

On the Influence of Gender and Race in Romantic Relationship Prediction from Large Language Models

Abhilasha Sancheti* Haozhe An* Rachel Rudinger

University of Maryland, College Park
{sancheti, haozhe, rudinger}@umd.edu

Abstract

We study the presence of heteronormative biases and prejudice against interracial romantic relationships in large language models by performing controlled name-replacement experiments for the task of relationship prediction. We show that models are less likely to predict romantic relationships for (a) same-gender character pairs than different-gender pairs; and (b) intra/inter-racial character pairs involving Asian names as compared to Black, Hispanic, or White names. We examine the contextualized embeddings of first names and find that gender for Asian names is less discernible than non-Asian names. We discuss the social implications of our findings, underlining the need to prioritize the development of inclusive and equitable technology.

1 Introduction

Identifying romantic relationships from a given dialogue presents a challenging task in natural language understanding (Jia et al., 2021; Tiginova et al., 2021). The perceived gender, race, or ethnicity of the speakers, often inferred from their names, may inadvertently lead a model to predict a relationship type that conforms to conventional societal views. We hypothesize that, when predicting romantic relationships, models may mirror *heteronormative biases* (Pollitt et al., 2021; Vásquez et al., 2022) and prejudice against interracial romantic relationships (Lewandowski and Jackson, 2001; Miller et al., 2004) present in humans and society. Heteronormative biases assume and favor traditional gender roles, heterosexual relationships, and nuclear families, often marginalizing other gender expressions, sexuality, and family dynamics. In the US, legal protections for interracial and gay marriages were not achieved nationwide until 1967

*These authors contributed equally to this work.

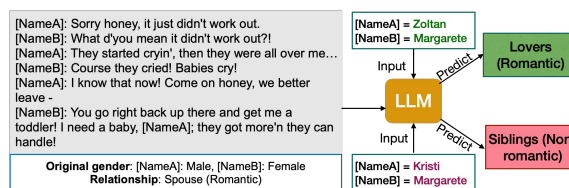


Figure 1: Sample conversation from DDRel (Jia et al., 2021) dataset and relationships predicted by Llama2-7B when characters are replaced by names with **different-gender** and **same-gender**. LLM tends to predict differently despite the same conversation.

and 2015, respectively. These relationships continue to face prejudice and discrimination in the present days (Buist, 2019; Knauer, 2020; Zambelli, 2023; Pittman et al., 2024; Daniel, 2024).

In this paper, we consider the task of predicting romantic relationships from dialogues in movie scripts to study *whether LLMs make such predictions based on the demographic attributes associated with a pair of character names, in ways that reflect heteronormative biases and prejudice against interracial romantic relationships*. For instance, Figure 1 shows a conversation between a female and a male spouse pair, for which Llama2-7B predicts a romantic relationship when the names in the conversation are replaced with a pair of different-gender names, but predicts a non-romantic relationship when replaced by same-gender names.

Ideally, name-replacement should not significantly alter the predictions of a fair and robust model, as the utterance content plays a more substantial role in language understanding, despite the potential interdependence between utterances and original names. Different predictions suggest that a model may be prone to overlooking romantic relationships that diverge from societal norms, thus raising ethical concerns. Such behavior would indicate that language models inadequately represent certain societal groups (Blodgett et al., 2020), potentially exacerbating stigma surrounding relation-

ships (Rosenthal and Starks, 2015; Reczek, 2020) and sidelining underrepresented groups (Nozza et al., 2022; Felkner et al., 2023).

Through controlled character name-replacement experiments, we find that relationships between (a) same-gender character pairs; and (b) intra/interracial character pairs involving Asian names are less likely to be predicted as romantic. These findings reveal how some LLMs may stereotypically interpret interactions between people, potentially reducing the recognition of non-mainstream relationship types. While prior work studies gender and racial biases by identifying stereotypical attributes of *individuals* (Cao et al., 2022; Cheng et al., 2023; An et al., 2023), this paper investigates the role of gender and race in LLMs’ inferences about *relationships* between *two* individuals using a relationship prediction dataset (Jia et al., 2021).

2 Experimental Setup

We define the following task. Given a conversation C which consists of a sequence of turns $((S_1, u_1), (S_2, u_2), \dots, (S_n, u_n))$ between characters A and B , where $S_i \in \{S_A, S_B\}$ indicates that the speaker of an utterance ($u_i, i \in \{1 : n\}$) is either A or B , the task is to identify the relationship represented as a categorical label from a pre-defined set. We carry out controlled name-replacement experiments by prompting LLMs (zero-shot) to predict the relationship type between A and B given C .

Models We study Llama2 ({7B, 13B}-chat) (Touvron et al., 2023) with its official implementation,¹ and Mistral-7B-Instruct (Jiang et al., 2023) using its huggingface implementation. Hyperparameters are specified in §A.

Dataset We use the test set of DDRel (Jia et al., 2021) which consists of movie scripts from IMSDb, with annotations for relationship labels between the characters according to 13 pre-defined types (Table 3 in appendix). We consider Lovers, Spouse, or Courtship predictions as romantic and the rest as non-romantic. For our experiments, we use 327 instances of the test set in which characters originally have different genders (manually annotated) because the test set has no dialogues between same-gender characters with the romantic label. We discuss the limitations of this study due to data source representation issues at the end of this paper.

¹<https://github.com/facebookresearch/llama>

Prompt Selection As LLMs are sensitive to prompts (Min et al., 2022), we experimented with several prompt formulations on the original data (test set) for accuracy, and selected the prompt (see Figure 4 in appendix) resulting in the highest accuracy which was closest to scores reported by others (Jia et al., 2021; Ou et al., 2024). We note that our prompt selection is done prior to running the name-replacement experiments.

Evaluation We compare the average recall of predicting romantic relationships across different gender assignments and races/ethnicities. We study recall as we hypothesize heteronormative and interracial relationship biases would manifest as low (romantic) recall for same-gender and interracial groups. For completeness, we also report the mean precision, F1, and accuracy scores in §D.

2.1 Studying the Influence of Gender Pairings

We ask whether the models are equally likely to recognize romantic relationships for character pairs of varying gender assignments and if this behavior is the same across different races. We hypothesize that models are prone to heteronormative bias and are more likely to predict romantic relationships for contrastive gender assignments. To test this, we collect 30 names per race,² dividing them into 10 non-linearly segmented bins that cover gender-neutral names (shown in Figure 2) based on the percentage of population that has been assigned as female at birth. Detailed name inclusion criteria and data sources are elaborated in §C.1. We replace the original name-pair in each conversation with all pairs of distinct names per race.

As dialogues may reveal gender identities (*e.g.*, “sir”, “ma’am”, “father”, etc.), we manually identify a subset (271 instances) where such explicit cues are absent (to the best of our judgement) to minimize gender information leakage and avoid explicit gender inconsistency between the dialogue and the gender associated with the replaced name. In these dialogues, gendered pronouns typically refer to a third person who is not part of the conversation. As a result, they do not reveal the speakers’ gender identity. However, pronouns can indicate the sexual orientation of a speaker (*e.g.*, “Betty: *You do love him, don’t you?*”). Such cues, along with other implicit cues about gender identity that are harder to detect, may confound our analysis. However, our

²Except for Hispanic wherein we did not get any names in 5 – 10% bin and only 1 name in 25 – 50% bin.




Figure 2: Recall of predicting romantic relationships from Llama2-7B for subset of the dataset where characters originally have different genders. Horizontal and vertical axes denote % female of the name replacing an originally female and male character name from the dialogue. The upper-triangle (lower-triangle) shows the scores when names are replaced preserving (swapping) the genders of characters’ names as-is in the original conversation. We consider the names with lesser % female as male names for determining gender preservation for name-replacement.

findings as discussed in §3 reveal that implicit cues are not a major confounding factor. We discuss this aspect further in the Limitations section.

2.2 Studying Intra/Inter-Racial Pairings

We examine whether the models exhibit prejudice against interracial romantic relationships when making predictions. We collect another set of 80 first names that are both strongly race- and gender-indicative, evenly distributed among four races/ethnicities and two genders (details described in §C.2). We perform pairwise name-replacements using these 80 names for the 327 test samples to analyze the relationship predictions among different intra/inter-racial name pairs.

We defer details related to full prompt used and model output parsing to §A.

3 Findings

Same-gender relationships are less likely to be predicted as romantic than different-gender ones. We observe a significant variation in recall of romantic relationship predictions from Llama2-7B (see Figure 2) for name-replacements involving different (top-right, and bottom-left)- versus same-gender pairs. This reveals that the model conservatively predicts romantic relationships when both the characters have names associated with the same gender (top-left – both male; bottom-right – both female). However, the precision across all races ranges between 0.78 – 0.84 (see Figure 5 in appendix). Such (relatively) low difference indicates that, while the model makes precise romantic predictions across all gender assignments and races, romantic predictions are more likely for contrastive gender assignments. Higher recall (Figure 2) for both female (bottom-right) replacements than both

male (top-left) across all races indicates a potential **stronger heteronormative bias against both male than both female pairs**. This could potentially be an effect of associating female names with romantic relationships as indicated by higher recall for female-neutral than male-neutral pairs. To test this hypothesis, we substitute one speaker’s name with a male, female or neutral name while keeping the other anonymized (substituting with “X”). We find that name pairs containing one female name tend to have higher recall than those containing one male name (Table 4 in appendix). This could either be due to a stronger association of female names with romantic relationships in general, or stronger heteronormative bias against male-male romantic relationships if models are (effectively) marginalizing probabilities over the anonymous character. A possible explanation for the former is that women tend to be portrayed only as objects of romance in fictional works, *e.g.*, as popularly evidenced by the failure of many movies to pass the Bechdel test (Agarwal et al., 2015).

The smaller gap in the recall between both female (bottom-right) name-replacements and different-gender (top-right and bottom-left) ones for Asian and Hispanic as compared to White and Black may result from model’s inability to discern gender from Asian and Hispanic names as accurately as for White and Black names. Figures 6 and 7 (appendix) show similar trends for Llama2-13B and Mistral-7B, respectively.

The unnaturalness of movie scripts with name and gender substitutions could, in theory, provide an alternative explanation for the observed biases, but the evidence shows this is not the cause. As female characters may speak differ-

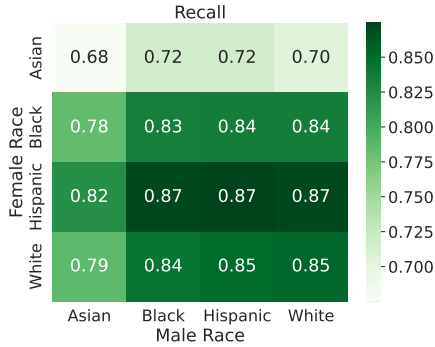


Figure 3: Recall of predicting romantic relationships from Llama2-7B for subset of the dataset where characters have different genders and are replaced with names associated with different races/ethnicities.

ently from male characters, our name-replacements can introduce statistical inconsistency between the gender associated with a character name and the style or content of the lines they speak, potentially confounding our observations. However, comparable recall between name-replacements that preserve the gender (upper-triangle; specifically top-right) associated with the original speakers and the swapped variants (lower-triangle; specifically bottom-left) in Figure 2, indicates that swapping both characters’ genders has minimal impact on model’s performance in the conversations we used. Hence, we conclude the potential inconsistency between gender and linguistic content is not a major confounding factor.

Character pairs involving Asian names have lower romantic recall; however, we do not find strong evidence against interracial pairings. While Llama2-7B has similar precision of predicting a romantic relationship across all racial pairs (0.80 – 0.82, shown in Figure 8 in appendix), Figure 3 shows name pairs involving at least one Asian name have significantly lower recall. Noticeably, the recall is the lowest (0.68) when both character names are associated with Asian. Although there are variations in recall values among different racial setups, we do not observe disparate differences between interracial and intraracial name pairs for non-Asian names. Results for Llama2-13B and Mistral-7B, shown respectively in Figure 9 and 10 in the appendix, demonstrate a similar trend that Asian names lead to substantially lower recall values. Such systematically worse performance on Asian names potentially perpetuates known algorithmic biases (Chander, 2016; Akter et al., 2021; Papakyriakopoulos and Mboya, 2023).

Race/Ethnicity		Asian	Black	Hispanic	White
Gender	Logistic regression	53.3±12.7	96.4±2.9	80.5±13.0	99.9±0.2
	Majority baseline	54.2±0.0	54.2±0.0	54.2±0.0	53.9±0.3
Race	Logistic regression	97.6±1.9	70.5±6.3	89.5±4.1	94.2±3.8
	Majority baseline	50.6±0.2	50.6±0.4	50.9±0.4	50.9±0.3

Table 1: Logistic regression classification accuracy (%) of predicting the demographic attributes associated with a name from Llama2-7B contextualized embeddings.

4 Analysis and Discussion

We perform additional experiments to understand the observed model behavior.

Why does a model tend to predict fewer romantic relationships for racial pairings that involve Asian names?

Although we select names for each race that have strong real-world statistical associations with one gender, we hypothesize that low recall on pairs with one or more Asian names may be due to model’s inability to discern gender from Asian names. To test this hypothesis, we retrieve the contextualized embeddings from Llama2-7B for each first name (collected in §2.2) occurrence in 15 romantic and 15 non-romantic random dialogues. We obtain 209, 800 embeddings, which are used to train logistic regression models that classify the gender or race associated with a name (details in §A). As we compare the average classification accuracy (across 5 different train-test splits) against a majority baseline, we observe, in Table 1, that gender could be effectively predicted for non-Asian name embeddings, and the embeddings are distinguishable by race for all races/ethnicities in a One-vs-All setting. However, Asian name embeddings encode minimal gender information, decreasing the likelihood of a model leveraging the inferred gender identity when making relationship predictions that reflect heteronormative biases.

Does gender association have a stronger influence on model’s prediction than race/ethnicity?

We hypothesize that models’ tendency to associate gender with names influences their relationship predictions. To test this, we substitute names with generic placeholders (“X” and “Y”) to get a baseline where a model has no access to character names (more details in §B). After name-replacements, any deviation from these results (Table 2) would indicate that a model exploits the implicit information from first names. In Figure 2, multiple settings have recall values that significantly differ from those in the anonymized

Model	Precision	Recall	F1	Accuracy
Gender Pairings				
Llama2-7B	0.7978	0.6887	0.7392	0.6125
Llama2-13B	0.8649	0.3019	0.4476	0.4170
Mistral-7B	0.8269	0.2028	0.3258	0.3432
Racial Pairings				
Llama2-7B	0.8063	0.7131	0.7569	0.6422
Llama2-13B	0.8696	0.3287	0.4665	0.4404
Mistral-7B	0.8406	0.2311	0.3625	0.3761

Table 2: Evaluation scores for anonymous name-replacements (character replaced with “X” or “Y”) for different models under study. These results depict the model’s performance solely based on the context.

setting (0.6887). This disparity suggests name-replacements introduce gender information that significantly influences the model behavior. Such trends are less prominent for Asian names due to the model’s apparent inability to distinguish gender information in Asian names (Table 1). By contrast, racial information encoded in first names exerts a lesser impact. Non-Asian heterosexual intra/inter-racial pairs give rise to similar recall in Figure 3. We thus do not observe strong prejudice against interracial romantic relationships here.

5 Social Implications

It has been a prolonged and arduous struggle to recognize and accept gay marriages in the US (Andersen, 2016; Duberman, 2019). Legal recognition of these relationships remains a challenge in many other countries (Lee and Ostergard Jr, 2017; Chia, 2019; Ramdas, 2021). Even within the US, LGBTQIA+ people still encounter discrimination (Buist, 2019; Knauer, 2020; Naylor, 2020).

We believe heteronormative biases we have observed could impact various downstream LLM use cases, potentially causing both representational and allocational harms (Blodgett et al., 2020). For example, when LLMs are used for story generation based on social media posts as the premise (Te et al., 2018; Li et al., 2024a), the life events of members of the LGBTQIA+ community may be overlooked or misrepresented. If LLMs struggle to recognize same-gender romantic relationships, they may further marginalize the LGBTQIA+ community by diminishing their social visibility and representation. In addition, such model behavior may result in uneven allocation of resources or opportunities. Consider an online advertising system that

promotes low-interest home loans for married couples based on social media interactions. A model unable to identify same-gender marriages would exclude these couples from the promotion. Therefore, building inclusive technology that respects minority rights is essential.

6 Related Work

Prior works (Wang et al., 2022; Jeoung et al., 2023; Sandoval et al., 2023; Wan et al., 2023; An et al., 2023, 2024; Nghiem et al., 2024) show that language models often treat first names differently, even with controlled input contexts, due to factors like frequency and demographic attributes associated with names (Maudslay et al., 2019; Shwartz et al., 2020; Wolfe and Caliskan, 2021; Czarnowska et al., 2021; An and Rudinger, 2023). Our work uses models’ interpretations of gender associated with first names to reveal heteronormative biases in some LLMs.

Further, NLP systems often fail in interpreting various social factors (*e.g.*, social norms, cultures, and relations) of language (Hovy and Yang, 2021). One such factor of interest is the representation of social relationships in these systems, including power dynamics (Prabhakaran et al., 2012), friendship (Krishnan and Eisenstein, 2015), and romantic relationships (Seraj et al., 2021). Recently, Stewart and Mihalcea (2024) show failure of popular machine translation systems in translating sentences concerning relationships between nouns of same-gender. Leveraging the task of relationship prediction and using an existing dataset (Jia et al., 2021), our work contributes to the assessment of social relationship-related biases in LLMs arising from gender and race associations with first names.

7 Conclusion

Through controlled name-replacement experiments, we find that LLMs predict romantic relationships between characters based on the demographic identities associated with their first names. Specifically, relationship predictions between same-gender and intra/inter-racial character pairs involving Asian names are less likely to be romantic. Our analysis of contextualized name embeddings sheds light on the cause of our findings. We also highlight the social implications of this potentially harmful model behavior for the LGBTQIA+ community. We urge advocates to build technology that respects the rights of marginalized social groups.

Limitations

Prompt sensitivity and in-context learning.

LLMs are sensitive to prompt formats (Min et al., 2022; Li et al., 2024b) therefore the accuracy of predictions may vary within or across models. While we had experimented with several prompts before converging to the one we use (gave the best prediction accuracy on the original dataset as well as close to that reported in Jia et al. (2021)), future work may investigate the impact of different prompt formulations and if in-context learning can help in reducing the influence of biases on the downstream tasks.

Inadequate coverage of names associated with different identities.

We recognize that our paper has limitations regarding the number of races and genders studied. This is due to the unavailability of data sources to compile a sufficiently large number of names strongly associated with a wide range of underrepresented races and gender identities.

Linguistic usage might be significantly different in same-gender romantic relationships.

The test set we have utilized (Jia et al., 2021) does not contain dialogues between same-gender character pairs in romantic relationships. As a consequence, we lack conversations that effectively depict interactions between same-gender partners. We acknowledge this limitation in our data source. However, in cases where same-gender partners exhibit behavior similar to different-gender couples, our results indicate that LLMs tend to demonstrate heteronormative biases in the intersection of these interaction styles.

Conversations might contain implicit gender-revealing cues.

While we ensure consistency between gender associated with an utterance (based on how a male speaks vs a female) and the gender associated with a name by only considering the conversations that do not have explicit gender-revealing cues as described in §2.1, we acknowledge the possibility of the presence of implicit gender-revealing cues which is harder to detect. However, we believe that our findings stand valid even if the implicit cues are present as demonstrated by comparable recall between name-replacements that preserve the gender (upper-triangle; specifically top-right) associated with the original speaker and the swapped variants (lower-triangle; specifically bottom-left) in Figure 2. We

leave further analysis of the nuances with implicit cues to future work.

Ethical Considerations

Inconsistency between self-identification and demographic attributes associated with a name.

Our categorization of names into subgroups of race/ethnicity and gender is based on real-world data as we observe a strong statistical association between names and demographic attributes (race/ethnicity and gender). However, it is crucial to realize that a person with a particular name may identify themselves differently from the majority, and we should respect their individual preferences and embrace the differences. We have attempted to accommodate diverse possibilities in self-identification by incorporating gender-neutral names into our experimental setup. While there is still ample room for improvement in addressing this issue, we have taken a step forward in promoting the inclusion of additional forms of self-identification in ethical NLP research.

Ethical concerns about the task of relationship prediction.

Predicting interpersonal relationships from conversations may require access to private and sensitive data. If no proper consent from a user is obtained, using personal data could lead to serious ethical and legal concerns. Although building systems that identify the relationship type between speakers could contribute to the development of AI agents that better understand human interactions, it is crucial to be transparent about what data is collected and how it is processed in such systems. Even if data privacy is properly handled when using a model to predict relationship types, people often exercise caution when revealing romantic relationships. Therefore, the deployment of an NLP system to identify such relationships should be disclosed to users who may be affected, and any predictions should remain confidential unless the user’s consent is obtained for public disclosure.

Acknowledgements

We would like to thank the anonymous reviewers for their valuable feedback. Rachel Rudinger is supported by NSF CAREER Award No. 2339746. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

- Apoorv Agarwal, Jiehan Zheng, Shruti Kamath, Sri-rankumar Balasubramanian, and Shirin Ann Dey. 2015. [Key female characters in film have more to talk about besides men: Automating the Bechdel test](#). In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 830–840, Denver, Colorado. Association for Computational Linguistics.
- Shahriar Akter, Grace McCarthy, Shahriar Sajib, Katina Michael, Yogesh K. Dwivedi, John D’Ambra, and K.N. Shen. 2021. [Algorithmic bias in data-driven innovation in the age of ai](#). *International Journal of Information Management*, 60:102387.
- Haozhe An, Christabel Acquaye, Colin Wang, Zongxia Li, and Rachel Rudinger. 2024. [Do large language models discriminate in hiring decisions on the basis of race, ethnicity, and gender?](#) In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 386–397, Bangkok, Thailand. Association for Computational Linguistics.
- Haozhe An, Zongxia Li, Jieyu Zhao, and Rachel Rudinger. 2023. [SODAPOP: Open-ended discovery of social biases in social commonsense reasoning models](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 1573–1596, Dubrovnik, Croatia. Association for Computational Linguistics.
- Haozhe An and Rachel Rudinger. 2023. [Nichelle and nancy: The influence of demographic attributes and tokenization length on first name biases](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 388–401, Toronto, Canada. Association for Computational Linguistics.
- Ellen Ann Andersen. 2016. Transformative events in the lgbtq rights movement. *Ind. JL & Soc. Equal.*, 5:441.
- Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. 2020. [Language \(technology\) is power: A critical survey of “bias” in NLP](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5454–5476, Online. Association for Computational Linguistics.
- Carrie L Buist. 2019. Lgbtq rights in the fields of criminal law and law enforcement. *U. Rich. L. Rev.*, 54:877.
- Yang Trista Cao, Anna Sotnikova, Hal Daumé III, Rachel Rudinger, and Linda Zou. 2022. [Theory-grounded measurement of U.S. social stereotypes in English language models](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1276–1295, Seattle, United States. Association for Computational Linguistics.
- Anupam Chander. 2016. The racist algorithm. *Mich. L. Rev.*, 115:1023.
- Myra Cheng, Esin Durmus, and Dan Jurafsky. 2023. [Marked personas: Using natural language prompts to measure stereotypes in language models](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1504–1532, Toronto, Canada. Association for Computational Linguistics.
- Joy L Chia. 2019. Lgbtq rights in china: Movement-building in uncertain times. In *Handbook on human rights in China*, pages 657–680. Edward Elgar Publishing.
- Paula Czarnowska, Yogarshi Vyas, and Kashif Shah. 2021. [Quantifying social biases in NLP: A generalization and empirical comparison of extrinsic fairness metrics](#). *Transactions of the Association for Computational Linguistics*, 9:1249–1267.
- Shaji Daniel. 2024. [Negotiating the challenges of an interracial marriage: An interpretive phenomenological analysis of the perception of diaspora indian partners](#). *Family Relations*, 73(1):282–297.
- Martin Duberman. 2019. *Stonewall: The definitive story of the LGBTQ rights uprising that changed America*. Penguin.
- Virginia Felkner, Ho-Chun Herbert Chang, Eugene Jang, and Jonathan May. 2023. [WinoQueer: A community-in-the-loop benchmark for anti-LGBTQ+ bias in large language models](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9126–9140, Toronto, Canada. Association for Computational Linguistics.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2019. The curious case of neural text degeneration. *arXiv preprint arXiv:1904.09751*.
- Dirk Hovy and Diyi Yang. 2021. [The importance of modeling social factors of language: Theory and practice](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 588–602, Online. Association for Computational Linguistics.
- Sullam Jeoung, Jana Diesner, and Halil Kilicoglu. 2023. [Examining the causal impact of first names on language models: The case of social commonsense reasoning](#). In *Proceedings of the 3rd Workshop on Trustworthy Natural Language Processing (TrustNLP 2023)*, pages 61–72, Toronto, Canada. Association for Computational Linguistics.
- Qi Jia, Hongru Huang, and Kenny Q Zhu. 2021. [Ddrel: A new dataset for interpersonal relation classification](#)

- in dyadic dialogues. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 13125–13133.
- 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. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.
- Nancy J Knauer. 2020. The lgbtq equality gap and federalism. *Am. UL Rev.*, 70:1.
- Vinodh Krishnan and Jacob Eisenstein. 2015. “you’re mr. lebowski, i’m the dude”: Inducing address term formality in signed social networks. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1616–1626, Denver, Colorado. Association for Computational Linguistics.
- Chelsea Lee and Robert L Ostergard Jr. 2017. Measuring discrimination against lgbtq people: A cross-national analysis. *Human Rights Quarterly*, pages 37–72.
- Donna A. Lewandowski and Linda A. Jackson. 2001. Perceptions of interracial couples: Prejudice at the dyadic level. *Journal of Black Psychology*, 27(3):288–303.
- Junyi Li, Tianyi Tang, Wayne Xin Zhao, Jian-Yun Nie, and Ji-Rong Wen. 2024a. Pre-trained language models for text generation: A survey. *ACM Comput. Surv.*, 56(9).
- Zongxia Li, Ishani Mondal, Yijun Liang, Huy Nghiem, and Jordan Lee Boyd-Graber. 2024b. Pedants: Cheap but effective and interpretable answer equivalence.
- Rowan Hall Maudslay, Hila Gonen, Ryan Cotterell, and Simone Teufel. 2019. It’s all in the name: Mitigating gender bias with name-based counterfactual data substitution. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5267–5275, Hong Kong, China. Association for Computational Linguistics.
- Suzanne C. Miller, Michael A. Olson, and Russell H. Fazio. 2004. Perceived reactions to interracial romantic relationships: When race is used as a cue to status. *Group Processes & Intergroup Relations*, 7(4):354–369.
- Sewon Min, Xinxu Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022. Rethinking the role of demonstrations: What makes in-context learning work? In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11048–11064, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Lorenda A Naylor. 2020. *Social equity and LGBTQ rights: Dismantling discrimination and expanding civil rights*. Routledge.
- Huy Nghiem, John Prindle, Jieyu Zhao, and Hal Daumé III. 2024. "you gotta be a doctor, lin": An investigation of name-based bias of large language models in employment recommendations. *arXiv preprint arXiv:2406.12232*.
- Debora Nozza, Federico Bianchi, Anne Lauscher, and Dirk Hovy. 2022. Measuring harmful sentence completion in language models for LGBTQIA+ individuals. In *Proceedings of the Second Workshop on Language Technology for Equality, Diversity and Inclusion*, pages 26–34, Dublin, Ireland. Association for Computational Linguistics.
- Jiao Ou, Junda Lu, Che Liu, Yihong Tang, Fuzheng Zhang, Di Zhang, and Kun Gai. 2024. Dialogbench: Evaluating llms as human-like dialogue systems. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 6137–6170.
- Orestis Papakyriakopoulos and Arwa M. Mboya. 2023. Beyond algorithmic bias: A socio-computational interrogation of the google search by image algorithm. *Social Science Computer Review*, 41(4):1100–1125.
- Patricia S. Pittman, Claire Kamp Dush, Keeley J. Pratt, and Jen D. Wong. 2024. Interracial couples at risk: Discrimination, well-being, and health. *Journal of Family Issues*, 45(2):303–325.
- Amanda M Pollitt, Sara E Mernitz, Stephen T Russell, Melissa A Curran, and Russell B Toomey. 2021. Heteronormativity in the lives of lesbian, gay, bisexual, and queer young people. *Journal of Homosexuality*, 68(3):522–544.
- Vinodkumar Prabhakaran, Owen Rambow, and Mona Diab. 2012. Predicting overt display of power in written dialogs. In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 518–522, Montréal, Canada. Association for Computational Linguistics.
- Kamalini Ramdas. 2021. Negotiating lgbtq rights in singapore: The margin as a place of refusal. *Urban Studies*, 58(7):1448–1462.
- Corinne Reczek. 2020. Sexual-and gender-minority families: A 2010 to 2020 decade in review. *Journal of Marriage and Family*, 82(1):300–325.
- Evan TR Rosenman, Santiago Olivella, and Kosuke Imai. 2023. Race and ethnicity data for first, middle, and surnames. *Scientific Data*.
- Lisa Rosenthal and Tyrel J Starks. 2015. Relationship stigma and relationship outcomes in interracial and same-sex relationships: Examination of sources and buffers. *Journal of Family Psychology*, 29(6):818.

- Sandra Sandoval, Jieyu Zhao, Marine Carpuat, and Hal Daumé III. 2023. [A rose by any other name would not smell as sweet: Social bias in names mistranslation](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 3933–3945, Singapore. Association for Computational Linguistics.
- Sarah Seraj, Kate G Blackburn, and James W Pennebaker. 2021. Language left behind on social media exposes the emotional and cognitive costs of a romantic breakup. *Proceedings of the National Academy of Sciences*, 118(7):e2017154118.
- Vered Shwartz, Rachel Rudinger, and Oyvind Tafjord. 2020. [“you are grounded!”: Latent name artifacts in pre-trained language models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6850–6861, Online. Association for Computational Linguistics.
- Ian Stewart and Rada Mihalcea. 2024. [Whose wife is it anyway? assessing bias against same-gender relationships in machine translation](#). In *Proceedings of the 5th Workshop on Gender Bias in Natural Language Processing (GeBNLP)*, pages 365–375, Bangkok, Thailand. Association for Computational Linguistics.
- Robee Khyra Mae J. Te, Janica Mae M. Lam, and Ethel Ong. 2018. [Using social media posts as knowledge resource for generating life stories](#). In *Proceedings of the 32nd Pacific Asia Conference on Language, Information and Computation*, Hong Kong. Association for Computational Linguistics.
- Anna Tiginova, Paramita Mirza, Andrew Yates, and Gerhard Weikum. 2021. [PRIDE: Predicting Relationships in Conversations](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4636–4650, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, et al. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *arXiv preprint arXiv:2307.09288*.
- Juan Vásquez, Gemma Bel-Enguix, Scott Thomas Andersen, and Sergio-Luis Ojeda-Trueba. 2022. [Hetero-Corpus: A corpus for heteronormative language detection](#). In *Proceedings of the 4th Workshop on Gender Bias in Natural Language Processing (GeBNLP)*, pages 225–234, Seattle, Washington. Association for Computational Linguistics.
- Yixin Wan, George Pu, Jiao Sun, Aparna Garimella, Kai-Wei Chang, and Nanyun Peng. 2023. [“kelly is a warm person, joseph is a role model”: Gender biases in LLM-generated reference letters](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 3730–3748, Singapore. Association for Computational Linguistics.
- Jun Wang, Benjamin Rubinstein, and Trevor Cohn. 2022. [Measuring and mitigating name biases in neural machine translation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2576–2590, Dublin, Ireland. Association for Computational Linguistics.
- Robert Wolfe and Aylin Caliskan. 2021. [Low frequency names exhibit bias and overfitting in contextualizing language models](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 518–532, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Elena Zambelli. 2023. [Interracial couples and the phenomenology of race, place, and space in contemporary england](#). *Identities*, 30(5):725–743.
- Xian Zhao and Monica Biernat. 2019. [Your name is your lifesaver: Anglicization of names and moral dilemmas in a trilogy of transportation accidents](#). *Social Psychological and Personality Science*, 10(8):1011–1018.

A Detailed Experimental Setup

We present additional information about our experimental setup.

Models We use recently introduced two popular language models for testing our hypothesis, namely Llama (Touvron et al., 2023) (7B, 13B chat), and Mistral-7B (Jiang et al., 2023). Each model uses nucleus sampling (Holtzman et al., 2019) with default parameters, a temperature of 0, and a maximum generation length of 512. Each experiment over 327 test instances takes ~ 30 mins for Llama2-7B, ~ 1 hr for Llama2-13B, and ~ 25 mins for Mistral-7B. We ran 870 experiments per race (560 for Hispanic) for studying gender bias and 1600 experiments (400 per race-pair) for racial bias.

Computing Evaluation Scores We first compute precision, recall, F1, and accuracy scores for each name-pair-replacement and report the average scores for each name-pair bin, and each race-pair for studying the influence of gender, and race associated with names, respectively.

Dataset Statistics Table 3 presents the frequency of each relationship label along with romantic and non-romantic categories used for the purpose of this study, in the test split of DDRel (Jia et al., 2021) dataset. Out of 327 conversations with different-gender characters in the dataset, 271 do not contain explicit gender information.

Prompts We provide the prompt template used in our experiments in Figure 4.

Parsing Outputs from LLMs We observe inconsistencies in the outputs predicted by LLMs despite clear instructions regarding formatting. We use regular expressions to extract the JSON outputs and the predictions from them. We consider invalid outputs (*i.e.*, non-pre-defined class) from LLMs as a separate class (invalid) for evaluation purposes across all experiments.

Logistic Regression for Name Embeddings We quantitatively study the amount of gender information encoded in these embeddings by training a logistic regression model, separately for each race, to classify the gender associated with a name, using embeddings of 70% of names in a race as the training set and the remaining as the test set. Similarly, we train a logistic regression model to conduct a “One-vs-All” classification for each race. We control the train and test set in the racial setup to have

Relationship Labels	Frequency	Romantic	#Gender Neutral
Lovers	182	✓	155
Courtship	15	✓	12
Spouse	57	✓	46
Siblings	15	✗	13
Child-Other Family Elder	13	✗	7
Child-Parent	39	✗	11
Colleague/Partners	70	✗	59
Workplace Superior-Subordinate	48	✗	24
Professional Contact	27	✗	10
Opponents	20	✗	11
Friends	95	✗	83
Roommates	21	✗	21
Neighbours	8	✗	7
Total	610	-	459

Table 3: Frequency of relationship types in the test split of DDRel dataset (Jia et al., 2021).

a balanced number of positive and negative samples by down-sampling the instances from other races (1/3 from each other race). We repeat the logistic regression training with 5 different random train-test splits. We set the random state of the logistic regression model to 0 and maximum iteration to 1000. In Table 1, we report the average results across 5 runs with their standard deviation.

B Anonymous Name-replacement Experiments

We perform two types of anonymous name-replacement experiments differing in whether both names are anonymized or only one.

B.1 Both Names Are Anonymized

We substitute names with generic placeholders (“X” and “Y”) to get a baseline where a model has no access to character names to test the hypothesis that models’ tendency to associate gender with the names influences their relationship predictions.

B.2 One Name Is Anonymized

We substitute one name and keep the other anonymized to analyze the impact of one character’s gender on romantic relationship predictions independent of the second. We replace one name with a male, female or a neutral name either preserving or swapping the original gender of the non-anonymized name while keeping the other name anonymized. Male, neutral, and female names belong to 0 – 25, 25 – 75, and 75 – 100% bins, respectively. We report the recall scores for romantic relationship prediction (same/swapped) for different models in Table 4.

<p>System Prompt: You are an avid novel reader and a code generator. Please output in JSON format. No preambles.</p> <p>Prompt: Your task is to read a conversation between two people and infer the type of relationship between the two people from the given list of relationship types.</p> <p>Input: Following is the conversation between <i>{char_a}</i> and <i>{char_b}</i>.</p> <p><i>{context}</i></p> <p>What is the type of the relationship between <i>{char_a}</i> and <i>{char_b}</i> according to the below list of type of relationships: [Child-Parent, Child-Other Family Elder, Siblings, Spouse, Lovers, Courtship, Friends, Neighbors, Roommates, Workplace Superior - Subordinate, Colleague/Partners, Opponents, Professional Contact]</p> <p>Constraint: Please answer in JSON format with the type of relationship and explanation for the inferred relationship. Type of relationship can only be from the provided list.</p> <p>Output in JSON format:</p>
--

Figure 4: Prompt template used in our experiments. “*{char_a}*”, “*{char_b}*”, and “*{context}*” are placeholders here and they are instantiated with character names and dialogues accordingly for model inference.

Model	Race	Male	Neutral	Female
Llama2-7B	Asian	0.6049/0.6128	0.6085/0.6203	0.6663/0.6517
	Black	0.6069/0.6230	0.6454/0.6392	0.6572/0.6458
	Hispanic	0.6292/0.6284	0.6486/0.6541	0.7093/0.6897
	White	0.6387/0.6372	0.6328/0.6297	0.6887/0.6761
Llama2-13B	Asian	0.2991/0.2940	0.2806/0.2798	0.3090/0.3043
	Black	0.3066/0.2854	0.3004/0.2909	0.3054/0.3105
	Hispanic	0.3021/0.2801	0.2956/0.2980	0.3206/0.3190
	White	0.3149/0.2952	0.2924/0.2878	0.3121/0.3121
Mistral	Asian	0.1789/0.1694	0.1808/0.1840	0.1895/0.1906
	Black	0.1855/0.1828	0.1902/0.1871	0.1922/0.1859
	Hispanic	0.1986/0.1955	0.1848/0.1776	0.2048/0.1973
	White	0.1895/0.1836	0.1887/0.1871	0.1942/0.1922

Table 4: Recall scores (same/swapped) for romantic relationship predictions when one name is anonymous while another is either a male, neutral, or female name as per bins marked in Figure 2. The results show that models are more likely to predict a romantic relationship when one of the names is a female name.

C First Names

We detail the name selection criteria in our experiments. We also list all first names we have used in our experiments to study the influence of different gender and racial/ethnic name pairing.

C.1 First Names Used to Study the Influence of Gender Pairing

We first collect names that have frequency over 200 and have more than 80% of the population having that name identify themselves as a particular race (Asian, Black, Hispanic, and White) from Rosenman et al. 2023. Then, we partition these names into 10 non-linearly segmented bins (shown in Figure 2) based on the percentage of population that

has been assigned as female at birth using statistics from the Social Security Application dataset (SSA³). We randomly sample 3 names per bin totaling to 30 names per race⁴ for performing the replacements. We consider names belonging to a spectrum of female gender associations to ensure coverage of gender-neutral names.

We list all the names used in this set of experiments. We include the percentage of the population assigned female gender at birth in parentheses.

Asian Seung (0.00%), Quoc (0.00%), Dat (0.00%), Nghia (2.30%), Thuan (2.40%), Thien (2.70%), Hoang (6.40%), Sang (6.60%), Jun (9.60%), Sung (13.50%), Jie (17.30%), Wei (21.80%), Hyun (39.00%), Khanh (41.90%), Wen (44.60%), Hien (51.70%), An (54.80%), Ji (61.40%), In (80.80%), Diem (88.60%), Quyen (88.90%), Ling (91.30%), Xiao (91.50%), Ngoc (92.40%), Su (95.40%), Hanh (95.60%), Vy (97.00%), Eun (98.30%), Trinh (100.00%), Huong (100.00%)

Black Deontae (0.00%), Antwon (0.10%), Javonte (1.00%), Dejon (2.90%), Jamell (3.40%), Dijon (4.60%), Dashawn (5.80%), Deshon (6.20%), Pernell (8.30%), Rashawn (10.10%), Torrance (13.20%), Semaj (22.60%), Demetris (25.60%), Kamari (33.60%), Amari (42.00%), Shamari (56.10%), Kenyatta (57.10%), Ivory (59.30%), Chaka (76.20%), Ashante (89.40%),

³<https://www.ssa.gov/oact/babynames/>

⁴Except for Hispanic wherein we did not get any names in 5 – 10% bin and only 1 name in 25 – 50% bin.

Unique (89.90%), Kenya (92.20%), Nikia (93.80%), Akia (94.30%), Kenyetta (95.50%), Shante (96.40%), Shaunta (97.00%), Laquandra (100.00%), Lakesia (100.00%), Daija (100.00%)

Hispanic Nestor (0.00%), Fidel (0.00%), Raul (0.60%), Leonides (2.70%), Yamil (4.50%), Reyes (10.80%), Cruz (13.10%), Neftali (14.90%), Noris (38.10%), Nieves (62.40%), Guadalupe (72.60%), Ivis (75.00%), Monserrate (78.20%), Ibis (82.60%), Johanny (89.40%), Elba (91.50%), Matilde (93.40%), Rocio (96.90%), Lucero (97.30%), Cielo (97.50%), Lucila (100.00%), Zuleyka (100.00%), Yaquelin (100.00%)

White Zoltan (0.00%), Leif (0.10%), Jack (0.40%), Ryder (3.30%), Carmine (3.40%), Haden (4.10%), Tate (5.30%), Dickie (5.50%), Logan (7.40%), Parker (17.50%), Sawyer (20.90%), Hayden (22.50%), Dakota (29.70%), Britt (38.30%), Harley (41.70%), Campbell (53.90%), Barrie (56.10%), Peyton (61.90%), Kelley (88.00%), Jodie (88.20%), Leigh (88.70%), Clare (90.90%), Rylee (92.20%), Meredith (94.70%), Baylee (97.00%), Lacey (97.30%), Ardith (97.70%), Kristi (99.80%), Galina (100.00%), Margarete (100.00%)

C.2 First Names Used to Study the Influence of Intra/Inter-racial Pairing

By referencing [Rosenman et al. 2023](#) and the SSA dataset again, we collect another set of both race- and gender-indicative first names with a minimum frequency of 200, applying a threshold of 90% for the percentage of the population assigned either female or male at birth. For race threshold, we set it to be 90% for Asian, Black, and Hispanic, and 70% for White. Although we choose a lower threshold for White to account for the phenomenon of name Anglicization ([Zhao and Biernat, 2019](#)), we still obtain empirical results that strongly indicate these names are represented differently from names associated with other races/ethnicities. In total, we obtain 80 names that are evenly distributed among four races/ethnicities and two genders. We replace name-pairs while preserving the gender associated with the names in the original dialogue.

Asian Female Thuy, Thu, Huong, Trang, Ngoc, Hanh, Hang, Xuan, Trinh, Eun

Asian Male Tuan, Hai, Sang, Hoang, Nam, Huy, Quang, Duc, Trung, Hieu

Black Female Latoya, Ebony, Latasha, Latonya, Tamika, Kenya, Tameka, Lakeisha, Tanisha, Precious

Black Male Tyrone, Cedric, Darius, Jermaine, Demetrius, Malik, Jalen, Roosevelt, Marquis, Deandre

Hispanic Female Luz, Mayra, Marisol, Maribel, Alejandra, Yesenia, Migdalia, Xiomara, Mariela, Yadira

Hispanic Male Luis, Jesus, Lazaro, Osvaldo, Heriberto, Jairo, Rigoberto, Adalberto, Ezequiel, Ulises

White Female Mary, Patricia, Jennifer, Linda, Elizabeth, Barbara, Susan, Jessica, Kimberly, Sandra

White Male James, Michael, John, Robert, William, David, Christopher, Richard, Joseph, Charles

D Additional Results

We report the results for Llama2-13B (Figures 6 and 9) and Mistral-7B (Figures 7 and 10). We also report the F1 and accuracy scores for Llama2-7B, for completeness, in Figure 5 and 8. We observe similar trends as Llama2-7B discussed in the main body of the paper.

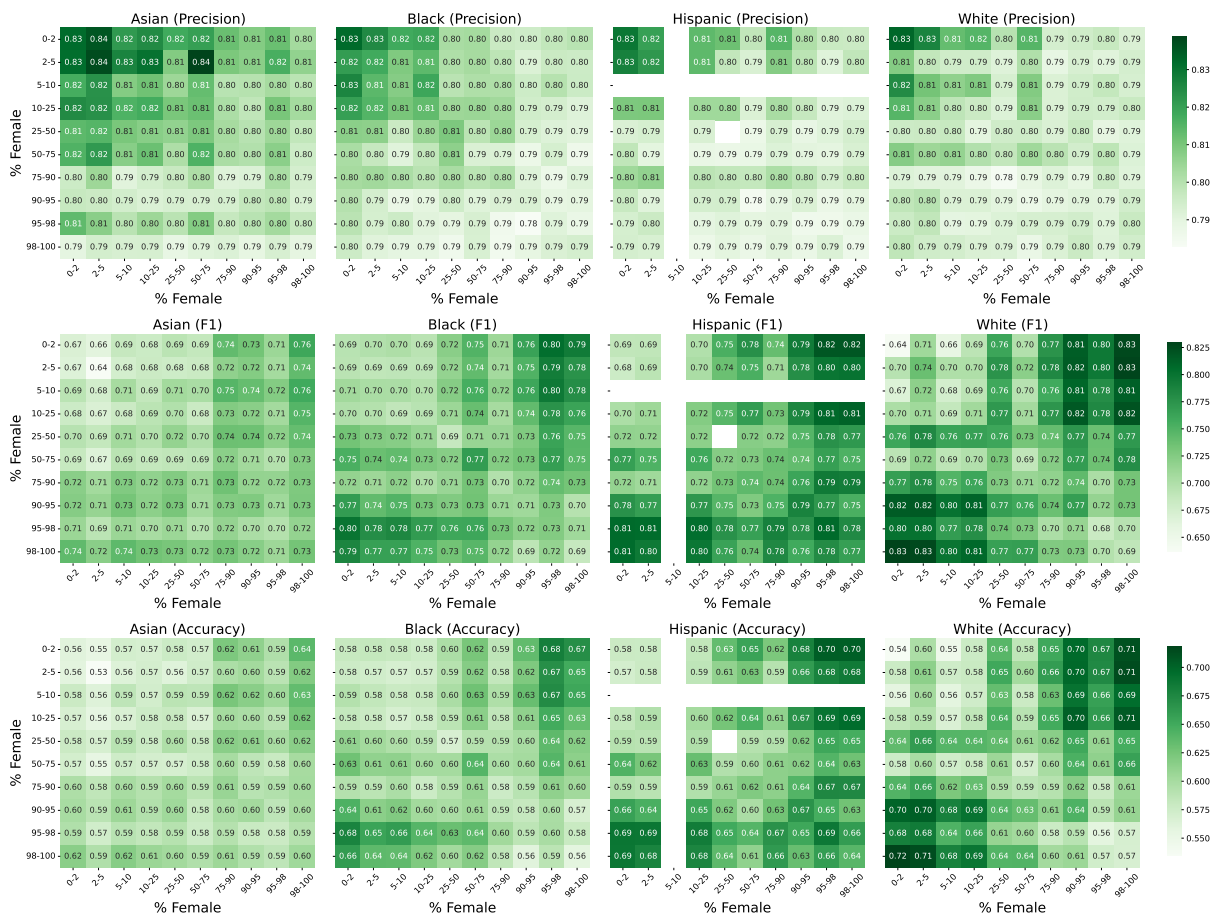


Figure 5: Precision, F1-score and Accuracy plots for romantic predictions from Llama2-7B model.




Figure 6: Precision, Recall, F1-score and Accuracy plots for romantic predictions from Llama2-13B model.

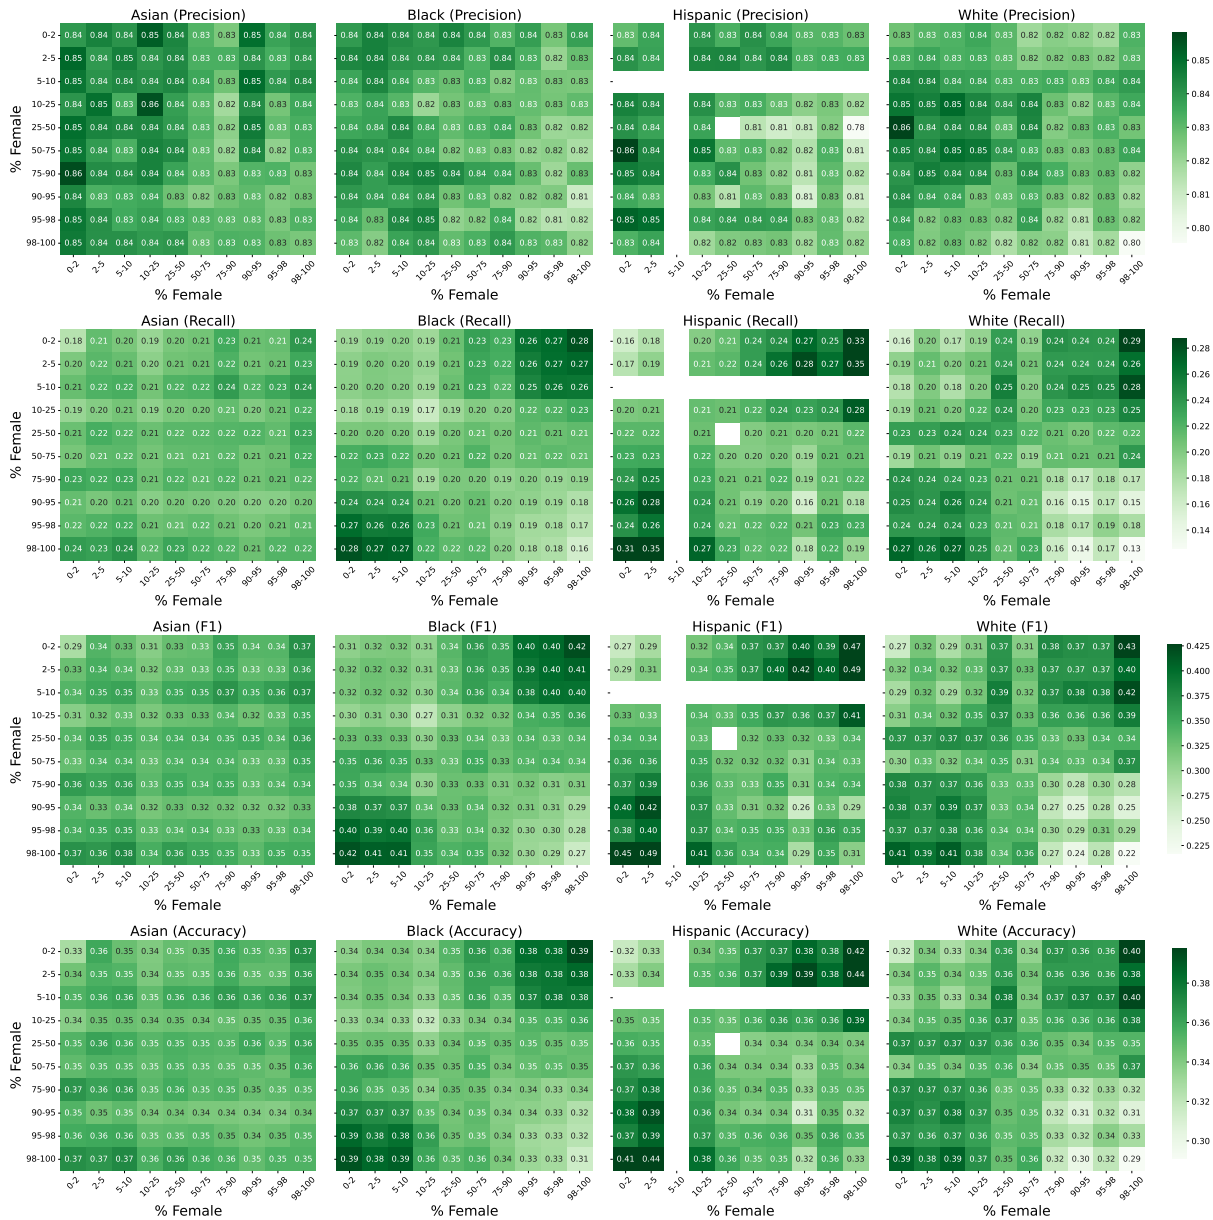


Figure 7: Precision, Recall, F1-score and Accuracy plots for romantic predictions from Mistral-7B model.




Figure 8: Precision, Recall, F1, and Accuracy of predicting romantic relationships from Llama2-7B for subset of the dataset where characters have different genders and are replaced with names associated with different races/ethnicities.




Figure 9: Precision, Recall, F1, and Accuracy of predicting romantic relationships from Llama2-13B for subset of the dataset where characters have different genders and are replaced with names associated with different races/ethnicities.




Figure 10: Precision, Recall, F1, and Accuracy of predicting romantic relationships from Mistral-7B for subset of the dataset where characters have different genders and are replaced with names associated with different races/ethnicities.