022.
- [TMK21] Nenad Tomasev, Kevin R McKee, Jackie Kay, and Shakir Mohamed. “Fairness for unobserved characteristics: Insights from technological impacts on queer communities.” In Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society, pp. 254–265, 2021.
- [TMS23] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. “Llama 2: Open foundation and fine-tuned chat models.” arXiv preprint arXiv:2307.09288, 2023.
- [TYB19] Nat Thorne, Andrew Kam-Tuck Yip, Walter Pierre Bouman, Ellen Marshall, and Jon Arcelus. “The terminology of identities between, outside and beyond the gender binary - A systematic review.” International Journal of Transgenderism, 2019.
- [VN21] Wietse de Vries and Malvina Nissim. “As Good as New. How to Successfully Recycle English GPT-2 to Make Models for Other Languages.” In Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021, pp. 836–846, 2021.
- [WDX22] Xiao Wang, Shihan Dou, Li Xiong, Yicheng Zou, Qi Zhang, Tao Gui, Liang Qiao, Zhanzhan Cheng, and Xuanjing Huang. “MINER: Improving Out-of-Vocabulary Named Entity Recognition from an Information Theoretic Perspective.” In Annual Meeting of the Association for Computational Linguistics, 2022.

- [Web19] Shannon Weber. Queer Media Images: LGBT Perspectives (Born This Way: Biology and Sexuality in Lady Gaga’s Pro-LGBT Media). Lexington Books, 2019.
- [WGU21] Johannes Welbl, Amelia Glaese, Jonathan Uesato, Sumanth Dathathri, John Mellor, Lisa Anne Hendricks, Kirsty Anderson, Pushmeet Kohli, Ben Coppin, and Po-Sen Huang. “Challenges in detoxifying language models.” arXiv preprint arXiv:2109.07445, 2021.
- [WKK19] Thomas Wiegand, Ramesh Krishnamurthy, Monique Kuglitsch, Naomi Lee, Sameer Pujari, Marcel Salathé, Markus Wenzel, and Shan Xu. “WHO and ITU establish benchmarking process for artificial intelligence in health.” The Lancet, **394**(10192):9–11, 2019.
- [WMR21] Laura Weidinger, John Mellor, Maribeth Rauh, Conor Griffin, Jonathan Uesato, Po-Sen Huang, Myra Cheng, Mia Glaese, Borja Balle, Atoosa Kasirzadeh, et al. “Ethical and social risks of harm from language models.” arXiv preprint arXiv:2112.04359, 2021.
- [WRA18] Kellie Webster, Marta Recasens, Vera Axelrod, and Jason Baldridge. “Mind the GAP: A Balanced Corpus of Gendered Ambiguous Pronouns.” Transactions of the Association for Computational Linguistics, **6**:605–617, 2018.
- [WRA20] Kellie Webster, Marta Recasens, Vera Axelrod, and Jason Baldridge. “Measuring and Reducing Gendered Correlations in Pre-trained Models.” In Proceedings of EMNLP, 2020.
- [WSK21] Tao Wang, Luo Si, Ryan Kennedy, and Su Lin Blodgett. “Fair Learning with Private Demographic Data.” In Proceedings of EMNLP, 2021.
- [WYS19] Hai Wang, Dian Yu, Kai Sun, Jianshu Chen, and Dong Yu. “Improving Pre-Trained Multilingual Model with Vocabulary Expansion.” In Proceedings of the 23rd



- Conference on Computational Natural Language Learning (CoNLL). Association for Computational Linguistics, 2019.
- [WZY19] Tianlu Wang, Jieyu Zhao, Mark Yatskar, Kai-Wei Chang, and Vicente Ordonez. “Balanced Datasets Are Not Enough: Estimating and Mitigating Gender Bias in Deep Image Representations.” In 2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019, pp. 5309–5318. IEEE, 2019.
- [WZY20] Tianlu Wang, Jieyu Zhao, Mark Yatskar, Kai-Wei Chang, and Vicente Ordonez. “Double-Hard Debias: Tailoring Word Embeddings for Gender Bias Mitigation.” In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pp. 5443–5453. Association for Computational Linguistics, 2020.
- [YP22] Shaked Yehezkel and Yuval Pinter. “Incorporating Context into Subword Vocabularies.” In Conference of the European Chapter of the Association for Computational Linguistics, 2022.
- [ZC20] Jieyu Zhao and Kai-Wei Chang. “LOGAN: Local Group Bias Detection by Clustering.” In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), pp. 1968–1977, 2020.
- [ZLA20] Haoran Zhang, Amy X Lu, Mohamed Abdalla, Matthew McDermott, and Marzyeh Ghassemi. “Hurtful words: quantifying biases in clinical contextual word embeddings.” In proceedings of the ACM Conference on Health, Inference, and Learning, pp. 110–120, 2020.
- [ZLM18] Brian Hu Zhang, Blake Lemoine, and Margaret Mitchell. “Mitigating Unwanted Biases with Adversarial Learning.” In Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society, pp. 335–340, 2018.

- [ZLX23] Chunting Zhou, Pengfei Liu, Puxin Xu, Srinivasan Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, LILI YU, Susan Zhang, Gargi Ghosh, Mike Lewis, Luke Zettlemoyer, and Omer Levy. “LIMA: Less Is More for Alignment.” In A. Oh, T. Neumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, editors, Advances in Neural Information Processing Systems, volume 36, pp. 55006–55021. Curran Associates, Inc., 2023.
- [ZMW19] Ran Zmigrod, Sabrina J Mielke, Hanna Wallach, and Ryan Cotterell. “Counterfactual data augmentation for mitigating gender stereotypes in languages with rich morphology.” arXiv preprint arXiv:1906.04571, 2019.
- [ZMW24] Chen Zhao, Feng Mi, Xintao Wu, Kai Jiang, Latifur Khan, and Feng Chen. “Dynamic Environment Responsive Online Meta-Learning with Fairness Awareness.” ACM Transactions on Knowledge Discovery from Data, **18(6)**:1–23, 2024.
- [ZRG22] Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. “Opt: Open pre-trained transformer language models.” arXiv preprint arXiv:2205.01068, 2022.
- [ZWY17a] Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. “Men Also Like Shopping: Reducing Gender Bias Amplification using Corpus-level Constraints.” In Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, pp. 2979–2989, 2017.
- [ZWY17b] Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. “Men Also Like Shopping: Reducing Gender Bias Amplification using Corpus-level Constraints.” Proceedings of the EMNLP 2017, Sep 2017.
- [ZWY18a] Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. “Gender Bias in Coreference Resolution: Evaluation and Debiasing Methods.” In Proceedings of the 2018 Conference of the North American Chapter of the Association

for Computational Linguistics: Human Language Technologies, volume 2 (Short Papers), pp. 15–20. Association for Computational Linguistics, 2018.

[ZWY18b] Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. “Gender Bias in Coreference Resolution: Evaluation and Debiasing Methods.” In Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers), pp. 15–20, New Orleans, Louisiana, June 2018. Association for Computational Linguistics.

[ZWY18c] Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. “Gender Bias in Coreference Resolution: Evaluation and Debiasing Methods.” In North American Chapter of the Association for Computational Linguistics, 2018.

[ZWY19] Jieyu Zhao, Tianlu Wang, Mark Yatskar, Ryan Cotterell, Vicente Ordonez, and Kai-Wei Chang. “Gender Bias in Contextualized Word Embeddings.” In Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers), pp. 629–634, Minneapolis, Minnesota, 2019. Association for Computational Linguistics.

[ZZL18] Jieyu Zhao, Yichao Zhou, Zeyu Li, Wei Wang, and Kai-Wei Chang. “Learning Gender-Neutral Word Embeddings.” In Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, pp. 4847–4853, 2018.