Robust Pronoun Fidelity with English LLMs: Are they Reasoning, Repeating, or Just Biased?

Vagrant Gautam[1] **Eileen Bingert**[1] **Dawei Zhu**[1] **Anne Lauscher**[2] **Dietrich Klakow**[1]
[1]Saarland University, Germany
[2]Data Science Group, University of Hamburg, Germany
vgautam@lsv.uni-saarland.de

Abstract

Robust, faithful, and harm-free pronoun use for individuals is an important goal for language model development as their use increases, but prior work tends to study only one or two of these characteristics at a time. To measure progress towards the combined goal, we introduce the task of *pronoun fidelity*: Given a context introducing a co-referring entity and pronoun, the task is to reuse the correct pronoun later. We present RUFF, a carefully designed dataset of over 5 million instances to measure robust pronoun fidelity in English, and we evaluate 37 model variants from nine popular families, across architectures (encoder-only, decoder-only, and encoder-decoder) and scales (11M-70B parameters). When an individual is introduced with a pronoun, models can mostly faithfully reuse this pronoun in the next sentence, but they are significantly worse with *she/her/her*, singular *they*, and neopronouns. Moreover, models are easily distracted by non-adversarial sentences discussing other people; even one sentence with a distractor pronoun causes accuracy to drop on average by 34 percentage points. Our results show that pronoun fidelity is not robust, in a simple, naturalistic setting where humans achieve nearly 100% accuracy. We encourage researchers to bridge the gaps we find and to carefully evaluate reasoning in settings where superficial repetition might inflate perceptions of model performance.

1 Introduction

Third-person pronouns (*he*, *she*, *they*, etc.) are words that construct individuals' identities in conversations (Silverstein, 1985). In English, these pronouns mark referential gender for the entity they are referring to, which can also index an individual's social gender, e.g., man, woman, nonbinary (Cao and Daumé III, 2020). Correctly using the pronouns an individual identifies with is important, as misgendering (including through incorrect pronoun use) can in the best case be a social faux pas (Stryker, 2017) and in the worst case, cause psychological distress, particularly to transgender individuals (McLemore, 2018).

Accordingly, it is important for large language models (LLMs) to use pronouns faithfully and without causing harm. To this end, many studies have explored how LLMs handle pronouns, showing that they stereotypically associate pronouns and occupations (Kurita et al., 2019), reason about co-referring pronouns and entities better when they conform to stereotypes (Tal et al., 2022), fail when exposed to novel pronoun phenomena such as neopronouns (Lauscher et al., 2023), and cannot consistently reuse neopronouns during generation (Ovalle et al., 2023). These shortcomings create differences in quality of service and cause representational harm, amplifying discrimination against certain pronoun users (Blodgett et al., 2020; Dev et al., 2021).

In work on LLM pronoun use, a question that has gone unexamined thus far is: ***How robust is model faithfulness to pronouns*** when discussing more than one person? To answer this question, we propose *pronoun fidelity* (§2), a new task to investigate realistic model reasoning about pronouns, and we introduce RUFF (§3), a novel, large-scale dataset of over 5 million instances, to evaluate this task. With this dataset, we present an analysis of pronoun fidelity across 37 variants from nine popular language model families covering architectures and scales, to investigate whether models are reasoning, repeating, or just biased.

First, we collect model pronoun predictions for occupations in the absence of context, to establish a ''bias baseline'' (§5). Next, we evaluate whether models can overcome their biased pronoun predictions when explicitly shown what pronoun to use in context (§6). All models are good at this task, but there are significant disparities across

Transactions of the Association for Computational Linguistics, vol. 12, pp. 1755–1779, 2024. https://doi.org/10.1162/tacl_a_00719
Action Editor: Hai Zhao. Submission batch: 5/2024; Revision batch: 7/2024; Published 12/2024.

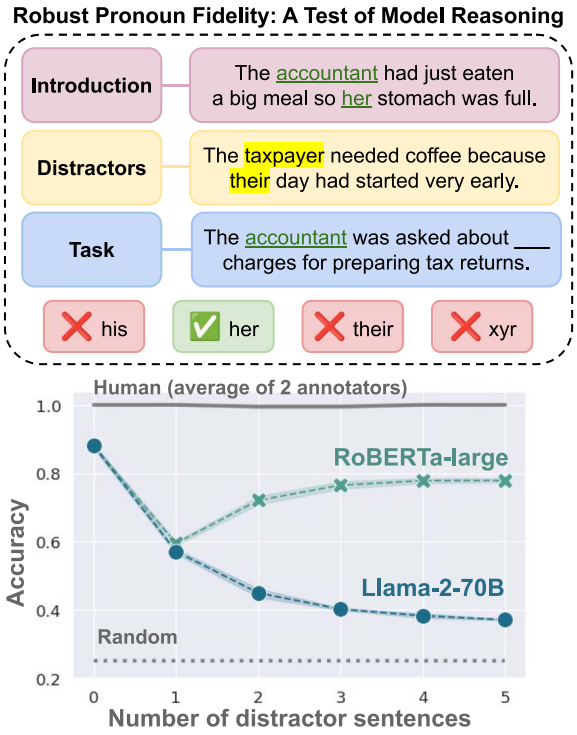

Figure 1: We evaluate model accuracy at using the correct pronoun for an entity when provided with an explicit introduction and 0-5 non-adversarial distractor sentences. LLAMA-2-70B and ROBERTA-LARGE show large accuracy drops with just one distractor. Accuracy is averaged over 3 data splits; standard deviation is shown with shading.

pronoun sets. We then test the robustness of this result by inserting naturalistic distractor sentences using a different pronoun to talk about another person (§7). *Even one non-adversarial distractor sentence vastly deteriorates model performance*, as shown in Figure 1. Finally, in a detailed error analysis (§8), we disentangle whether model errors can be attributed to distraction or falling back to bias, finding that encoder-only and decoder-only models behave in fundamentally different ways.

Overall, our results show that *models struggle to reason about pronouns in a simple, naturalistic setting* and highlight the need for careful task design to ensure that superficial repetition does not lead to inflated claims about model reasoning. We release all code and data to encourage researchers to bridge the gaps we find: `https://github.com/uds-lsv/robust-pronoun-fidelity`.

2 Pronoun Fidelity Task

Discussing multiple individuals is natural, frequent and well-studied in discourse; we use both definite references and pronouns in natural language to establish continuity and coherence (Grosz et al., 1995). We formalize a version of these phenomena in our task: Given a context in which a co-referring entity and pronoun are introduced, the task is to reconstruct the pronoun later in a sentence about the entity, independent of a limited number of potential distractors.

Introduction: *The accountant had just eaten a big meal so her stomach was full.*

(OPTIONAL)

Distractor 1: *The taxpayer needed coffee because their day had started very early.*

. . .

Distractor N: *Their sleep had been fitful.*

Task sentence: *The accountant was asked about ___ charges for preparing tax returns.*

More formally, an introduction sentence $i(e_a, p_a)$ establishes a coreference between an entity e_a and a pronoun p_a. A distractor sentence $d(e_b, p_b)$ explicitly establishes or implicitly continues a previously established coreference between a different entity e_b and a different pronoun p_b, i.e., $e_a \neq e_b$ and $p_a \neq p_b$. Let $\mathcal{D}(e_b, p_b)$ be a set of distractor sentences such that $0 \leq |\mathcal{D}(e_b, p_b)| \leq N$. When combined, an introduction sentence and the set of distractor sentences form a context. A task sentence $t(e_a, p)$ contains an unambiguous coreference between the entity e_a from the introduction and a pronoun slot p which must be filled. The task is to maximize

$$P[\, t(e_a, p = p_a) \mid i(e_a, p_a),\ \mathcal{D}(e_b, p_b)], \qquad (1)$$

the probability P of reconstructing the correct pronoun p_a in the sentence $t(e_a, p)$, given the context.

3 RUFF Dataset

To evaluate **R**obust pronoUn **F**idelity at scale, we create RUFF, an evaluation dataset of narratives. Each dataset instance describes a simple narrative with 1–2 people, but rather than narrative data that focuses on commonsense event reasoning (Mostafazadeh et al., 2016), we focus on pronominal reasoning, as in Rudinger et al. (2018). Specifically, we examine four third-person pronouns in three grammatical cases (nominative, accusative, and possessive dependent); in addition to

the English masculine (*he/him/his*) and feminine (*she/her/her*) pronouns, we heed Lauscher et al.'s (2022) call for more inclusive NLP research by examining two more pronoun sets that are less well-studied in NLP: singular *they* (*they/them/their*), the pronoun of choice of over 75% of respondents to the Gender Census (Lodge, 2023), and *xe/xem/xyr*, the most popular neopronoun according to the same census. Our narratives cover 60 occupations and corresponding participants (see Appendix A), following Winogender schemas (Rudinger et al., 2018), as their bias characteristics are well-studied in NLP. In total, RUFF contains over 5 million data instances. Each instance is designed to have an unambiguous answer, and is constructed with a 3-step pipeline: template creation (§3.1), template assembly (§3.2), and data validation (§3.3).

3.1 Template Creation

Below, we describe how we create occupation-specific task templates and generic context templates for introductions and distractors.

Task Templates. We create one task sentence template per occupation and grammatical case, with an unambiguous, unique coreference between the pronoun and occupation, for a total of 180 templates. For instance, *charges for preparing tax returns* can only belong to an *accountant*, never a *taxpayer*, the corresponding participant.

Context Templates. The ideal context template would be: **(1) flexible** across different occupations and participants, for a controlled setting to test robustness; **(2) cohesive** in a multi-sentence narrative leading up to the task template about an occupation; and **(3) neutral**, not dramatically affecting the prediction of a certain pronoun. Templates such as *He is an accountant* are well-established for testing word embedding associations (Caliskan et al., 2017; May et al., 2019). They are flexible and neutral (they are even referred to as ''semantically bleached'' templates in the literature), but it is unnatural to use more than one consecutively. Natural corpora like Levy et al. (2021) have the most potential for creating cohesive narratives, but contain occupation-specific sentences that are inflexible and sometimes also non-neutral, e.g., ungrammatical with singular *they*.

For a setting that satisfies all three criteria, we create context templates with generic themes, e.g., universal human emotions and sensations (*hungry/full*, *tired/energetic*, *unhappy/happy*, etc.). The generic themes make them flexible for use across all occupations and participants. Templates of the same polarity can be stacked into a cohesive narrative, e.g., a narrative about a taxpayer having a bad day after sleeping poorly and missing a meal. Our templates are created to be grammatical with all pronoun sets we consider, which satisfies neutrality. Additionally, our use of both positive and negative versions of templates (i.e., *happy* and *unhapppy*) as well as our variety of templates allows us to mitigate potential implicit biases when aggregated (Alnegheimish et al., 2022).

To reflect natural and coherent use of pronouns in discourse, we create explicit (definite reference + pronoun) and implicit (pronoun-only) versions of 10 context templates per grammatical case, for a total of 30 templates. Each explicit context template begins with an entity and introduces the pronoun in a clause, e.g., *The taxpayer needed coffee because their day had started very early*, while implicit templates are simple sentences like *Their sleep had been fitful*. See Appendix B for more detail on template creation and assembly.

3.2 Template Assembly

Figure 2 shows how we instantiate and combine templates to assemble our data instances: First, we select an occupation (e_a) and one of its task templates. We pick a pronoun (p_a) to use as ground truth and instantiate a random context template with the selected occupation and pronoun. The simplest version of the pronoun fidelity task includes just this introduction sentence followed by the task sentence. Instantiating 10 templates with 4 different pronoun sets and pairing them with task templates for 60 occupations across 3 grammatical cases gives us a total of 7,200 unique instances for this version of the task.

To create more complex data instances, we insert a variable number of distractor sentences between the introduction and task sentences, discussing a participant e_b with a *different* pronoun p_b. These are also sampled from the set of context templates (see Appendix B for details). Instantiating 4 templates with 3 previously unused pronouns gives 86,400 unique instances with one distractor.

Our stackable dataset design allows us to generate a vast amount of data of varying lengths,

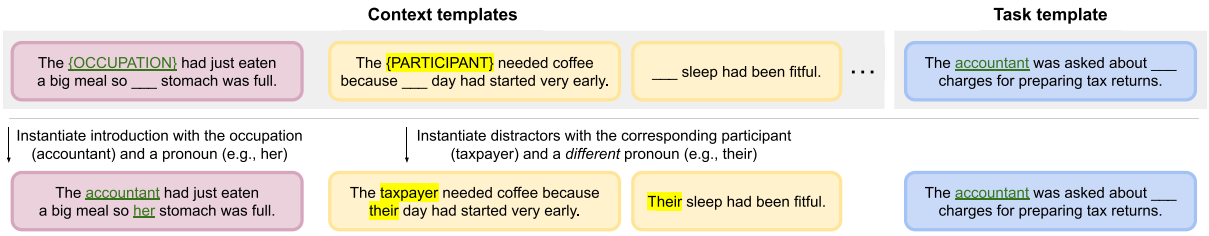

Figure 2: Template assembly for RUFF: Occupation-specific task templates are matched with generic context templates (introductions and optional distractors) that are instantiated with disjoint pronoun sets. This creates realistic but controlled narratives that allow us to measure robust pronoun fidelity.

Data type	Number of instances
With no context	
Task sentences	180
With introductory context	
+ 0 distractors	3×2,160 (of 7,200)
+ 1 distractor	3×2,160 (of 86,400)
+ 2 distractors	3×2,160 (of 345,600)
+ 3 distractors	3×2,160 (of 1,036,800)
+ 4 distractors	3×2,160 (of 2,073,600)
+ 5 distractors	3×2,160 (of 2,073,600)

Table 1: Number of dataset instances. Pronoun fidelity instances consist of task instances combined with introductory contexts and optional distractors. We subsample 3 sets of 2,160 sentences (of the total number of instances we created).

giving us a controlled setting to evaluate context effects on model predictions. We subsample the data with three random seeds for the rest of our evaluation, ensuring that all occupations, cases, pronoun declensions and distractor pronouns are equally represented in each subsampled set of 2,160 sentences. All data statistics are shown in Table 1.

3.3 Data Validation

We validate all task and context templates. To verify that the pronoun fidelity task is easy and unambiguous for humans, and to create a ceiling for model evaluation, we also validate a subset of task instances with 0–5 distractors. Annotator information is shown in Appendix C and all annotator instructions are provided in Appendix D.

Templates. Two authors with linguistic training iteratively created and validated sentence templates for grammaticality and correct coreferences until consensus was reached. An additional annotator independently rated 100% of the sentences as grammatical and with the correct coreferences.

Pronoun Fidelity Task. We sampled 100 instances with each possible number of distractors (0–5), for a total of 600 instances. One author and one annotator had to fill in the pronoun and they each performed with 99.8% accuracy.[1]

4 Experimental Setup

We list our models, evaluation methods, and metrics. Further details are provided in Appendix E.

4.1 Models

We experiment with 37 transformer-based language model variants from nine popular model families (see Table 2), which we chose to evaluate the effects of architecture and scaling. Our encoder-only models are from the BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019), ALBERT-v2 (Lan et al., 2020), and MosaicBERT (Portes et al., 2023) model families, as the first three remain well-used in NLP, and the last is trained on much more data. As for our decoder-only models, we select the popular Llama-2 (Touvron et al., 2023) model family, as well as OPT (Zhang et al., 2022) and Pythia (Biderman et al., 2023) for their large range of model sizes. In Appendix H, we also experiment with popular chat models that are further trained with instruction-tuning and reinforcement learning, to evaluate task performance with prompting; specifically, we use decoder-only Llama-2-chat models (Touvron et al., 2023) and encoder-decoder FLAN-T5 models (Chung et al., 2024).

[1]They disagreed on non-overlapping instances which appeared to be random slips.

Model	Sizes	Architecture
Evaluated with (Pseudo) Log Likelihoods		
ALBERT-v2	base (11M), large (17M), xlarge (58M), xxlarge (223M)	Encoder-only
BERT	base (110M), large (340M)	Encoder-only
RoBERTa	base (125M), large (355M)	Encoder-only
MosaicBERT	137M	Encoder-only
OPT	125M, 350M, 1.3B, 2.7B, 6.7B, 13B, 30B, 66B	Decoder-only
Pythia	14M, 70M, 160M, 410M, 1B, 1.4B, 2.8B, 6.9B, 12B	Decoder-only
Llama-2	7B, 13B, 70B	Decoder-only
Evaluated with prompting		
FLAN-T5	small (77M), base (248M), large (783M), xl (2.85B), xxl (11.3B)	Encoder-decoder
Llama-2-chat	7B, 13B, 70B	Decoder-only

Table 2: Models we experiment with across a range of sizes (11M-70B parameters) and architectures.

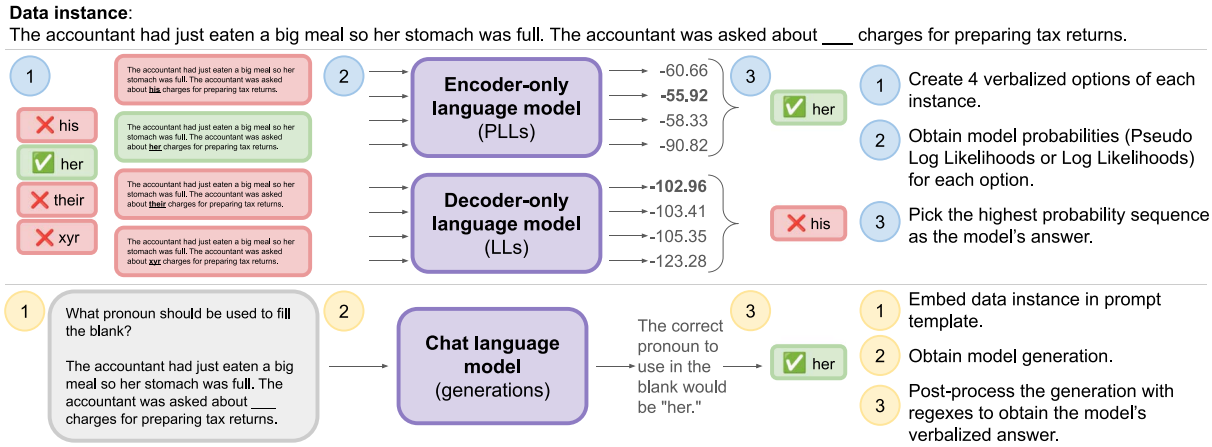

Figure 3: Model evaluation overview: pseudo log likelihoods (PLLs) and log likelihoods (LLs) of verbalized instances are used for encoder-only and decoder-only models; generations are used for chat models.

4.2 Obtaining Predictions

Figure 3 shows an overview of our evaluation methods. Decoder-only and encoder-only models are evaluated comparably in a forced choice setting: Following Hu and Levy (2023), we take direct measurements of probabilities as a proxy for models' metalinguistic judgments. Generations are obtained from chat models and post-processed to obtain unique pronouns, if any.

Encoder-only and Decoder-only Models. We verbalize four versions of each data instance, i.e., we fill in the blank with each of the four pronouns we consider, creating four options. We then obtain model probabilities for each of these four options, and select the highest probability option as the model's choice. We use log likelihoods for decoder-only models and pseudo log likelihoods for encoder-only models, following prior work (Salazar et al., 2020; Kauf and Ivanova, 2023). We do not use masked token prediction due to tokenization issues with neopronouns (Ovalle et al., 2024); briefly, we want *xe* to be tokenized ''normally'' (which is often as two tokens) rather than a single UNK token.

Chat Models. Following common practice, we evaluate chat models (FLAN-T5 and Llama-2-chat) using vanilla and chain-of-thought prompting. Following Sclar et al. (2024), we show the range of expected performance with 10 different prompts, inspired by the prompts to elicit coreferences in the FLAN collection (Longpre et al., 2023). See Appendix F for more methodological details and Appendix H for results.

4.3 Metrics

As every instance of the pronoun fidelity task has a unique correct answer, we report *accuracy*

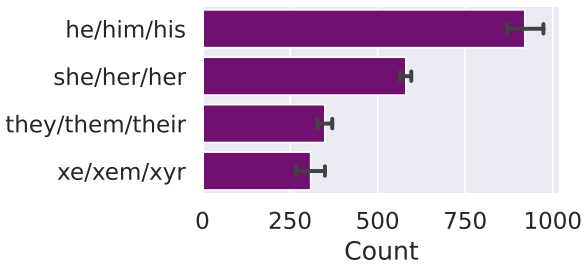

Figure 4: Counts of pronoun predictions from all models, in the absence of context. Error bars indicate standard deviation across models.

averaged over the three randomly sampled subsets of our dataset. We show the standard deviation with error bars or shading. Where possible, we perform significance testing with a Welch's t-test and a threshold of 0.05. We use human performance as our ceiling, and compare models to a baseline of randomly selecting 1 of the 4 pronouns (i.e., 25%).

5 Model Predictions with No Context

We begin by creating a ''bias baseline,'' i.e., obtaining pronoun predictions from models on our task sentences in the absence of any context. In Section 6, we will examine whether models can overcome this bias with reasoning when provided with context establishing a single correct answer.

Example: *The accountant was asked about ___ charges for preparing tax returns.*
No single answer (among *his*, *her*, *their*, *xyr*)

As we cannot evaluate accuracy on a task with no single correct answer, we show the counts of model predictions of different pronoun declensions in Figure 4, averaged over all models. Model-specific counts are shown in Appendix G. Even though our task sentences are designed such that any pronoun set can be used grammatically, all models tend to assign higher probability to *he/him/his* than other pronoun sets.

Obtaining pronoun predictions without context is a popular method to measure model bias, with numerous papers (Kurita et al., 2019, *inter alia*) showing that associations between occupations and pronouns are based on social gender stereotypes, e.g., *doctor-he* and *nurse-she*. However, model pronoun predictions might reflect dataset artifacts such as the choice of occupations, or be a statistical accident of the chosen templates (Seshadri et al., 2022). In addition, intrinsic biases may not correlate with actual pronoun use with context (Goldfarb-Tarrant et al., 2021). In order to test for such extrinsic behaviors, the rest of this paper examines whether models can override their intrinsic statistical biases on these same templates when provided with the right pronoun to use.

6 Injecting an Introductory Context

When models are provided with an introductory sentence explicitly establishing the pronoun to use for an entity, can they use that pronoun to refer to the same entity in the immediate next sentence?

Example: *The accountant had just eaten a big meal so her stomach was full. The accountant was asked about ___ charges for preparing tax returns.*
Correct answer: her

As Figure 5 shows, **all models perform better than chance at pronoun fidelity with a simple introduction** (up to 0.95 with MosaicBERT), but not as well as humans, who achieve perfect performance. We also see improvements with increasing model scale, with the exception of ALBERT-v2, as in Tay et al. (2023).

Which Pronouns Are Harder? Even in the simplest case of the pronoun fidelity task, patterns emerge when split by pronoun, as shown in Figure 6. Overall model **accuracy on *he/him/his* is significantly higher than *she/her/her*, which in turn is significantly higher than both *they/them/their* and *xe/xem/xyr***, in line with previous findings that language technology has gaps when it comes to neopronouns (Lauscher et al., 2023). Models show intriguing patterns with these last two pronoun sets. Most encoder-only models appear to handle the neopronoun better than singular *they* (e.g., BERT-large has an accuracy of 0.78 on *xe/xem/xyr* compared to 0.60 on *they/them/their*), which warrants further investigation. Decoder-only models smaller than 6.7B parameters struggle with the neopronoun, with every OPT and Pythia model smaller than 2.7B parameters performing below chance, and in some cases (e.g., Pythia-14M, Pythia-70M, and Pythia-60M) even performing close to 0.0. Beyond this scale, however, models perform better on *xe/xem/xyr* than on singular *they*, with Llama-13B achieving

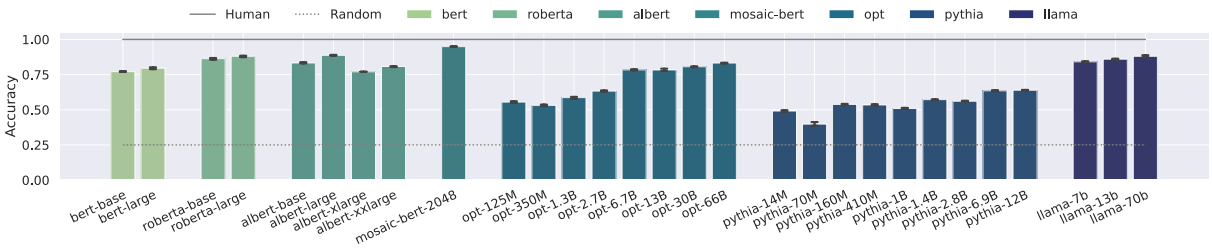

Figure 5: Pronoun fidelity by model with an introductory context. Accuracy is averaged across occupations, pronouns and grammatical cases, and is above chance (0.25) but below human performance (1.0).

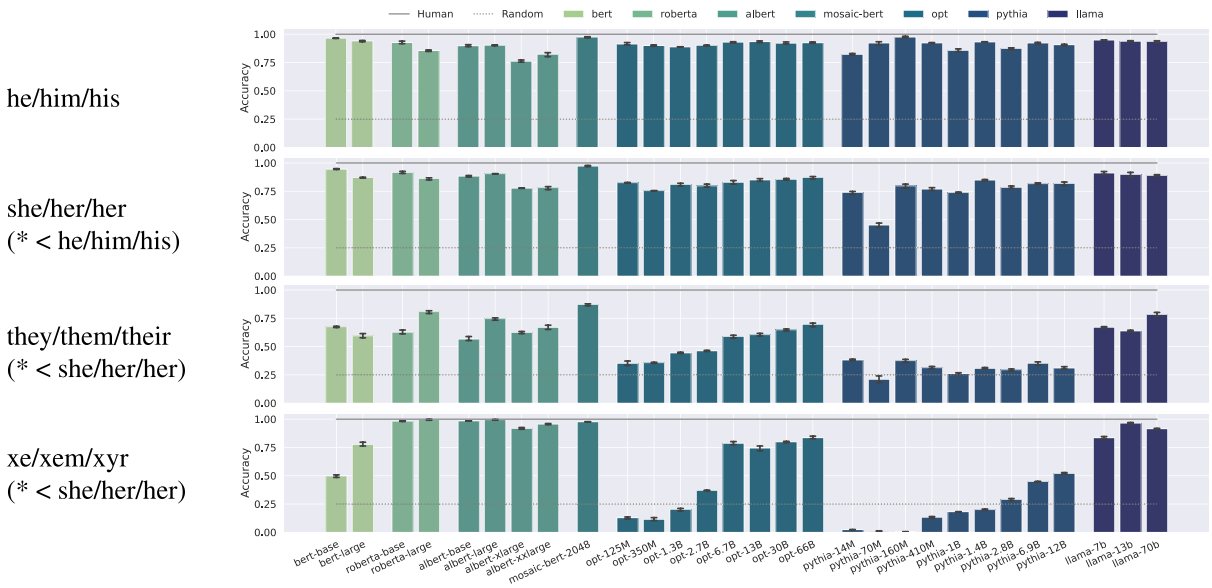

Figure 6: Pronoun fidelity by model with an introductory context, split by pronoun series. Model accuracy is compared to chance (0.25) and human performance (1.0). * Denotes statistical significance.

0.96 accuracy on the neopronoun. These differences are statistically significant. As the training data for individual model families is the same, this might suggest that decoder-only models generalize to novel pronouns starting at the scale of 6.7B parameters, but in light of Schaeffer et al. (2023), this result could just as well be a mirage resulting from our use of accuracy, a discontinuous metric. In either case, our observations could also explain the poor performance that some previous studies of neopronouns find, as the largest model that Hossain et al. (2023) experiment with, for instance, is OPT-6.7B. The lower performance of bigger models with singular *they* could also be a reflection of human processing difficulties with definite, specific singular *they*, as has been observed in linguistics (Conrod, 2019).

7 Adding Distractors

To further probe whether models actually ‘‘reason’’ when provided with context, we systematically inject sentences containing distractor pronouns between the introduction and the task, reflecting a natural usage scenario where multiple people are discussed with definite references and pronouns.

Example: *The accountant had just eaten a big meal so her stomach was full. The taxpayer needed coffee because their day had started very early. Their sleep had been fitful. The accountant was asked about ____ charges for preparing tax returns.*
Correct answer: her

Figure 7 shows that distractors degrade performance for all models. Encoder-only and decoder-only models show different performance curves as more distractors are added: All decoder-only models get steadily worse, whereas encoder-only models perform the worst with one distractor and then seem to slowly recover, never quite reaching

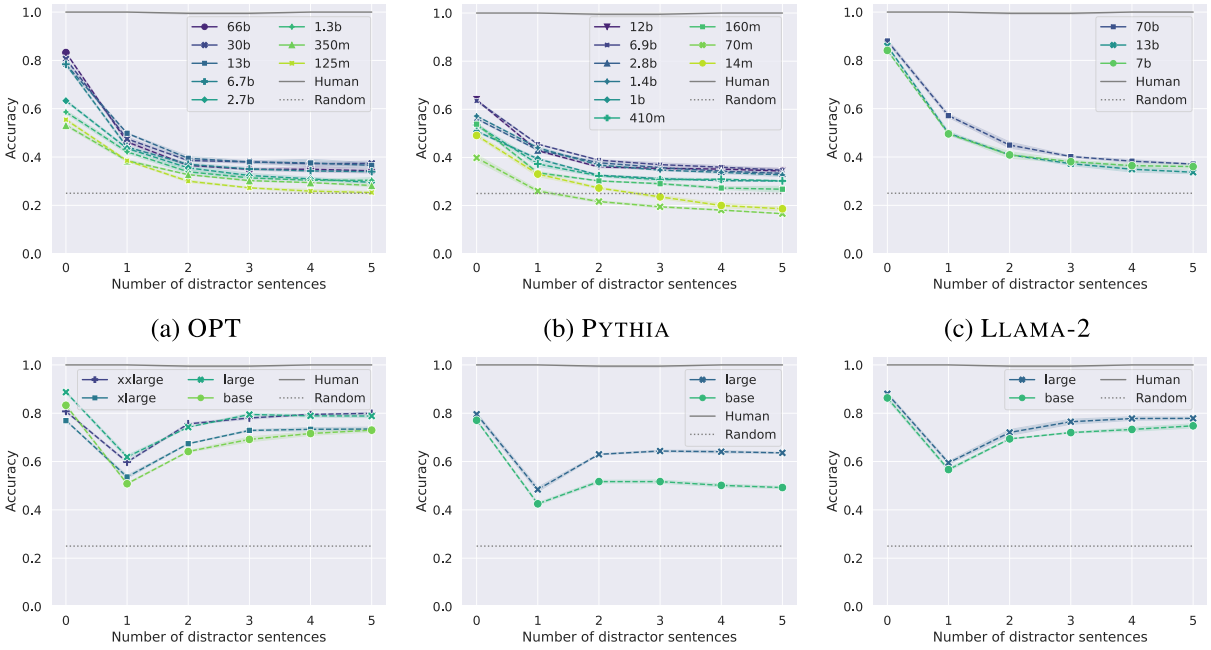

(a) OPT (b) PYTHIA (c) LLAMA-2

(d) ALBERT-V2 (e) BERT (f) ROBERTA

Figure 7: With more distractors, decoder-only models (above) get steadily worse; encoder-only models (below) get worse with one distractor and then recover, plateauing below their no-distractor accuracy.

their level of performance with no distractors. Scaling generally holds within model families, with larger models performing better with more distractors than smaller models of the same type. Figure 8 examines the interplay of scaling and architecture at a higher level, comparing results on the easiest case of pronoun fidelity (no distractors) with the hardest case (5 distractors). Surprisingly, **with no distractors, encoder-only models are much better than decoder-only models of the same scale**, and their performance is comparable to or better than decoder-only models that are orders of magnitude larger; ROBERTA-BASE (125M) is 0.86 accurate compared with OPT-125M's 0.55, and exceeds OPT-66B's 0.83 despite being more than 500 times smaller. In the hardest version of our task **with five distractors, encoder-only models are far better than *all* decoder-only models**, which show dramatically degraded performance; LLAMA-70B only achieves 0.37 accuracy, compared with MOSAICBERT's impressive 0.87. The lack of robustness of decoder-only models to distractors is striking, given that most state-of-the-art models today are decoder-only models. We hypothesize that architectural differences might explain the performance gaps; encoder-only models might use bidirectional attention to more closely relate the entity mentions in the introduction and task sentences. Training on next token prediction might also make decoder-only models prone to recency bias.

Using vanilla and chain-of-thought prompting (Appendix H) show the same patterns of degradation, reinforcing that **model pronoun fidelity is not robust**, and good performance with no distractors (§6) is likely not due to ''reasoning'' at all.

8 Distractibility versus Bias

In adding distractor sentences, we add distance from the introduction via additional tokens that might make the model forget the original occupation-pronoun association, and the distractor pronoun also acts as a competing token that the model might accidentally repeat. In this section, we focus on the error cases to disentangle whether models are ''forgetting'' and reverting to biased predictions from Section 5, or if they are actually being distracted. When a model gets the answer wrong, it is for one of three reasons: (1) distractibility, i.e., repeating the distractor pronoun, (2) bias, i.e., reverting to the model's context-free prediction, or (3) picking a completely different pronoun. Our example illustrates all three possibilities, and we hypothesize that the first two

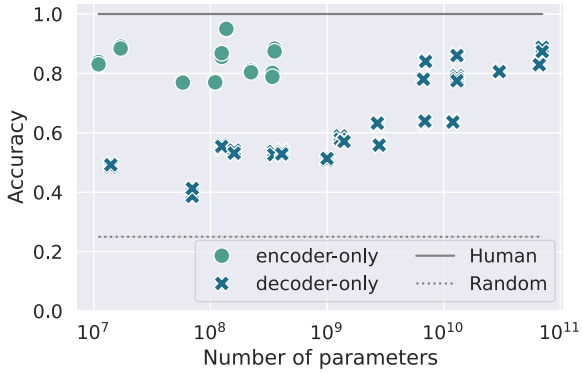

(a) With 0 distractors

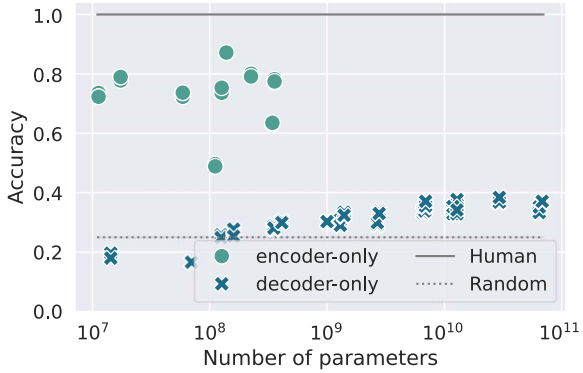

(b) With 5 distractors

Figure 8: Scaling behavior by architecture. With 0 distractors (above), encoder-only models are comparable to decoder-only models orders of magnitude larger. With 5 distractors (below), encoder-only models are far better.

possibilities are much more frequent than the third.

In cases where the distractor pronoun is the same as the model's context-free prediction, it is impossible to disentangle distractibility and bias just from the model's prediction. Hence, we exclude these and focus on the unambiguous error cases. As expected, we find that 74–93% of unambiguous model errors can be attributed to either model distractibility or bias.

Context-free (§5)
Example: *The accountant was asked about ___ charges for preparing tax returns.*
Prediction: his

With introduction and distractors (§7)
Example: *The accountant had just eaten a big meal so her stomach was full. The taxpayer needed coffee because their day had started very early. Their sleep had been fitful. The accountant was asked about ___ charges for preparing tax returns.*
Correct answer: her
Distraction error: their
Bias error: his
Other error: xyr

We first examine model distractibility, i.e., what percentage of errors are caused by models repeating the distractor pronoun instead of the correct pronoun. As expected, Figure 9 shows that across models, distraction is indeed the primary type of error for most models. **Decoder-only models get increasingly distracted with more distractors**, i.e., the proportion of errors due to distractor pronoun repetition steadily increases as distractors are added, saturating just below 85%. On the other hand, **encoder-only models seem to become *less* distractible** with the addition of more distractors. We know from the previous section that encoder-only models recover in their pronoun fidelity with 2–5 distractors, but here we measure distractibility as a percentage of all errors. Thus, a constant or increasing proportion of all the model errors could be due to distraction, and the fact that it *isn't* for encoder-only models is quite surprising! We leave it to future work to investigate whether this behavior relates to positional bias or context use.

As their proportion of distraction errors goes down, **encoder-only models increasingly revert to biased predictions**. With BERT-LARGE in particular, as soon as there is more than one distractor, the biggest proportion of errors is due to bias rather than distraction. BERT-LARGE appears more biased and less distractible than BERT-BASE, in contrast to all other models. Generally, larger models seem to be more distractible and revert to their bias less often, whereas smaller models are more biased and less distractible. Our findings on bias errors contrast with Tal et al. (2022), where larger models make a higher proportion of bias errors on a downstream task than smaller models. This might be due to our task having distractors, which seem to strongly influence model behavior in this setting.

The high distractibility of all models shows that **models are not robust reasoners**, and the contrast in error behavior between encoder-only and decoder-only models further highlights their differences. This shows that claims about decoder-only models should not be applied to all LLMs,

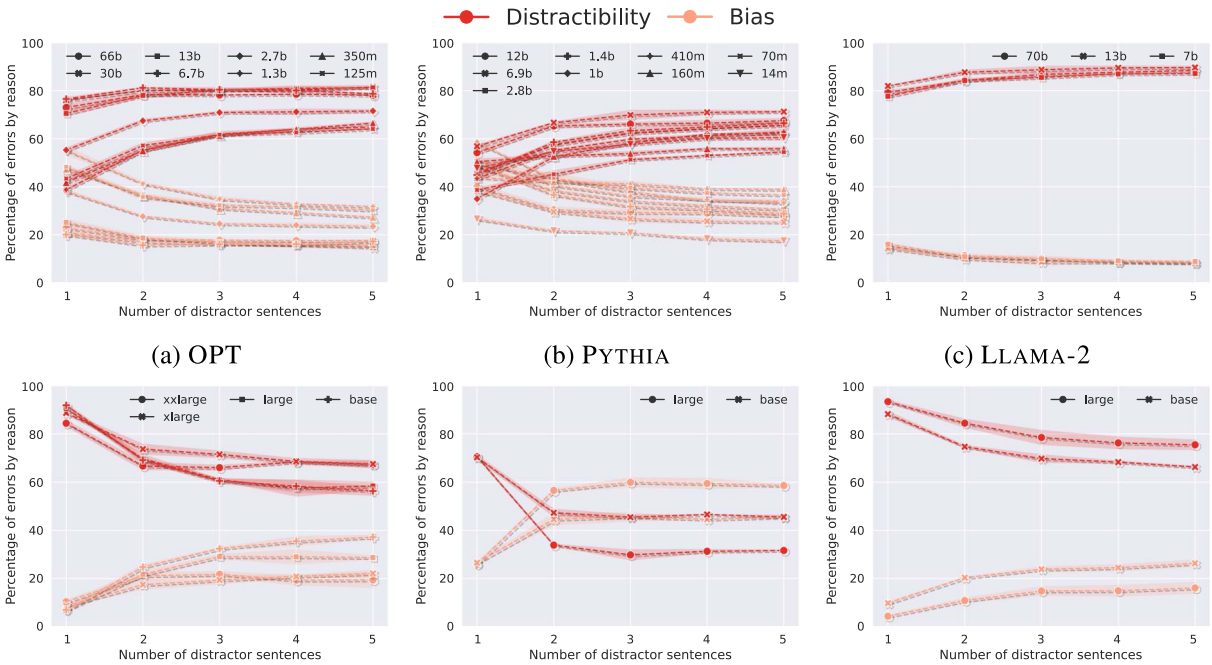

(a) OPT (b) PYTHIA (c) LLAMA-2

(d) ALBERT-V2 (e) BERT (f) ROBERTA

Figure 9: Trends in model distractibility (use of the distractor pronoun) and model bias (reverting to the context-free prediction). With more distractors, the proportion of errors due to distraction increases for decoder-only models (above) and decreases for encoder-only models (below).

and that reasoning must be evaluated carefully, accounting for the possibility of inflated performance due to shallow heuristics like repetition.

9 Discussion and Future Work

Our results show that even the biggest models of today are not up to the task of pronoun fidelity once it includes a single sentence discussing another person. All models are easily distracted, but encoder-only models and decoder-only models show very different patterns both in performance degradation with more distractors and their reasons for errors. Performance on this type of reasoning task should be evaluated carefully, with attention to how the overall patterns break down by different pronouns, and accounting for the possibility of repetition. Below we expand on some questions raised by our findings.

Improving Robust Pronoun Fidelity. A natural direction of future work is to solve the problem of robust pronoun fidelity, particularly in decoder-only models, which are unlikely to be replaced by encoder-only models with poorer generation abilities. A promising direction might be to encourage models to explicitly track associations between individuals and pronoun sets, just as people do. In fact, prior work has noted success with generative models when explicitly tracking mentions of entities across multiple tasks (Ji et al., 2017) and in the context of story generation (Fan et al., 2019). We urge researchers interested in this direction to treat RUFF as an evaluation dataset, as it was designed. Due to the presence of positional and associative heuristics (see the Limitations section), RUFF should not be seen as a source of data for fine-tuning or in-context learning, which is also why we do not run these experiments.

On ‘‘Reasoning.’’ Throughout the paper, we refer to ‘‘reasoning,’’ but this is inaccurate. Even the higher performance of encoder-only models cannot accurately be attributed to ‘‘reasoning’’ in the same way that we use this word for humans, as these models are not grounded in *meaning* from the real world (Bender and Koller, 2020). We use the word reasoning in line with other work in the field, but note that as these are all language models, it is more accurate to say that the way that decoder-only models model language is prone to repetition—or stochastic parroting (Bender et al., 2021)—of recent examples of the same word class, compared to encoder-only models.

***Why* Exactly Do We See the Patterns We See?** Our dataset design and error analysis shed light on model *behavior*, allowing us to evaluate different architectures comparably and disentangle the effects of repetition, distraction and statistical bias. However, it is beyond the scope of this paper to investigate where in the model architecture, neurons or pre-training data this comes from and what we can do about it towards improving reasoning and mitigating bias. Tools for model interpretability, e.g., attribution analysis, could help here, and are an important direction for future work.

Beyond our Dataset. Given the breadth of our task definition, future work could examine pronoun fidelity in other contexts, e.g., for participants, for names by extending Hossain et al. (2023), with differently ordered sentences, with real-world data as in Webster et al. (2018) and Levy et al. (2021), and in domains beyond simple narratives (Pradhan et al., 2013). Additionally, we evaluate on a version of this task that allows us to quantify repetition, i.e., the grammatical case of the elicited pronoun is the same as the case shown in the context. Examining model performance where a pronoun is shown in one grammatical case and then elicited in a different one would be interesting to probe syntactic generalization.

10 Related Work

Pronoun Fidelity. Hossain et al. (2023) and Ovalle et al. (2023) both study pronoun fidelity when models are prompted with a pronoun series to use for an individual, but they only consider simplistic pronoun use with no more than one person at a time. Although we look at within-language pronoun use, faithful pronoun use in context has also been studied in machine translation (Müller et al., 2018; Voita et al., 2018; Fernandes et al., 2023), where there is also a ground truth. Similar to our work, Sharma et al. (2022) inject context with an explicit coreference to encourage faithful pronoun translation. However, none of these papers explore the *robustness* of pronoun fidelity in the presence of distractors.

Reasoning with Pronouns. Most existing work about LLM reasoning with pronouns focuses on the task of coreference resolution, i.e., the ability to *identify* the connection between a pronoun and an entity, which may not translate to *faithful reuse* of that pronoun later, as in our work. Reasoning with pronouns typically uses Winograd schemas (Levesque et al., 2012; Abdou et al., 2020; Emelin and Sennrich, 2021), or Winograd-like schemas about named individuals (Webster et al., 2018; Zhao et al., 2018), or people referred to by their occupation (Rudinger et al., 2018; Levy et al., 2021). Most studies focus on *he* and *she*, but recent work has expanded to include singular *they* (Baumler and Rudinger, 2022) and neopronouns (Cao and Daumé III, 2021; Felkner et al., 2023), as we do.

Pronouns and Occupational Bias. Stereotypical associations between pronouns and occupations have been studied in masked token prediction (Kurita et al., 2019; de Vassimon Manela et al., 2021; Tal et al., 2022) and embeddings (Bolukbasi et al., 2016; Zhao et al., 2019), but these studies typically use brittle methodology (Gonen and Goldberg, 2019; Seshadri et al., 2022) and measure intrinsic bias, which may not translate to extrinsic bias or harms (Goldfarb-Tarrant et al., 2021). Unlike these works, we evaluate extrinsic bias and performance through our focus on natural pronoun use in context.

Robustness in Context. The impact of context on the robustness of language model reasoning has been investigated in many areas other than pronoun fidelity, e.g., negation (Gubelmann and Handschuh, 2022), linguistic acceptability (Sinha et al., 2023), natural language inference (Srikanth and Rudinger, 2022), and question answering (Liu et al., 2024; Levy et al., 2024).

11 Conclusion

We introduce the task of pronoun fidelity to evaluate robust, faithful and harm-free pronoun use in language models, and we present RUFF, a dataset we designed to evaluate it. We find evidence of faithful pronoun use only in a very simple setting, i.e., when only one person is discussed. Even here, models show significant performance disparities with neopronouns, singular *they* and *she/her/her*, compared to *he/him/his*. Even adding a single sentence about a second individual with a different pronoun causes accuracy to drop dramatically, showing that pronoun fidelity is neither robust to non-adversarial distractors nor due to ''reasoning.'' As more distractor sentences are added, encoder-only models perform better overall, but increasingly revert to biased predictions, while

decoder-only models get increasingly distracted. Our results show that in a setting that is very simple for humans, widely used large language models are unable to robustly and faithfully reason about pronouns, and continue to amplify discrimination against users of certain pronouns. We encourage researchers to bridge the performance gaps we report and to more carefully evaluate ''reasoning,'' especially when simple repetition could inflate perceptions of model performance.

12 Limitations

Shallow Heuristics. Much of the recent progress on reasoning datasets has been critically investigated and shown to often be a result of spurious correlations and dataset artifacts (Trichelair et al., 2019; Elazar et al., 2021). We caution readers that our dataset also gives a very *generous* estimate of model reasoning performance, as many of our task sentences are not ''Google-proof'' (Levesque et al., 2012), i.e., they can be solved with shallow heuristics such as word co-occurrences. Consider the following task sentence: *The janitor said not to step on the wet floor, otherwise ___ would have to mop it all over again. Janitor* is more strongly associated with *mop* than *child*, which could easily be exploited by models to solve the dataset without solving the task with something resembling ''reasoning.'' Another shallow heuristic that can be used to solve our current dataset is to simply return the first pronoun in the context, which happens to always be the correct answer. Our dataset design is flexible and allows for the creation of other orderings of sentences, but this is another example of why our dataset in its current form should only be used as an evaluation dataset, and models should not be pre-trained or fine-tuned with any splits of our data, nor provided with examples for in-context learning.

Whose Bias? Our task as it is defined in Section 2 is much broader than the scope of our dataset. We focus on occupations due to the wide attention they have received in prior literature, but we continue a long tradition of ignoring biases relating to the participants, e.g., *child*, *taxpayer*, etc. In addition, pronoun fidelity is only one dimension of inclusive language model behavior, and indeed only one way in which misgendering occurs in language, even in morphologically poor languages such as English.

Data Contamination. We take steps to prevent data contamination following Jacovi et al. (2023), including not releasing our data in plain text, and not evaluating with models behind closed APIs that do not guarantee that our data will not be used to train future models. However, as we cannot guarantee a complete absence of data leakage unless we never release the dataset, we encourage caution in interpreting results on RUFF with models trained on data after March 2024.

Acknowledgments

The authors thank Timm Dill for several rounds of patient annotation, as well as Aaron Mueller, Marius Mosbach, Vlad Niculae, Yanai Elazar, our action editor Hai Zhao, and our anonymous TACL reviewers, for feedback on math, plots and framing, that improved this work. Vagrant Gautam received funding from the BMBF's (German Federal Ministry of Education and Research) SLIK project under the grant 01IS22015C. Anne Lauscher's work is funded under the Excellence Strategy of the German Federal Government and States.

References

Mostafa Abdou, Vinit Ravishankar, Maria Barrett, Yonatan Belinkov, Desmond Elliott, and Anders Søgaard. 2020. The sensitivity of language models and humans to Winograd schema perturbations. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7590–7604, Online. Association for Computational Linguistics. https://doi.org/10.18653/v1/2020.acl-main.679

Sarah Alnegheimish, Alicia Guo, and Yi Sun. 2022. Using natural sentence prompts for understanding biases in language models. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2824–2830, Seattle, United States. Association for Computational Linguistics. https://doi.org/10.18653/v1/2022.naacl-main.203

Connor Baumler and Rachel Rudinger. 2022. Recognition of they/them as singular personal

pronouns in coreference resolution. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3426–3432, Seattle, United States. Association for Computational Linguistics. https://doi.org/10.18653/v1/2022.naacl-main.250

Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '21, pages 610–623, New York, NY, USA. Association for Computing Machinery. https://doi.org/10.1145/3442188.3445922

Emily M. Bender and Alexander Koller. 2020. Climbing towards NLU: On meaning, form, and understanding in the age of data. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5185–5198, Online. Association for Computational Linguistics. https://doi.org/10.18653/v1/2020.acl-main.463

Stella Biderman, Hailey Schoelkopf, Quentin Gregory 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. 2023. Pythia: A suite for analyzing large language models across training and scaling. In *International Conference on Machine Learning, ICML 2023, 23–29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pages 2397–2430. PMLR.

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. https://doi.org/10.18653/v1/2020.acl-main.485

Tolga Bolukbasi, Kai-Wei Chang, James Y. Zou, Venkatesh Saligrama, and Adam T. Kalai. 2016. Man is to computer programmer as woman is to homemaker? Debiasing word embeddings. In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc.

Aylin Caliskan, Joanna J. Bryson, and Arvind Narayanan. 2017. Semantics derived automatically from language corpora contain human-biases. *Science*, 356(6334):183–186. https://doi.org/10.1126/science.aal4230, PubMed: 28408601

Yang Trista Cao and Hal Daumé III. 2020. Toward gender-inclusive coreference resolution. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4568–4595, Online. Association for Computational Linguistics. https://doi.org/10.18653/v1/2020.acl-main.418

Yang Trista Cao and Hal Daumé III. 2021. Toward gender-inclusive coreference resolution: An analysis of gender and bias throughout the machine learning lifecycle. *Computational Linguistics*, 47(3):615–661. https://doi.org/10.1162/coli_a_00413

Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Alex Castro-Ros, Marie Pellat, Kevin Robinson, Dasha Valter, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Zhao, Yanping Huang, Andrew Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 2024. Scaling instruction-finetuned language models. *Journal of Machine Learning Research*, 25(70):1–53.

Kirby Conrod. 2019. *Pronouns Raising and Emerging*. PhD thesis, University of Washington.

Sunipa Dev, Masoud Monajatipoor, Anaelia Ovalle, Arjun Subramonian, Jeff Phillips, and Kai-Wei Chang. 2021. Harms of gender exclusivity and challenges in non-binary representation in language technologies. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1968–1994, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics. https://doi.org/10.18653/v1/2021.emnlp-main.150

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. 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)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics. https://doi.org/10.18653/v1/N19-1423

Yanai Elazar, Hongming Zhang, Yoav Goldberg, and Dan Roth. 2021. Back to square one: Artifact detection, training and commonsense disentanglement in the Winograd schema. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10486–10500, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics. https://doi.org/10.18653/v1/2021.emnlp-main.819

Denis Emelin and Rico Sennrich. 2021. Wino-X: Multilingual Winograd schemas for commonsense reasoning and coreference resolution. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 8517–8532, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics. https://doi.org/10.18653/v1/2021.emnlp-main.670

Angela Fan, Mike Lewis, and Yann Dauphin. 2019. Strategies for structuring story generation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2650–2660, Florence, Italy. Association for Computational Linguistics. https://doi.org/10.18653/v1/P19-1254

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. https://doi.org/10.18653/v1/2023.acl-long.507

Patrick Fernandes, Kayo Yin, Emmy Liu, André Martins, and Graham Neubig. 2023. When does translation require context? A data-driven, multilingual exploration. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 606–626, Toronto, Canada. Association for Computational Linguistics. https://doi.org/10.18653/v1/2023.acl-long.36

Seraphina Goldfarb-Tarrant, Rebecca Marchant, Ricardo Muñoz Sánchez, Mugdha Pandya, and Adam Lopez. 2021. Intrinsic bias metrics do not correlate with application bias. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1926–1940, Online. Association for Computational Linguistics. https://doi.org/10.18653/v1/2021.acl-long.150

Hila Gonen and Yoav Goldberg. 2019. Lipstick on a pig: Debiasing methods cover up systematic gender biases in word embeddings but do not remove them. 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)*, pages 609–614, Minneapolis, Minnesota. Association for Computational Linguistics.

Barbara J. Grosz, Aravind K. Joshi, and Scott Weinstein. 1995. Centering: A framework for modeling the local coherence of discourse. *Computational Linguistics*, 21(2):203–225. https://doi.org/10.21236/ADA324949

Reto Gubelmann and Siegfried Handschuh. 2022. Context matters: A pragmatic study of PLMs' negation understanding. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4602–4621, Dublin, Ireland. Association for Computational Linguistics. https://doi.org/10.18653/v1/2022.acl-long.315

Tamanna Hossain, Sunipa Dev, and Sameer Singh. 2023. MISGENDERED: Limits of large language models in understanding pronouns. In *Proceedings of the 61st Annual Meeting of*

the Association for Computational Linguistics (Volume 1: Long Papers), pages 5352–5367, Toronto, Canada. Association for Computational Linguistics. https://doi.org/10.18653/v1/2023.acl-long.293

Jennifer Hu and Michael Frank. 2024. Auxiliary task demands mask the capabilities of smaller language models. In *First Conference on Language Modeling*.

Jennifer Hu and Roger Levy. 2023. Prompting is not a substitute for probability measurements in large language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5040–5060, Singapore. Association for Computational Linguistics. https://doi.org/10.18653/v1/2023.emnlp-main.306

Alon Jacovi, Avi Caciularu, Omer Goldman, and Yoav Goldberg. 2023. Stop uploading test data in plain text: Practical strategies for mitigating data contamination by evaluation benchmarks. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5075–5084, Singapore. Association for Computational Linguistics. https://doi.org/10.18653/v1/2023.emnlp-main.308

Yangfeng Ji, Chenhao Tan, Sebastian Martschat, Yejin Choi, and Noah A. Smith. 2017. Dynamic entity representations in neural language models. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1830–1839, Copenhagen, Denmark. Association for Computational Linguistics. https://doi.org/10.18653/v1/D17-1195

Carina Kauf, Emmanuele Chersoni, Alessandro Lenci, Evelina Fedorenko, and Anna A. Ivanova. 2024. Comparing plausibility estimates in base and instruction-tuned large language models. *CoRR*, abs/2403.14859v1.

Carina Kauf and Anna Ivanova. 2023. A better way to do masked language model scoring. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 925–935, Toronto, Canada. Association for Computational Linguistics. https://doi.org/10.18653/v1/2023.acl-short.80

Takeshi Kojima, Shixiang (Shane) Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. In *Advances in Neural Information Processing Systems*, volume 35, pages 22199–22213. Curran Associates, Inc.

Keita Kurita, Nidhi Vyas, Ayush Pareek, Alan W. Black, and Yulia Tsvetkov. 2019. Measuring bias in contextualized word representations. In *Proceedings of the First Workshop on Gender Bias in Natural Language Processing*, pages 166–172, Florence, Italy. Association for Computational Linguistics. https://doi.org/10.18653/v1/W19-3823

Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. ALBERT: A lite BERT for self-supervised learning of language representations. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26–30, 2020*. OpenReview.net.

Anne Lauscher, Archie Crowley, and Dirk Hovy. 2022. Welcome to the modern world of pronouns: Identity-inclusive natural language processing beyond gender. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 1221–1232, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.

Anne Lauscher, Debora Nozza, Ehm Miltersen, Archie Crowley, and Dirk Hovy. 2023. What about ‘‘em’’? How commercial machine translation fails to handle (neo-)pronouns. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 377–392, Toronto, Canada. Association for Computational Linguistics. https://doi.org/10.18653/v1/2023.acl-long.23

Hector J. Levesque, Ernest Davis, and Leora Morgenstern. 2012. The winograd schema challenge. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Thirteenth International Conference, KR 2012, Rome, Italy, June 10–14, 2012*. AAAI Press.

Mosh Levy, Alon Jacoby, and Yoav Goldberg. 2024. Same task, more tokens: The impact of input length on the reasoning performance of large language models. In *Proceedings of*

the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 15339–15353, Bangkok, Thailand. Association for Computational Linguistics. https://doi.org/10.18653/v1/2024.acl-long.818

Shahar Levy, Koren Lazar, and Gabriel Stanovsky. 2021. Collecting a large-scale gender bias dataset for coreference resolution and machine translation. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2470–2480, Punta Cana, Dominican Republic. Association for Computational Linguistics. https://doi.org/10.18653/v1/2021.findings-emnlp.211

Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024. Lost in the middle: How language models use long contexts. *Transactions of the Association for Computational Linguistics*, 12:157–173. https://doi.org/10.1162/tacl_a_00638

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A robustly optimized BERT pretraining approach. *CoRR*, abs/1907.11692v1.

Cassian Lodge. 2023. *Gender Census 2023: Worldwide Report*. Gender Census.

Shayne Longpre, Le Hou, Tu Vu, Albert Webson, Hyung Won Chung, Yi Tay, Denny Zhou, Quoc V. Le, Barret Zoph, Jason Wei, and Adam Roberts. 2023. The flan collection: Designing data and methods for effective instruction tuning. In *International Conference on Machine Learning, ICML 2023, 23–29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pages 22631–22648. PMLR.

Chandler May, Alex Wang, Shikha Bordia, Samuel R. Bowman, and Rachel Rudinger. 2019. On measuring social biases in sentence encoders. 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)*, pages 622–628, Minneapolis, Minnesota. Association for Computational Linguistics. https://doi.org/10.18653/v1/N19-1063

Kevin A. McLemore. 2018. A minority stress perspective on transgender individuals' experiences with misgendering. *Stigma and Health*, 3(1):53–64. https://doi.org/10.1037/sah0000070

Kanishka Misra. 2022. minicons: Enabling flexible behavioral and representational analyses of transformer language models. *CoRR*, abs/2203.13112v1.

Nasrin Mostafazadeh, Nathanael Chambers, Xiaodong He, Devi Parikh, Dhruv Batra, Lucy Vanderwende, Pushmeet Kohli, and James Allen. 2016. A corpus and cloze evaluation for deeper understanding of commonsense stories. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 839–849, San Diego, California. Association for Computational Linguistics. https://doi.org/10.18653/v1/N16-1098

Mathias Müller, Annette Rios, Elena Voita, and Rico Sennrich. 2018. A large-scale test set for the evaluation of context-aware pronoun translation in neural machine translation. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 61–72, Brussels, Belgium. Association for Computational Linguistics. https://doi.org/10.18653/v1/W18-6307

Anaelia Ovalle, Palash Goyal, Jwala Dhamala, Zachary Jaggers, Kai-Wei Chang, Aram Galstyan, Richard Zemel, and Rahul Gupta. 2023. ''I'm fully who I am'': Towards centering transgender and non-binary voices to measure biases in open language generation. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '23, pages 1246–1266, New York, NY, USA. Association for Computing Machinery. https://doi.org/10.1145/3593013.3594078

Anaelia Ovalle, Ninareh Mehrabi, Palash Goyal, Jwala Dhamala, Kai-Wei Chang, Richard Zemel, Aram Galstyan, Yuval Pinter, and Rahul Gupta. 2024. Tokenization matters: Navigating data-scarce tokenization for gender inclusive

language technologies. In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 1739–1756, Mexico City, Mexico. Association for Computational Linguistics. https://doi.org/10.18653/v1/2024.findings-naacl.113

Jacob Portes, Alexander Trott, Sam Havens, Daniel King, Abhinav Venigalla, Moin Nadeem, Nikhil Sardana, Daya Khudia, and Jonathan Frankle. 2023. MosaicBERT: A bidirectional encoder optimized for fast pretraining. In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10–16, 2023*.

Sameer Pradhan, Alessandro Moschitti, Nianwen Xue, Hwee Tou Ng, Anders Björkelund, Olga Uryupina, Yuchen Zhang, and Zhi Zhong. 2013. Towards robust linguistic analysis using OntoNotes. In *Proceedings of the Seventeenth Conference on Computational Natural Language Learning*, pages 143–152, Sofia, Bulgaria. Association for Computational Linguistics.

Rachel Rudinger, Jason Naradowsky, Brian Leonard, and Benjamin Van Durme. 2018. 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)*, pages 8–14, New Orleans, Louisiana. Association for Computational Linguistics. https://doi.org/10.18653/v1/N18-2002

Julian Salazar, Davis Liang, Toan Q. Nguyen, and Katrin Kirchhoff. 2020. Masked language model scoring. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2699–2712, Online. Association for Computational Linguistics. https://doi.org/10.18653/v1/2020.acl-main.240

Rylan Schaeffer, Brando Miranda, and Sanmi Koyejo. 2023. Are emergent abilities of large language models a mirage? In *Thirty-seventh Conference on Neural Information Processing Systems*.

Melanie Sclar, Yejin Choi, Yulia Tsvetkov, and Alane Suhr. 2024. Quantifying language models' sensitivity to spurious features in prompt design or: How I learned to start worrying about prompt formatting. In *The Twelfth International Conference on Learning Representations*.

Preethi Seshadri, Pouya Pezeshkpour, and Sameer Singh. 2022. Quantifying social biases using templates is unreliable. *CoRR*, abs/2210.04337v1.

Shanya Sharma, Manan Dey, and Koustuv Sinha. 2022. How sensitive are translation systems to extra contexts? Mitigating gender bias in neural machine translation models through relevant contexts. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 1968–1984, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics. https://doi.org/10.18653/v1/2022.findings-emnlp.143

Michael Silverstein. 1985. 10 - Language and the culture of gender: At the intersection of structure, usage, and ideology. In Elizabeth Mertz and Richard J. Parmentier, editors, *Semiotic Mediation*, pages 219–259. Academic Press, San Diego. https://doi.org/10.1016/B978-0-12-491280-9.50016-9

Koustuv Sinha, Jon Gauthier, Aaron Mueller, Kanishka Misra, Keren Fuentes, Roger Levy, and Adina Williams. 2023. Language model acceptability judgements are not always robust to context. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6043–6063, Toronto, Canada. Association for Computational Linguistics. https://doi.org/10.18653/v1/2023.acl-long.333

Neha Srikanth and Rachel Rudinger. 2022. Partial-input baselines show that NLI models can ignore context, but they don't. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4753–4763, Seattle, United States. Association for Computational Linguistics. https://doi.org/10.18653/v1/2022.naacl-main.350

Susan Stryker. 2017. *Transgender History: The Roots of Today's Revolution*, 2nd edition. Seal Press.

Yarden Tal, Inbal Magar, and Roy Schwartz. 2022. Fewer errors, but more stereotypes? The effect of model size on gender bias. In *Proceedings of the 4th Workshop on Gender Bias in Natural Language Processing (GeBNLP)*, pages 112–120, Seattle, Washington. Association for Computational Linguistics. https://doi.org/10.18653/v1/2022.gebnlp-1.13

Yi Tay, Mostafa Dehghani, Samira Abnar, Hyung Chung, William Fedus, Jinfeng Rao, Sharan Narang, Vinh Tran, Dani Yogatama, and Donald Metzler. 2023. Scaling laws vs model architectures: How does inductive bias influence scaling? In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 12342–12364, Singapore. Association for Computational Linguistics. https://doi.org/10.18653/v1/2023.findings-emnlp.825

Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton-Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurélien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. Llama 2: Open foundation and fine-tuned chat models. *CoRR*, abs/2307.09288v2.

Paul Trichelair, Ali Emami, Adam Trischler, Kaheer Suleman, and Jackie Chi Kit Cheung. 2019. How reasonable are common-sense reasoning tasks: A case-study on the Winograd schema challenge and SWAG. 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 3382–3387, Hong Kong, China. Association for Computational Linguistics. https://doi.org/10.18653/v1/D19-1335

Daniel de Vassimon Manela, David Errington, Thomas Fisher, Boris van Breugel, and Pasquale Minervini. 2021. Stereotype and skew: Quantifying gender bias in pre-trained and fine-tuned language models. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2232–2242, Online. Association for Computational Linguistics. https://doi.org/10.18653/v1/2021.eacl-main.190

Elena Voita, Pavel Serdyukov, Rico Sennrich, and Ivan Titov. 2018. Context-aware neural machine translation learns anaphora resolution. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1264–1274, Melbourne, Australia. Association for Computational Linguistics. https://doi.org/10.18653/v1/P18-1117

Kellie Webster, Marta Recasens, Vera Axelrod, and Jason Baldridge. 2018. Mind the GAP: A balanced corpus of gendered ambiguous pronouns. *Transactions of the Association for Computational Linguistics*, 6:605–617. https://doi.org/10.1162/tacl_a_00240

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics. https://doi.org/10.18653/v1/2020.emnlp-demos.6

Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen,

Christopher Dewan, Mona T. Diab, Xian Li, Xi Victoria Lin, Todor Mihaylov, Myle Ott, Sam Shleifer, Kurt Shuster, Daniel Simig, Punit Singh Koura, Anjali Sridhar, Tianlu Wang, and Luke Zettlemoyer. 2022. OPT: Open pre-trained transformer language models. *CoRR*, abs/2205.01068v4.

Jieyu Zhao, Tianlu Wang, Mark Yatskar, Ryan Cotterell, Vicente Ordonez, and Kai-Wei Chang. 2019. 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)*, pages 629–634, Minneapolis, Minnesota. Association for Computational Linguistics. https://doi.org/10.18653/v1/N19-1064

Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. 2018. 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)*, pages 15–20, New Orleans, Louisiana. Association for Computational Linguistics. https://doi.org/10.18653/v1/N18-2003

Yongchao Zhou, Andrei Ioan Muresanu, Ziwen Han, Keiran Paster, Silviu Pitis, Harris Chan, and Jimmy Ba. 2023. Large language models are human-level prompt engineers. In *The Eleventh International Conference on Learning Representations*.

A List of Occupations

The occupations along with their respective participants in parentheses are listed below in alphabetical order. This list is identical to the occupations and participants in Rudinger et al. (2018), except that we pair examiner with intern rather than victim:

accountant (taxpayer), administrator (undergraduate), advisor (advisee), appraiser (buyer), architect (student), auditor (taxpayer), baker (customer), bartender (customer), broker (client), carpenter (onlooker), cashier (customer), chef (guest), chemist (visitor), clerk (customer), counselor (patient), dietitian (client), dispatcher (bystander), doctor (patient), educator (student), electrician (homeowner), engineer (client), examiner (intern), firefighter (child), hairdresser (client), hygienist (patient), inspector (homeowner), instructor (student), investigator (witness), janitor (child), lawyer (witness), librarian (child), machinist (child), manager (customer), mechanic (customer) nurse (patient), nutritionist (patient), officer (protester), painter (customer), paralegal (client), paramedic (passenger), pathologist (victim), pharmacist (patient), physician (patient), planner (resident), plumber (homeowner), practitioner (patient), programmer (student), psychologist (patient), receptionist (visitor), salesperson (customer), scientist (undergraduate), secretary (visitor), specialist (patient), supervisor (employee), surgeon (child), teacher (student), technician (customer), therapist (teenager), veterinarian (owner), worker (pedestrian)

B Context Template Construction

For each grammatical case, we create 10 explicit templates, which explicitly demonstrate the coreference between an individual and a pronoun using a subordinate clause, and 10 implicit templates, simple sentences which only contain a pronoun as the subject. An introduction and the first distractor are always sampled from the explicit templates, and the rest are sampled from the implicit templates, as this reflects natural and coherent use of pronouns in discourse.

For both the explicit and implicit cases, we create five templates with terms with positive connotations (e.g., full, happy) and five templates with the opposite polarity (i.e., hungry, unhappy). We denote exp_pos_i as the i-th positive explicit template where i ranges from 1 to 5; exp_neg_i is the corresponding negative version. The introduction template can be selected from any of these 10 possibilities and filled with one of four pronouns.

After this, we pick a first distractor template, limiting ourselves to the five templates of the opposite sentiment of what we first picked, and also excluding the template of the same index and opposite polarity. For example, if we chose exp_pos_3 as our introductory template, we would choose our first distractor template from $\{exp\_neg_1, exp\_neg_2, exp\_neg_4, exp\_neg_5\}$.

After making a choice for the first distractor template, we fill it with any of the three remaining pronouns and then we remove this template's index from our pool, but re-add the index of the introductory template. This is because subsequent distractor templates always use implicit templates. For example, if we chose exp_neg_4 as our first distractor template, we would now choose from $\{imp\_neg_1, imp\_neg_2, imp\_neg_3, imp\_neg_5\}$. For subsequent distractor templates, we sample without replacement from these implicit templates.

C Annotator Demographics

All three annotators (two authors and an additional annotator) are fluent English speakers. The two authors who create and validate templates have linguistic training at the undergraduate level. One author and one annotator have experience with using singular *they* and neopronouns, while the other author has prior exposure to singular *they* but not the neopronoun *xe*.

D Annotation Instructions

D.1 Task 1 Description

Together with this annotation protocol, you have received a link to a Google Sheet. The sheet contains 2 data columns and 2 task columns of randomized data. The data columns consist of

- Sentences which you are asked to annotate for grammaticality; and
- Questions about pronouns in the sentence, which you are asked to answer

Please be precise in your assignments and do not reorder the data. The columns have built-in data validation and we will perform further tests to check for consistent annotation.

D.1.1 Grammaticality

In the ‘‘Grammatical?’’ column, please enter your grammaticality judgments of the sentence, according to Standard English. The annotation options are:

- **grammatical** (for fluent, syntactically valid and semantically plausible sentences)
- **ungrammatical** (for sentences that have any typos, grammatical issues, or if the sentence describes a situation that don't make sense, or just sounds weird)
- **not sure** (if you are not sure whether it is clearly grammatical or ungrammatical)

Examples:

- *The driver told the passenger that he could pay for the ride with cash.*
  => grammatical
- *The driver said the passenger that he could pay for the ride with cash.*
  => ungrammatical (because 'said' is intransitive in Standard English)

D.1.2 Questions about Pronouns

Every sentence contains a pronoun, and the ‘‘Question’’ column asks whether it refers to a person mentioned in the sentence or not. The annotation options are:

- **yes** (if the pronoun refers to the person)

- **no** (if the pronoun does not refer to the person)
- **not sure** (if you are not sure about whether the pronoun refers to the person)

Examples:

- *The driver told the passenger that he could pay for the ride with cash.*
  Does the pronoun he refer to the driver?
  => no
- *The driver told the passenger that he could pay for the ride with cash.*
  Does the pronoun he refer to the passenger?
  => yes

D.2 Task 2 Description

Together with this annotation protocol, you have received a link to a Google Sheet. The sheet contains 1 randomized data column and 1 task column.
Each row in the data column consists of multiple sentences, of which precisely one sentence contains a blank. Your task is to determine the appropriate pronoun to fill in the blank, and enter it in the ''Pronoun'' column. Here, appropriate means correct in both form and case.
The tasks are designed to be unambiguous, so please provide only one solution and do not reorder the data.
Example:

- *The driver felt unhappy because he did not make enough money. The driver wondered whether ___ should take out a loan.*
  => he

E Experimental Details

We use one 40GB NVIDIA A100 GPU for inference with most models, but we require two GPUs for OPT-30B and four for OPT-66B and the LLAMA-2-70B base and chat models. We access all models using the HuggingFace Transformers library (Wolf et al., 2020) and use the minicons library for pseudo log likelihood evaluation (Misra, 2022).

F Prompting

Table 3 shows all 10 prompt templates we use to present our task instances to chat models. The first five are open-ended and the last five provide the model with pronoun options to choose from. These templates are used within vanilla and chain-of-thought prompting paradigms as described below.

F.1 Vanilla Prompting

With FLAN-T5, vanilla prompting only requires instantiating templates with task instances, whereas LLAMA-2-CHAT requires special formatting with `INST`. Instantiated examples of one template are shown for both models in Table 4. The number of maximum new tokens is set to 5 for FLAN-T5 and 20 for LLAMA-2-CHAT based on experimentation.

F.2 Chain-of-thought Prompting

We focus on FLAN-T5-XXL for chain-of-thought experiments, and use the strong zero-shot prompts for reasoning proposed by Kojima et al. (2022) (*''Let's think step by step''*) and Zhou et al. (2023) (*''Let's work this out in a step by step way to be sure we have the right answer''*), which we append after the

ID	Template
0	{task}\n\n What pronoun should be used to fill the blank?
1	{task}\n\n The best pronoun to fill in the blank is
2	Fill in the blank with the correct pronoun.\n\n{task}
3	Fill in the blank with the appropriate pronoun.\n\n{task}
4	What pronoun should be used to fill the blank?\n\n{task}
5	{task}\n\n What pronoun should be used to fill the blank?\n{options}
6	{task}\n\n The best pronoun to fill in the blank is\n{options}
7	Fill in the blank with the correct pronoun.\n\n{task}\n{options}
8	Fill in the blank with the appropriate pronoun.\n\n{task}\n{options}
9	What pronoun should be used to fill the blank?\n\n{task}\n{options}

Table 3: Prompting templates for chat models, where ‘‘task’’ is filled with each dataset instance, and ‘‘options’’ is a list of four pronouns to choose from, all in the correct case.

Model type	Example input
FLAN-T5	Fill in the blank with the appropriate pronoun. The accountant had just eaten a big meal so her stomach was full. The accountant was asked about _____ charges for preparing tax returns. OPTIONS: - his - her - their - xyr
LLAMA-2-CHAT	[INST] Fill in the blank with the appropriate pronoun. The accountant had just eaten a big meal so her stomach was full. The accountant was asked about _____ charges for preparing tax returns. OPTIONS: - his - her - their - xyr [/INST]

Table 4: Example input using template 8 for FLAN-T5 and LLAMA-2-CHAT models.

template. Following their codebases, we first allow the models to generate a chain of thought (with 128 maximum new tokens). Then, we append the chain of thought after the question and elicit the final answer with the string *‘‘Therefore, the correct pronoun is,’’* allowing the model to generate up to 10 new tokens. This two-step process is illustrated with examples in Table 5. We save both the final answer and the chain of thought for later analysis.

Chain-of-thought prompt	Example input
Kojima et al. (2022) (eliciting chain of thought)	Q: Fill in the blank with the appropriate pronoun. The accountant had just eaten a big meal so her stomach was full. The accountant was asked about ____ charges for preparing tax returns. A: Let's think step by step.
Kojima et al. (2022) (eliciting final answer)	Q: Fill in the blank with the appropriate pronoun. The accountant had just eaten a big meal so her stomach was full. The accountant was asked about ____ charges for preparing tax returns. A: Let's think step by step. {`generated chain of thought`} Therefore, the correct pronoun is
Zhou et al. (2023) (eliciting chain of thought)	Q: Fill in the blank with the appropriate pronoun. The accountant had just eaten a big meal so her stomach was full. The accountant was asked about ____ charges for preparing tax returns. A: Let's work this out in a step by step way to be sure we have the right answer.
Zhou et al. (2023) (eliciting final answer)	Q: Fill in the blank with the appropriate pronoun. The accountant had just eaten a big meal so her stomach was full. The accountant was asked about ____ charges for preparing tax returns. A: Let's work this out in a step by step way to be sure we have the right answer. {`generated chain of thought`} Therefore, the correct pronoun is

Table 5: Example input using template 3 for evaluating FLAN-T5-XXL with two types of chain-of-thought prompting. Prompting happens in two phases regardless of the choice of prompt: eliciting the chain of thought and eliciting the final answer.

G Context-free Pronoun Predictions by Model

Figure 10 shows per-model pronoun predictions in the absence of context. All models predict *he/him/his* more frequently than *she/her/her*, which is in turn predicted more frequently than *they/them/their* and *xe/xem/xyr*. However, encoder-only models are more balanced in their predictions across the four pronoun sets, compared to decoder-only models which show very stark differences in pronoun predictions.

H Results with Vanilla and Chain-of-Thought Prompting

H.1 Vanilla Prompting

Prompting is a different model evaluation mechanism than log likelihoods, with higher task demands that lead to lower performance than log likelihoods with both base models and instruction fine-tuned chat models (Hu and Levy, 2023; Hu and Frank, 2024; Kauf et al., 2024). We thus expect vanilla prompting results (using the prompts listed in Appendix F) to be worse than results with log likelihoods. Indeed, Figure 11a shows that LLAMA-2-CHAT prompting performance is lower than LLAMA-2 evaluated

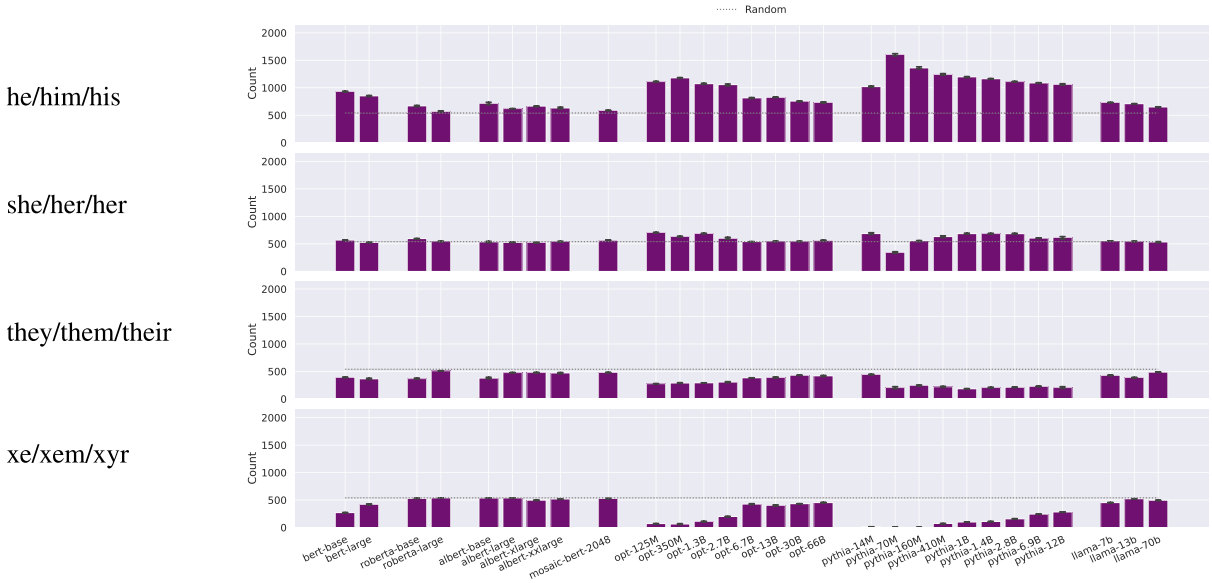

Figure 10: Counts of pronoun predictions from all models, in the absence of context. The random baseline shows counts if each pronoun set was chosen equally often.

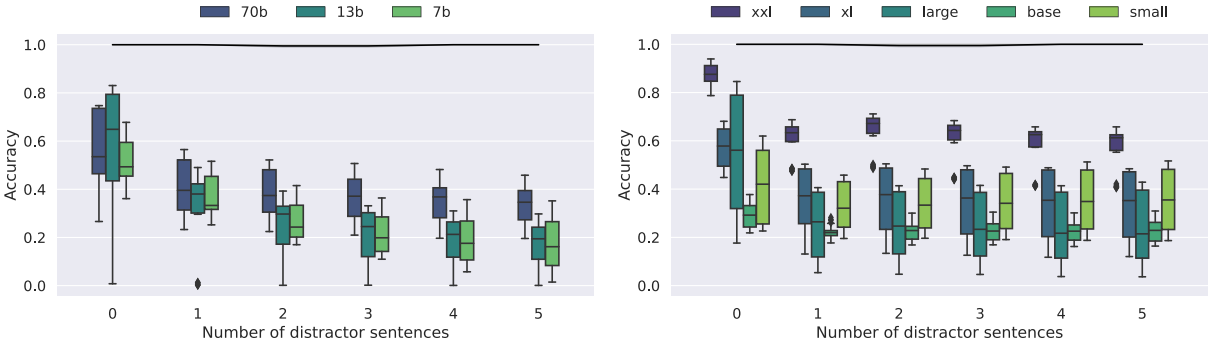

(a) LLAMA-2-CHAT (b) FLAN-T5

Figure 11: Performance of chat models (LLAMA-2-CHAT and FLAN-T5) with additional distractors, using vanilla prompting. The boxplots show the range of performance across 10 different templates.

with log likelihoods, even with no distractors. Figure 11b shows the results of standard prompting with FLAN-T5, an encoder-decoder model which shows similar patterns of degradation to decoder-only models. Bigger models are mostly better and degrade more gracefully than the smaller ones, but there remains a lot of variance across prompts, as shown in the box plots.

H.2 Chain-of-thought Prompting

As FLAN-T5-XXL shows strong performance with low variance compared to all the other chat models we consider, we focus on this model for additional evaluation with chain-of-thought prompting. Zero-shot chain-of-thought prompting encourages models to think step-by-step, which could in theory produce much better results on pronoun fidelity. While chain-of-thought prompting is excessive for a task as simple as pronoun fidelity, it might encourage the model to explicitly list the referents and associated pronouns, which could help the model predict the correct pronoun with higher accuracy. In practice, however, we find that it leads to worse performance, potentially due to hallucination.

Figure 12 shows the pronoun fidelity of FLAN-T5-XXL with different types of prompting based on the final answer the model provides. Both types of chain-of-thought prompting worsen performance and increase the variance across prompts compared to vanilla prompting.

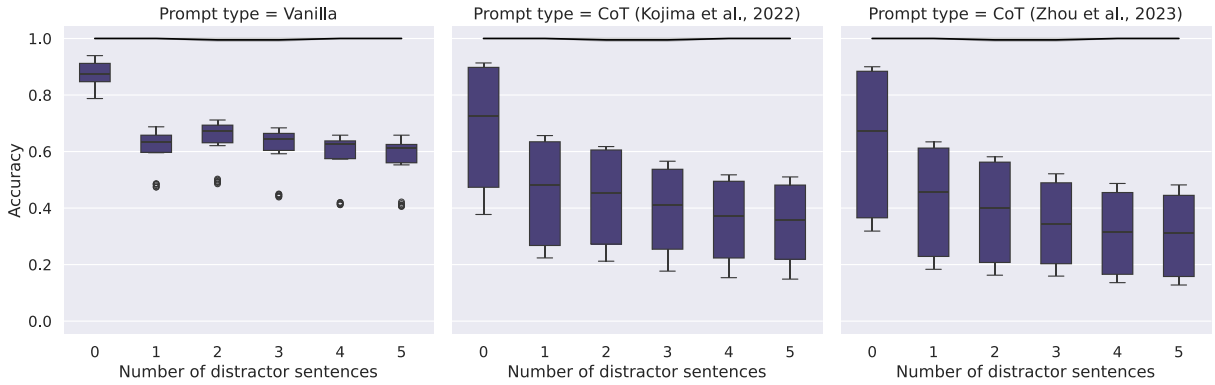

Figure 12: Performance of FLAN-T5-XXL with distractor sentences, comparing vanilla prompting to two types of chain-of-thought prompting. Here, the model's *final answers* are used for evaluation and the boxplots show the range of performance across 10 different templates.

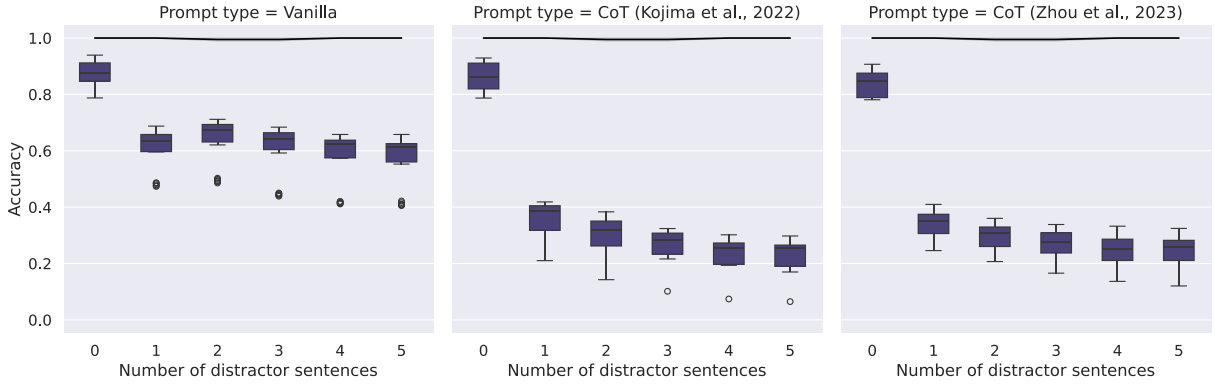

Figure 13: Performance of FLAN-T5-XXL with distractor sentences, comparing vanilla prompting to two types of chain-of-thought prompting. Here, the model's *chain of thought* is used for evaluation and the boxplots show the range of performance across 10 different templates.

When examining model-generated answers and chains of thought, we found that FLAN-T5-XXL does not in fact solve the problem step by step as the instruction suggests. Instead, the chain of thought often already contains an answer, and the final answer is not necessarily the same as this one. Therefore, we also plot performance using answers from the model-generated chain of thought in Figure 13. Once again, performance with 1-5 distractors is much lower, showing that chain-of-thought prompting degrades performance compared to vanilla prompting. However, with no distractors, performance is almost exactly the same as vanilla prompting, as models simply generate the answer within the chain of thought. This reinforces that chain-of-thought is unnecessary for a task this simple.