"You are grounded!": Latent Name Artifacts in Pre-trained Language Models

Vered Shwartz[1,2], **Rachel Rudinger**[1,2,3], **and Oyvind Tafjord**[1]
[1]Allen Institute for Artificial Intelligence
[2]Paul G. Allen School of Computer Science & Engineering, University of Washington
[3]University of Maryland, College Park, MD
{vereds,oyvindt}@allenai.org, rudinger@umd.edu

Abstract

Pre-trained language models (LMs) may perpetuate biases originating in their training corpus to downstream models. We focus on artifacts associated with the representation of given names (*e.g.*, Donald), which, depending on the corpus, may be associated with specific entities, as indicated by next token prediction (*e.g.*, Trump). While helpful in some contexts, grounding happens also in underspecified or inappropriate contexts. For example, endings generated for 'Donald is a' substantially differ from those of other names, and often have more-than-average negative sentiment. We demonstrate the potential effect on downstream tasks with reading comprehension probes where name perturbation changes the model answers. As a silver lining, our experiments suggest that additional pre-training on different corpora may mitigate this bias.

1 Introduction

Pre-trained language models (LMs) have transformed the NLP landscape. State-of-the-art performance across tasks is achieved by fine-tuning the latest LM on task-specific data. LMs provide an effective way to represent contextual information, including lexical and syntactic knowledge as well as world knowledge (Petroni et al., 2019).

LMs conflate generic facts (*e.g.*"the US has a president") with grounded knowledge regarding specific entities and events (*e.g.*"the (current) president is a male"), occasionally leading to gender and racial biases (*e.g.*"women can't be presidents") (May et al., 2019; Sheng et al., 2019).

In this work we focus on the representations of given names in pre-trained LMs (Table 1). Prior work showed that the representations of named entities incorporate sentiment (Prabhakaran et al., 2019), which is often transferable across entities via a shared given name (Field and Tsvetkov, 2019).

Model	Main Corpus Type	Gen.	Cls.
BERT (Devlin et al., 2019)	Wikipedia	×	∨
RoBERTa (Liu et al., 2019)	Web	×	∨
GPT (Radford et al., 2018)	Fiction	∨	×
GPT2 (Radford et al., 2019)	Web	∨	×
XLNet (Yang et al., 2019)	Web	∨	∨
TransformerXL (Dai et al., 2019)	Wikipedia	∨	×

Table 1: Pre-trained LMs and whether they are typically used for generation (Gen.) or classification (Cls.).

In a series of experiments we show that, depending on the corpus, some names tend to be grounded to specific entities, *even in generic contexts*.

The most striking effect is of politicians in GPT2. For example, the name *Donald*: 1) predicts *Trump* as the next token with high probability; 2) generated endings of "*Donald is a*" are easily distinguishable from any other given name; 3) their sentiment is substantially more negative; and 4) this bias can potentially perpetuate to downstream tasks.

Although these results are expected, their extent is surprising. Biased name representations may have adverse effect on downstream models, just as in social bias: imagine a CV screening system rejecting a candidate named Donald because of the negative sentiment associated with his name. Our experiments may be used to evaluate the extent of name artifacts in future LMs.[1]

2 Last Name Prediction

As an initial demonstration of the tendency of pre-trained LMs to ground given names to prominent named entities in the media, we examine the next-word probabilities assigned by the LM. If high probability is placed on a named entity's last name conditioned on observing their given name (*e.g.*, $P(\text{Trump}|\text{Donald}) = 0.99$), we take this as evidence that the LM is, in effect, interpreting the first-name mention as a reference to the named entity. We note that this is a lower bound on evidence

[1]Data and code available at: github.com/vered1986/LM_NE_bias

Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, pages 6850–6861,

Model	Named Entities from News					Named Entities from History				
	Minimal	News	History	Infrml	Avg	Minimal	News	History	Infrml	Avg
GPT	0.0	7.0	12.7	1.4	5.3	0.0	21.9	39.1	7.8	17.2
GPT2-small	22.5	63.4	50.7	15.5	38.0	12.5	29.7	56.2	12.5	27.7
GPT2-medium	33.8	64.8	49.3	12.7	40.2	21.9	32.8	62.5	4.7	30.5
GPT2-large	43.7	66.2	47.9	16.9	43.7	29.7	29.7	56.2	12.5	32.0
GPT2-XL	50.7	62.0	45.1	21.1	44.7	28.1	31.2	60.9	14.1	33.6
TransformerXL	14.1	18.3	15.5	12.7	15.2	35.9	43.8	51.6	37.5	42.2
XLNet-base	4.2	33.8	12.7	4.2	13.7	0.0	34.4	23.4	3.1	15.2
XLNet-large	11.3	40.8	23.9	9.9	21.5	6.2	29.7	31.2	7.8	18.7
Average	22.5	44.5	32.2	11.8	27.7	16.8	31.7	47.6	12.5	27.1

Table 2: Percentage of named entities such that each LM greedily generates their last name conditioned on a prompt ending with their given name. Named entities are (1) frequently mentioned people in the U.S. news, or (2) prominent people from history.

Named Entity	Media Freq.	Rank	Minimal Prompt		News Prompt		History Prompt		Informal Prompt	
			Next Word	%	Next Word	%	Next Word	%	Next Word	%
Donald Trump	2,844,894	15	**Trump**	70.8	**Trump**	99.0	**Trump**	93.2	**Trump**	34.1
Hillary Clinton	373,952	788	**Clinton**	80.9	**Clinton**	91.6	**Clinton**	82.9	**Clinton**	46.5
Robert Mueller	322,466	3	B[. Reich]	2.1	**Mueller**	82.2	F[. Kennedy]	13.5	.	16.6
Bernie Sanders	97,104	757	**Sanders**	66.8	**Sanders**	95.9	**Sanders**	84.8	**Sanders**	24.9
Benjamin Netanyahu	65,863	66	**Netanyahu**	10.8	**Netanyahu**	78.9	Franklin	61.3	.	15.7
Elizabeth Warren	58,370	5	,	4.7	**Warren**	90.1	Taylor	17.1	.	21.4
Marco Rubio	56,224	363	**Rubio**	15.2	**Rubio**	98.1	Polo	68.4	.	2.3
Richard Nixon	55,911	7	B[. Spencer]	2.1	**Nixon**	17.3	**Nixon**	76.8	.	20.0

Table 3: Maximum next-word probabilities from GPT2-XL conditioned on prompts with first names of select people frequently mentioned in the media. Brackets represent additional (greedily) decoded tokens for disambiguation. **Rank**: aggregate 1990 U.S. Census data of most common male and female names.

for grounding: while it is reasonable to assume that nearly all mentions of, *e.g.*, “Hillary Clinton” in text *are* references to (the entity) Hillary Clinton, other references may use different strings (“Hillary Rodham Clinton,” “H.R.C.,” or just “Hillary”). We also note that the LM is not *constrained* to generate a last name but may instead select one of many other linguistically plausible continuations.

We examine greedy decoding of named entity last names systematically for each generative LM. To this end, we compile two sets of prominent named entities from the media and from history.[2] We construct four prompt templates ending with a given name to feed to each LM: (1) **Minimal**: “[NAME]”, (2) **News**: “A new report from CNN says that [NAME]”, (3) **History**: “A newly published biography of [NAME]”, and (4) **Informal**: “I want to introduce you to my best friend, [NAME]”. Table 2 shows, for each LM, the percentage of named entities for which the LM greedily generated that entity’s last name[3] conditioned on one of the four prompt templates.

[2]Media: public.tableau.com/views/2018Top100/1_Top100. Name frequency source: 1990 U.S. Census statistics. See Section A for full list of names.

[3]Or a middle initial followed by the last name.

Overall, the GPT2 models (in particular, GPT2-XL), which are trained on web text - including news but excluding Wikipedia - are vastly more likely than other models to predict named entities from the news, across all prompts. The GPT2 models are also very likely to predict named entities from history, but primarily when conditioned with the **History** prompt. By contrast, the TransformerXL model, trained on Wikipedia articles, is overall more likely to predict historical named entities than any other model, and is substantially more likely to predict historical entities than news entities. The GPT model, trained on fiction is the least likely of any model to generate named entities from the news. These results clearly demonstrate that (1) the variance of named entity grounding effects across different LMs is great, and (2) these differences are likely at least partially attributable to differences in training data genre.

Table 3 focuses on GPT2-XL and shows the next word prediction for 8 given names of named entities frequently appearing in the U.S. news media, which are also common in the general population. Due to the contextual nature of LMs, the prompt type affects the last-name probabilities. Intuitively, generating the last name of an entity seems appro-

GPT		GPT2-small		GPT2-medium		GPT2-large		GPT2-XL		TransformerXL		XLNet-base		XLNet-large	
Name	$\mathbf{F_1}$	**Name**	$\mathbf{F_1}$	**Name**	$\mathbf{F_1}$	**Name**	$\mathbf{F_1}$	**Name**	$\mathbf{F_1}$	**Name**	$\mathbf{F_1}$	**Name**	$\mathbf{F_1}$	**Name**	$\mathbf{F_1}$
Philip	0.739	**Bernie**	0.853	**Bernie**	0.884	**Bernie**	0.815	**Bernie**	0.966	Virginia	0.761	Grace	0.793	Brittany	0.808
Bryan	0.683	**Donald**	0.800	**Donald**	0.845	**Barack**	0.800	**Donald**	0.922	Dylan	0.742	Rose	0.705	Matthew	0.803
Beverly	0.670	Victoria	0.772	**Irma**	0.834	**Theresa**	0.773	**Hillary**	0.869	**Hillary**	0.731	Martha	0.702	Amber	0.788
Louis	0.641	Virginia	0.771	Christian	0.822	**Donald**	0.759	**Barack**	0.832	**Jeff**	0.715	Victoria	0.700	**Hillary**	0.782
Danielle	0.639	Gloria	0.763	**Hillary**	0.782	Victoria	0.702	Virginia	0.767	Alice	0.693	Alice	0.692	Teresa	0.771
Kelly	0.631	**Hillary**	0.756	**Barack**	0.774	Matthew	0.688	Christian	0.749	Thomas	0.690	**Hillary**	0.661	Grace	0.764
Nicholas	0.631	Cheryl	0.755	Victoria	0.766	Jacob	0.688	Jose	0.746	Judy	0.681	Mary	0.657	Virginia	0.762
Brenda	0.630	**Jeff**	0.733	Virginia	0.760	Billy	0.677	**Irma**	0.739	Gregory	0.677	Kenneth	0.656	Jordan	0.755
Vincent	0.628	Ann	0.697	Joyce	0.757	Virginia	0.676	Joseph	0.732	Samantha	0.676	Bobby	0.653	Madison	0.754
Russell	0.625	Christina	0.693	Alice	0.753	**Paul**	0.668	Sophia	0.717	Amber	0.675	Virginia	0.651	**Barack**	0.751
0.526 ± 0.157		0.568 ± 0.173		0.572 ± 0.182		0.545 ± 0.166		0.549 ± 0.181		0.552 ± 0.169		0.525 ± 0.162		0.548 ± 0.175	

Table 4: Top 10 most predictable names from the "is a" endings for each model, using Nucleus sampling with $p = 0.9$ and limiting the number of generated tokens to 150. Bold entries mark given names that appear frequently in the media. Bottom: mean and STD of scores.

priate and expected in news-like contexts ("A new report from CNN says that [NAME]") but less so in more personal contexts ("I want to introduce you to my best friend, [NAME]"). Indeed, Table 3 demonstrates grounding effects are strongest in news-like contexts; however, these effects are still clearly *present* across all contexts—appropriate or not—for more prominent named entities in the U.S. media (Donald, Hillary, and Bernie). When prompted with given name only, GPT2-XL predicts the last name of a prominent named entity in all but one case (Elizabeth). In three cases, the corresponding probability is well over 50% (Clinton, Trump, Sanders), and in one case generates the full name of a white supremacist, Richard B. Spencer.

3 Given Name Recovery

Given a text discussing a certain person, can we recover their (masked) given name? Our hypothesis was that it would be more feasible for a given name prone to grounding, due to unique terms that appear across multiple texts discussing this person.

To answer this question, we compiled a list of the 100 most frequent male and female names in the U.S.,[4] to which we added the first names of the most discussed people in the media (Section 2). Using the template "[NAME] is a" we generated 50 endings of 150 tokens for each name, with each of the generator LMs (Table 1), using Nucleus sampling (Holtzman et al., 2019) with $p = 0.9$. For each pair of same-gender given names,[5] we trained a binary SVM classifier using the Scikit-learn library (Pedregosa et al., 2011) to predict the given name from the TF-IDF representation of the endings, excluding the name. Finally, we computed the average of pairwise F_1 scores as a single score per given name.

[4] www.ssa.gov/oact/babynames/decades/century.html.

[5] To avoid confounding gender bias.

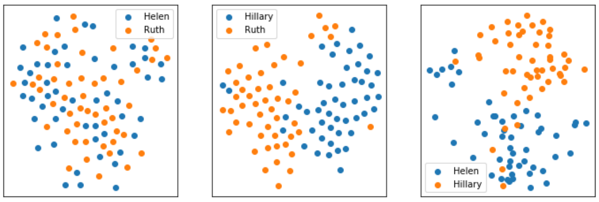

Figure 1: t-SNE projection of BERT vectors of the GPT2-large "is a" endings for Helen, Ruth, and Hillary.

Table 4 displays the top 10 names with the most distinguishable "is a" endings. Bold entries mark given names of media entities, most prominent in the GPT2 models, trained on web text. Apart from U.S. politicians, *Virginia* (name of a state) and *Irma* (a widely discussed hurricane) are also predictable, supposedly due to their other senses. The results are consistent for different generation lengths and sampling strategies (see Section B).

Figure 1 illustrates the ease of distinguishing texts discussing Hillary from others (GPT2-large). We masked the name ("[MASK] is a..."), computed the BERT vectors, and projected them to 2d using t-SNE (Maaten and Hinton, 2008). Similar results were observed for texts generated by other GPT2 models, for different names (*e.g.*, Donald, Bernie), and with other input representations (TF-IDF).

4 Sentiment Analysis

Following Prabhakaran et al. (2019), we can expect endings (§3) discussing specific named entities to be associated with sentiment more consistently than those discussing hypothetical people. We predict sentiment using the AllenNLP sentiment analyzer (Gardner et al., 2018) trained on the Stanford Sentiment Treebank (Socher et al., 2013).

Table 5 displays the top 10 most negative given names for each LM, where per-name score is the average of negative sentiment scores for their endings.

GPT		GPT2-small		GPT2-medium		GPT2-large		GPT2-XL		TransformerXL		XLNet-base		XLNet-large	
Name	**Score**	**Name**	**Score**	**Name**	**Score**	**Name**	**Score**	**Name**	**Score**	**Name**	**Score**	**Name**	**Score**	**Name**	**Score**
Noah	0.808	**Bernie**	0.619	**Donald**	0.629	**Bernie**	0.556	Alice	0.620	Sean	0.526	Judy	0.382	Kyle	0.324
John	0.802	**Donald**	0.591	**Bernie**	0.565	**Hillary**	0.537	**Donald**	0.546	**Mitch**	0.525	Albert	0.375	**Rudy**	0.318
Keith	0.800	Ryan	0.560	Jerry	0.559	Johnny	0.505	**Chuck**	0.526	Jack	0.512	Johnny	0.370	Johnny	0.318
Kenneth	0.795	**Hillary**	0.547	**Kevin**	0.546	Alice	0.490	Ryan	0.524	Johnny	0.507	**Hillary**	0.357	Sean	0.304
Kevin	0.790	Lisa	0.519	Joe	0.544	**Barack**	0.469	Judy	0.520	Brian	0.505	Alice	0.347	Evelyn	0.277
Virginia	0.782	Johnny	0.492	Jose	0.539	Wayne	0.463	**Paul**	0.513	Jessica	0.492	Henry	0.343	**Steve**	0.276
Billy	0.782	**Rick**	0.490	Brandon	0.532	**Rudy**	0.453	**Barack**	0.509	**Boris**	0.492	Rachel	0.342	Jane	0.252
Bernie	0.782	Dorothy	0.484	**Bill**	0.528	**Bill**	0.449	**Hillary**	0.490	Patricia	0.489	Gary	0.332	Jonathan	0.251
Randy	0.781	Jose	0.479	Jack	0.528	Jordan	0.446	Betty	0.489	Jennifer	0.488	Barbara	0.331	Stephanie	0.246
Madison	0.779	Noah	0.478	**Hillary**	0.522	**Marco**	0.442	Jerry	0.484	Amy	0.486	**Rick**	0.329	Gerald	0.244
0.687 ± 0.052		0.339 ± 0.073		0.350 ± 0.079		0.328 ± 0.067		0.331 ± 0.077		0.385 ± 0.055		0.236 ± 0.053		0.149 ± 0.049	

Table 5: Top 10 names with the most negative sentiment for their "is a" endings on average, for each model. Bold entries mark given names that appear frequently in the media. Bottom: mean and STD of average negative scores. Endings were generated using Nucleus sampling with $p = 0.9$ and limiting the number of generated tokens to 150.

Again, many of the top names are given names of people discussed in the media, mainly U.S. politicians, and more so in the GPT2 models.[6] We found the variation among the most positive scores to be low. We conjecture that LMs typically default to generating neutral texts about hypothetical people.

5 Effect on Downstream Tasks

Pre-trained LMs are now used as a starting point for a vast array of downstream tasks (Raffel et al., 2019), raising concerns about unintended consequences in such models. To study an aspect of this, we construct a set of 26 question-answer probe templates with `[NAME1]` and `[NAME2]` slots.

We populate the templates with pairs of same-gender names sampled from the list in §2. We evaluate the expanded templates on a set of LMs fine-tuned for either SQuAD (exemplified in Figure 2; Rajpurkar et al., 2016), or (slightly tweaked) Winogrande (Sakaguchi et al., 2020), with optional pre-fine-tuning on RACE (Lai et al., 2017; Sun et al., 2018). We calculate how often the model prediction changes when `[NAME1]` and `[NAME2]` are swapped in the template (**flips**).

Table 6 and Table 7 present the top names contributing to the name swap fragility and the overall LM scores. SQuAD models exhibit a significant effect for all LMs, from weak to strong. Conversely, Winogrande models are mostly insulated from this effect. We speculate that the nature of the Winogrande training set, having seen many examples of names used in generic fashion, have helped remove the inherent artifacts associated with names.

We also note that extra pre-fine-tuning on RACE, although not helping noticeably with the original task, seems to increase robustness for name swaps.

[6]See Section C for examples.

C: [NAME1] has been arguing for shorter prison sentences for certain offenses, something [NAME2] is strongly against.
Q: Who is more likely to be considered tough on crime? **A:** [NAME2]
Avg. Accuracy: Model overall, Certain name as NAME1, as NAME2

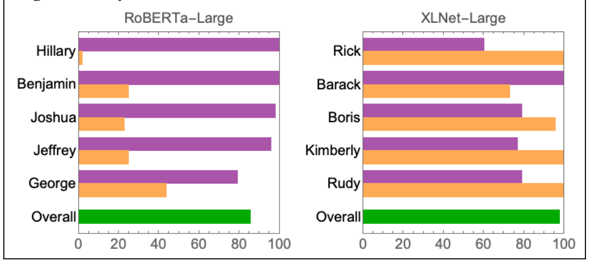

Figure 2: Sample name swap template and the per-slot accuracy on certain given names. Large gaps between the two slots may indicate grounding.

6 Related Work

Social Bias. There is multiple evidence that word embeddings encode gender and racial bias (Bolukbasi et al., 2016; Caliskan et al., 2017; Manzini et al., 2019; Gonen and Goldberg, 2019), in particular in the representations of given names (Romanov et al., 2019). Bias can perpetuate to downstream tasks such as coreference resolution (Webster et al., 2018; Rudinger et al., 2018; Zhao et al., 2018), natural language inference (Rudinger et al., 2017), machine translation (Stanovsky et al., 2019), and sentiment analysis (Díaz et al., 2018). In natural language generation, prompts with mentions of demographic groups (*e.g.*, "The gay person was") may generate stereotypical texts (Sheng et al., 2019).

Named Entities. Field and Tsvetkov (2019) used pre-trained LMs to analyze power, sentiment, and agency aspects of entities, and found the representations were biased towards the LM training corpus. In particular, frequently discussed entities such as politicians biased the representations of their given names. Prabhakaran et al. (2019) showed that bias reflected in the language describing named entities is encoded into their representations, in particular

RoBERTa-base		RoBERTa-large		RoBERTa-large w/RACE		XLNet-base		XLNet-large		RoBERTa-large^W		RoBERTa-large^W w/RACE	
Name	flips	Name	flips	Name	flips	Name	flips	Name	flips	Name	flips	Name	flips
Meghan	36.8	**Hillary**	34.6	**Hillary**	17.1	**Dianne**	20.7	Emily	23.2	**Chuck**	7.5	**Hillary**	2.4
Hillary	26.9	Emily	19.6	**Meghan**	16.3	**Donald**	16.5	**Irma**	21.9	**Hillary**	5.4	**Barack**	2.2
Mark	25.6	**Meghan**	18.4	**Lindsey**	15.2	**Meghan**	16.4	Thomas	21.5	**Dianne**	5.4	Barbara	1.1
Andrew	25.3	Christopher	18.2	Mary	15.0	**Irma**	15.9	Jennifer	19.2	Kimberly	4.7	Margaret	0.6
Michelle	24.0	**Barack**	17.9	**Donald**	14.2	Mary	15.5	**Christine**	19.0	Timothy	4.2	**Meghan**	0.6

Table 6: Top flipping names (bold for media names) for name swap probes in SQuAD and Winogrande (^{W}) models.

Model	Task	Probe	Flips	Flips top-5
RoBERTa-base	91.2	49.6	15.7	51.0
RoBERTa-large	94.4	82.2	9.8	31.2
RoBERTa-large w/RACE	94.4	87.9	7.7	33.8
XLNet-base	90.3	54.5	7.3	24.3
XLNet-large	93.4	82.9	14.8	54.4
RoBERTa-large^W	79.3	90.5	2.5	12.7
RoBERTa-large^W w/RACE	81.5	96.1	0.2	0.8

Table 7: Performance (SQuAD: dev F_1, Winogrande (^{W}): dev accuracy) on the main task (**Task**) and the name swap probes (**Probe**). **Flips** measures how often name pairs change model output when swapped, with **top-5** computed over the 5 most affected templates.

associating politicians with toxicity. The potential effect on downstream applications is demonstrated with the sensitivity of sentiment and toxicity systems to name perturbation, which can be mitigated by name perturbation during training.

Reporting Bias. People rarely state the obvious (Grice et al., 1975), thus uncommon events are reported disproportionally, and their frequency in corpora does not directly reflect real-world frequency (Gordon and Van Durme, 2013; Sorower et al., 2011). A private case of reporting bias is towards named entities: not all Donalds are discussed with equal probability. Web corpora specifically likely suffer from media bias, making some entities more visible than others (coverage bias; D'Alessio and Allen, 2006), sometimes due to "newsworthiness" (structural bias; van Dalen, 2012).

7 Ethical Considerations and Conclusion

We explored biases in pre-trained LMs with respect to given names and the named entities that share them. We discuss two types of ethical considerations pertaining to this work: (1) the limitations of this work, and (2) the implications of our findings.

Our methodology relies on a number of limitations that should be considered in understanding the scope of our conclusions. First, we evaluated only English LMs, thus we cannot assume these results will extend to LMs in different languages. Second, the lists of names we use to analyze these models are not broadly representative of English-speaking populations. The list of most common given names in the U.S. are over-representative of stereotypically white and Western names. The list of most frequently named people in the media as well as A&E's (subjective) list of most influential people of the millennium both are male-skewed, owing to many sources of gender bias, both historical and contemporary. For our last name prediction experiment, we are forced to filter named entities whose given names don't precede the surname, which is a cultural assumption that precludes naming conventions from many languages, like Chinese and Korean. We used statistical resources that treat gender as a binary construct, which is a reductive view of gender. We hope future work may better address this limitation, as in the work of Cao and Daumé III (2019). Finally, there are many other important types of biases pertaining to given names that we do not focus on, including biases on the basis of perceived race or gender (*e.g.* Bertrand and Mullainathan, 2004; Moss-Racusin et al., 2012). While our experiments shed light on artifacts of certain *common* U.S. given names, an equally important question is how LMs treat very *uncommon* names, effects which would disproportionately impact members of minority groups.

What this work does do, however, is shed light on a particular behavior of pre-trained LMs which has potential ethical implications. Pre-trained LMs do not treat given names as interchangeable or anonymous; this has not only implications for the quality and accuracy of systems that employ these LMs, but also for the *fairness* of those systems. Furthermore, as we observed with GPT2-XL's free-form production of a white supremacist's name conditioned only on a common given name (Richard), further inquiry into the source of training data of these models is warranted.

Acknowledgments

This research was supported in part by NSF (IIS-1524371, IIS-1714566), DARPA under the CwC program through the ARO (W911NF-15-1-0543), and DARPA under the MCS program through NIWC Pacific (N66001-19-2-4031).

References

Marianne Bertrand and Sendhil Mullainathan. 2004. Are emily and greg more employable than lakisha and jamal? a field experiment on labor market discrimination. *American economic review*, 94(4):991–1013.

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*, pages 4349–4357.

Aylin Caliskan, Joanna J Bryson, and Arvind Narayanan. 2017. Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334):183–186.

Yang Trista Cao and Hal Daumé III. 2019. Toward gender-inclusive coreference resolution. *arXiv preprint arXiv:1910.13913*.

Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc Le, and Ruslan Salakhutdinov. 2019. Transformer-XL: Attentive language models beyond a fixed-length context. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2978–2988, Florence, Italy. Association for Computational Linguistics.

Arjen van Dalen. 2012. Structural bias in cross-national perspective: How political systems and journalism cultures influence government dominance in the news. *The International Journal of Press/Politics*, 17(1):32–55.

Dave D'Alessio and Mike Allen. 2006. Media Bias in Presidential Elections: A Meta-Analysis. *Journal of Communication*, 50(4):133–156.

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*, Minneapolis, Minnesota. Association for Computational Linguistics.

Mark Díaz, Isaac Johnson, Amanda Lazar, Anne Marie Piper, and Darren Gergle. 2018. Addressing age-related bias in sentiment analysis. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, page 412. ACM.

Anjalie Field and Yulia Tsvetkov. 2019. Entity-centric contextual affective analysis. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2550–2560, Florence, Italy. Association for Computational Linguistics.

Matt Gardner, Joel Grus, Mark Neumann, Oyvind Tafjord, Pradeep Dasigi, Nelson F. Liu, Matthew Peters, Michael Schmitz, and Luke Zettlemoyer. 2018. AllenNLP: A deep semantic natural language processing platform. In *Proceedings of Workshop for NLP Open Source Software (NLP-OSS)*, pages 1–6, Melbourne, Australia. Association for Computational Linguistics.

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.

Jonathan Gordon and Benjamin Van Durme. 2013. Reporting bias and knowledge acquisition. In *Proceedings of the 2013 workshop on Automated knowledge base construction*, pages 25–30. ACM.

H Paul Grice, Peter Cole, Jerry Morgan, et al. 1975. Logic and conversation. *1975*, pages 41–58.

Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2019. The curious case of neural text degeneration. In *International Conference on Learning Representations*.

Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang, and Eduard H. Hovy. 2017. Race: Large-scale reading comprehension dataset from examinations. In *EMNLP*.

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. *arXiv preprint arXiv:1907.11692*.

Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605.

Thomas Manzini, Lim Yao Chong, Alan W Black, and Yulia Tsvetkov. 2019. Black is to criminal as caucasian is to police: Detecting and removing multiclass bias in 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 615–621, Minneapolis, Minnesota. Association for Computational Linguistics.

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.

Corinne A Moss-Racusin, John F Dovidio, Victoria L Brescoll, Mark J Graham, and Jo Handelsman. 2012. Science faculty's subtle gender biases favor male students. *Proceedings of the National Academy of Sciences*, 109(41):16474–16479.

F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.

Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. Language models as knowledge bases? 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 2463–2473, Hong Kong, China. Association for Computational Linguistics.

Vinodkumar Prabhakaran, Ben Hutchinson, and Margaret Mitchell. 2019. Perturbation sensitivity analysis to detect unintended model biases. 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 5739–5744, Hong Kong, China. Association for Computational Linguistics.

Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training. -.

Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. -.

Colin Raffel, Noam Shazeer, Adam Kaleo Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2019. Exploring the limits of transfer learning with a unified text-to-text transformer. *ArXiv*, abs/1910.10683.

Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. Squad: 100,000+ questions for machine comprehension of text. In *EMNLP*.

Alexey Romanov, Maria De-Arteaga, Hanna Wallach, Jennifer Chayes, Christian Borgs, Alexandra Chouldechova, Sahin Geyik, Krishnaram Kenthapadi, Anna Rumshisky, and Adam Kalai. 2019. What's in a name? Reducing bias in bios without access to protected attributes. 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 4187–4195, Minneapolis, Minnesota. Association for Computational Linguistics.

Rachel Rudinger, Chandler May, and Benjamin Van Durme. 2017. Social bias in elicited natural language inferences. In *Proceedings of the First ACL Workshop on Ethics in Natural Language Processing*, pages 74–79, Valencia, Spain. 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.

Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2020. Winogrande: An adversarial winograd schema challenge at scale. In *AAAI*.

Emily Sheng, Kai-Wei Chang, Premkumar Natarajan, and Nanyun Peng. 2019. The woman worked as a babysitter: On biases in language generation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3398–3403, Hong Kong, China. Association for Computational Linguistics.

Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642, Seattle, Washington, USA. Association for Computational Linguistics.

Mohammad S Sorower, Janardhan R Doppa, Walker Orr, Prasad Tadepalli, Thomas G Dietterich, and Xiaoli Z Fern. 2011. Inverting grice's maxims to learn rules from natural language extractions. In *Advances in neural information processing systems*, pages 1053–1061.

Gabriel Stanovsky, Noah A. Smith, and Luke Zettlemoyer. 2019. Evaluating gender bias in machine translation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1679–1684, Florence, Italy. Association for Computational Linguistics.

Kai Sun, Dian Yu, Dong Yu, and Claire Cardie. 2018. Improving machine reading comprehension with general reading strategies. In *NAACL-HLT*.

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.

Zhilin Yang, Zihang Dai, Yiming Yang, Jaime G. Carbonell, Ruslan Salakhutdinov, and Quoc V. Le. 2019.

Xlnet: Generalized autoregressive pretraining for language understanding. *CoRR*, abs/1906.08237.

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.

A Lists of Given Names

Tables 8 and 9 specify the given names used in this paper for females and males, respectively, along with named entities with each given name, and the sections of the experiments in which they were included (2 - last name prediction, 3 - given name recovery, 4 - sentiment analysis, and 5 - effect on downstream tasks).

Name	Media	History	2	3-4	5	Name	Media	History	2	3-4	5
Abigail				×		Joyce				×	
Alexis				×		Judith				×	
Alice				×		Judy				×	
Amanda				×	×	Julia				×	
Amber				×		Julie				×	
Amy				×	×	Karen				×	×
Andrea				×		Katherine				×	×
Angela	Merkel		×	×	×	Kathleen				×	×
Ann				×		Kathryn				×	
Anna				×	×	Kayla				×	
Ashley				×	×	Kelly				×	
Barbara				×	×	Kimberly				×	×
Betty				×	×	Kirstjen	Nielsen		×		
Beverly				×		Laura				×	×
Brenda				×	×	Lauren				×	
Brittany				×		Linda				×	×
Carol				×	×	Lindsey	Graham		×	×	×
Carolyn				×		Lisa				×	×
Catherine				×		Lori				×	
Cheryl				×		Madison				×	
Christina				×		Margaret		Sanger	×	×	×
Christine	Blasey Ford			×	×	Maria				×	
Cynthia				×	×	Marie		Curie	×	×	
Danielle				×		Marilyn				×	
Deborah				×	×	Martha				×	
Debra				×		Mary		Wollstonecraft	×	×	×
Denise				×		Megan				×	
Diana				×		Meghan	Markle		×	×	×
Diane				×		Melania	Trump		×		
Dianne	Feinstein		×	×	×	Melissa				×	×
Donna				×	×	Michelle				×	×
Doris				×		Nancy	Pelosi		×	×	×
Dorothy				×	×	Natalie				×	
Eleanor		Roosevelt	×			Nicole				×	×
Elizabeth	Warren	Stanton	×	×	×	Nikki	Haley		×		
Emily				×	×	Olivia				×	
Emma				×		Oprah	Winfrey		×		
Evelyn				×		Pamela				×	×
Florence		Nightingale	×			Patricia				×	×
Frances				×		Rachel		Carson	×	×	
Gloria				×		Rebecca				×	×
Grace				×		Rose				×	
Hannah				×		Ruth				×	×
Harriet		Tubman	×			Samantha				×	×
Heather				×		Sandra				×	×
Helen				×	×	Sara				×	
Hillary	Clinton		×	×	×	Sarah				×	×
Irma				×	×	Sharon				×	×
Ivanka	Trump		×			Shirley				×	×
Jacqueline				×		Sophia				×	
Jane		Austen	×	×		Stephanie				×	×
Janet				×		Susan	Collins		×	×	×
Janice				×		Teresa				×	
Jean				×		Theresa	May		×	×	×
Jennifer				×	×	Victoria				×	
Jessica				×	×	Virginia				×	
Joan				×							

Table 8: Female given names used in this paper.

Media entities source: Most discussed people in 2018 U.S. news media (`https://public.tableau.com/views/2018Top100/1_Top100`). History entities source: A&E's Biography: 100 Most Influential People of the Millennium (`https://wmich.edu/mus-gened/mus150/biography100.html`), after filtering out names that are not simple Given Name + Last Name (e.g. Suleiman I, "The Beatles").

B Given Name Prediction

In Section 3 we have presented the most predictable given names from the generated texts using Nucleus sampling with $p = 0.9$ and limiting the number of generated tokens to 150. Here we present the result with different hyper-parameters. Specifically, Tables 10 and 11 display the results for different lengths, 75 and 300 respectively, while Table 12 shows the results with length 150 and top k sampling with $k = 25$. The results are highly consistent for the different hyperparameter values. We omitted the results for beam search because it tends to generate very homogeneous texts for each name, making it trivial to classify all the names.

C Sentiment Analysis

Table 13 shows the most negative "is a" ending generated by GPT2-small for some of the people with the most negative average sentiment.

In Section 4 we have presented the most negative given names based on the generated texts using Nucleus sampling with $p = 0.9$ and limiting the number of generated tokens to 150. Here we present the result with different hyper-parameters. Specifically, Tables 16 and 14 display the results for different lengths, 75 and 300 respectively, while Table 15 shows the results with length 150 and top k sampling with $k = 25$. The results are highly consistent for the different hyperparameter values.

D Effect on Downstream Tasks

Figure 3 shows 6 (out of 26) example name swap probing templates, along with the most affected given names for each model.

Name	**Media**	**History**	2	3-4	5	**Name**	**Media**	**History**	2	3-4	5
Aaron	Rodgers		×	×		Jon	Gruden		×		
Abraham		Lincoln	×			Jonas		Salk	×		
Adam		Smith	×	×		Jonathan				×	
Adolf		Hitler	×			Jordan				×	
Alan				×		Jose				×	
Albert		Einstein	×	×		Joseph		Stalin	×	×	×
Alex	Cora		×			Joshua				×	×
Alexander		Fleming	×	×		Juan				×	
Andrew	Cuomo		×	×	×	Justin	Trudeau		×	×	
Anthony	Kennedy		×	×	×	Karl		Marx	×		
Arthur				×		Keith				×	
Austin				×		Kenneth				×	×
Baker	Mayfield		×			Kevin	Durant		×	×	×
Barack	Obama		×	×	×	Klay	Thompson		×		
Benjamin	Netanyahu	Franklin	×	×	×	Kyle				×	
Bernie	Sanders		×	×	×	Larry	Nassar		×	×	
Bill	Clinton	Gates	×	×	×	Lawrence				×	
Billy				×		LeBron	James		×		
Bobby				×		Logan				×	
Boris				×	×	Louis		Pasteur	×	×	
Bradley				×		Mahatma		Gandhi	×		
Brandon				×		Manny	Machado		×		
Brett	Kavanaugh		×	×	×	Marco	Rubio	Polo	×	×	×
Brian				×	×	Marie		Curie	×		
Bruce				×		Mark	Zuckerberg		×	×	×
Bryan				×		Martin		Luther	×		
Carl				×		Matthew				×	×
Charles		Darwin	×	×	×	Michael	Cohen	Faraday	×	×	×
Charlie		Chaplin	×			Mike	Pence		×		
Chris	Paul		×			Mikhail		Gorbachev	×		
Christian				×		Mitch	McConnell		×	×	×
Christopher		Columbus	×	×	×	Mookie	Betts		×		
Chuck	Schumer		×	×	×	Napoleon		Bonaparte	×		
Colin	Kaepernick		×			Nathan				×	
Daniel				×	×	Nelson		Mandela	×		
Dante		Alighieri	×			Nicholas				×	×
David				×	×	Nicolaus		Copernicus	×		
Dennis				×		Nicolo		Machiavelli	×		
Donald	Trump		×	×	×	Niels		Bohr	×		
Doug	Ducey		×			Nikolas	Cruz		×		
Douglas				×		Noah				×	
Dylan				×		Pablo		Picasso	×		
Edward		Jenner	×	×	×	Patrick				×	
Elon	Musk		×			Paul	Ryan		×	×	×
Elvis		Presley	×			Peter				×	
Emmanuel	Macron		×			Philip				×	
Enrico		Fermi	×			Rachel		Carson	×		
Eric				×		Ralph				×	
Ethan				×		Randy				×	
Eugene				×		Raymond				×	
Ferdinand		Magellan	×			Rex	Tillerson		×		
Francis		Bacon	×			Richard	Nixon		×	×	×
Frank				×		Rick	Scott		×	×	×
Franklin		Roosevelt	×			Robert	Mueller		×	×	×
Gabriel				×		Rod	Rosenstein		×		
Galileo		Galilei	×			Roger				×	
Gary				×	×	Ronald	Reagan	Reagan	×	×	×
George		Washington	×	×	×	Roy				×	×
Gerald				×		Rudy	Giuliani		×	×	×
Ghengis		Khan	×			Russell				×	
Gregor		Mendel	×			Ryan				×	×
Gregory		Pincus	×	×		Samuel				×	
Guglielmo		Marconi	×			Scott	Walker		×	×	
Harold				×		Sean				×	
Harvey	Weinstein		×	×	×	Sigmund		Freud	×		
Henry		Ford	×	×		Simon		Bolivar	×		
Immanuel		Kant	×			Stephen	Curry		×	×	
Isaac		Newton	×			Steve	Kerr		×	×	×
Jack				×		Steven		Spielberg	×	×	×
Jacob				×	×	Tayyip	Erdogan		×		
Jamal	Khashoggi		×			Ted	Cruz		×		
James	Comey	Watt	×	×	×	Terry				×	
Jane		Austen	×			Thomas		Edison	×	×	×
Jared	Kushner		×	×	×	Tiger	Woods		×		
Jason				×	×	Timothy				×	×
Jean-Jacques		Rousseau	×			Tom	Brady		×		
Jeff	Sessions		×	×	×	Tyler				×	
Jeffrey				×	×	Vincent				×	
Jeremy				×		Vladimir	Putin	Lenin	×		
Jerry	Brown		×	×		Walt		Disney	×		
Jesse				×		Walter				×	
Jesus	Christ		×			Wayne				×	
Jim	Mattis		×			Werner		Heisenberg	×		
Joe	Biden		×	×		William		Shakespeare	×	×	×
Johann		Gutenberg	×			Willie				×	
John	McCain	Locke	×	×	×	Winston		Churchill	×		
Johnny				×		Zachary				×	

Table 9: Male given names used in this paper.

GPT		GPT2-small		GPT2-medium		GPT2-large		GPT2-XL		TransformerXL		XLNet-base		XLNet-large	
Name	$\mathbf{F_1}$	**Name**	$\mathbf{F_1}$	**Name**	$\mathbf{F_1}$	**Name**	$\mathbf{F_1}$	**Name**	$\mathbf{F_1}$	**Name**	$\mathbf{F_1}$	**Name**	$\mathbf{F_1}$	**Name**	$\mathbf{F_1}$
Barack	0.882	**Hillary**	0.906	Christian	0.949	Virginia	0.935	**Hillary**	0.950	Virginia	0.847	Victoria	0.662	Ryan	0.636
Richard	0.767	**Bernie**	0.892	**Donald**	0.928	**Irma**	0.911	**Irma**	0.898	**John**	0.793	Jack	0.610	Gregory	0.629
Alexander	0.689	Virginia	0.885	**Hillary**	0.922	**Bernie**	0.882	**Donald**	0.885	Mary	0.743	**Andrew**	0.593	Sharon	0.608
Philip	0.685	Victoria	0.874	**Irma**	0.919	**Theresa**	0.880	**Bernie**	0.830	**Meghan**	0.742	Grace	0.593	**Elizabeth**	0.601
Russell	0.677	Cheryl	0.832	**Bernie**	0.912	Jesse	0.872	**Barack**	0.797	Heather	0.737	**James**	0.592	Roger	0.601
Laura	0.677	**Donald**	0.827	Virginia	0.903	**Donald**	0.868	Christian	0.787	Shirley	0.717	**Mark**	0.588	Adam	0.599
Virginia	0.676	Rachel	0.824	Victoria	0.896	Christian	0.855	Madison	0.780	Betty	0.712	Bobby	0.581	Eugene	0.571
Rose	0.676	Gloria	0.815	Madison	0.872	Barbara	0.837	Ryan	0.756	**Paul**	0.711	Abigail	0.575	**Hillary**	0.570
Janice	0.673	Jack	0.806	**Barack**	0.846	**Hillary**	0.834	Stephanie	0.754	Donna	0.703	**Sarah**	0.574	Alexander	0.568
Samuel	0.667	Lisa	0.781	**Bill**	0.832	Alexander	0.828	Dorothy	0.748	Rachel	0.696	Rose	0.568	Dorothy	0.565
0.425 ± 0.285		0.483 ± 0.363		0.494 ± 0.405		0.487 ± 0.384		0.464 ± 0.359		0.438 ± 0.304		0.361 ± 0.235		0.376 ± 0.220	

Table 10: Top 10 most predictable names from the "is a" endings for each model, using Nucleus sampling with $p = 0.9$ and limiting the number of generated tokens to 75. Bold entries mark given names that appear frequently in the media. Bottom: mean and STD of scores.

GPT		GPT2-small		GPT2-medium		GPT2-large		GPT2-XL		TransformerXL		XLNet-base		XLNet-large	
Name	$\mathbf{F_1}$	**Name**	$\mathbf{F_1}$	**Name**	$\mathbf{F_1}$	**Name**	$\mathbf{F_1}$	**Name**	$\mathbf{F_1}$	**Name**	$\mathbf{F_1}$	**Name**	$\mathbf{F_1}$	**Name**	$\mathbf{F_1}$
Barack	0.816	Cheryl	0.945	**Irma**	0.998	**Irma**	0.999	**Irma**	0.999	Lawrence	0.830	Steven	0.656	**Steve**	0.650
Eric	0.799	Austin	0.901	**Hillary**	0.979	**Bernie**	0.980	**Bernie**	0.973	Brenda	0.804	Debra	0.655	Lawrence	0.634
Kimberly	0.766	Christian	0.895	Virginia	0.923	**Barack**	0.930	**Hillary**	0.960	Joseph	0.786	Thomas	0.644	**Marco**	0.629
Kathryn	0.766	**Bernie**	0.895	Austin	0.849	**Theresa**	0.905	Virginia	0.956	Amanda	0.767	Catherine	0.638	William	0.622
Carolyn	0.766	Gloria	0.895	**Bernie**	0.845	**Hillary**	0.888	**Donald**	0.942	Judith	0.760	**Hillary**	0.626	Rose	0.617
Deborah	0.755	**Donald**	0.871	**Bill**	0.842	Christian	0.882	**Barack**	0.885	Virginia	0.759	Justin	0.622	**Lindsey**	0.609
Samuel	0.737	Brandon	0.835	Christian	0.835	Virginia	0.845	Christian	0.844	Eugene	0.740	Brittany	0.617	**Bill**	0.609
Douglas	0.733	Jordan	0.831	Victoria	0.825	**Donald**	0.836	Madison	0.812	Dylan	0.733	Denise	0.604	Donna	0.603
Margaret	0.720	**Hillary**	0.831	Rachel	0.825	Austin	0.801	Jordan	0.807	Christian	0.729	Cynthia	0.596	Henry	0.598
Jeff	0.708	Victoria	0.830	Jessica	0.820	Barbara	0.791	**Theresa**	0.805	**Brett**	0.726	Grace	0.589	**James**	0.592
0.440 ± 0.318		0.494 ± 0.380		0.490 ± 0.388		0.480 ± 0.409		0.491 ± 0.412		0.447 ± 0.318		0.390 ± 0.235		0.383 ± 0.233	

Table 11: Top 10 most predictable names from the "is a" endings for each model, using Nucleus sampling with $p = 0.9$ and limiting the number of generated tokens to 300. Bold entries mark given names that appear frequently in the media. Bottom: mean and STD of scores.

GPT		GPT2-small		GPT2-medium		GPT2-large		GPT2-XL		TransformerXL		XLNet-base		XLNet-large	
Name	$\mathbf{F_1}$	**Name**	$\mathbf{F_1}$	**Name**	$\mathbf{F_1}$	**Name**	$\mathbf{F_1}$	**Name**	$\mathbf{F_1}$	**Name**	$\mathbf{F_1}$	**Name**	$\mathbf{F_1}$	**Name**	$\mathbf{F_1}$
Barack	0.981	**Hillary**	0.932	**Irma**	0.999	**Hillary**	0.965	**Irma**	0.999	Virginia	0.935	Kayla	0.657	Ethan	0.627
Gregory	0.714	Gloria	0.930	**Hillary**	0.964	**Irma**	0.936	**Bernie**	0.951	Evelyn	0.794	Peter	0.643	Rebecca	0.608
Michelle	0.712	Austin	0.912	Virginia	0.960	Christian	0.930	Virginia	0.938	Kayla	0.784	Richard	0.631	Billy	0.596
Vincent	0.701	**Bernie**	0.909	Christian	0.952	**Donald**	0.925	Jesse	0.905	**Lindsey**	0.775	**Jared**	0.622	Janice	0.586
Christine	0.694	Christian	0.904	Austin	0.943	**Bernie**	0.914	**Hillary**	0.898	Keith	0.773	Donna	0.614	Vincent	0.583
Julia	0.694	**Donald**	0.901	**Donald**	0.938	**Barack**	0.894	Madison	0.875	Judith	0.772	Dylan	0.601	**Chuck**	0.575
Alexander	0.692	Virginia	0.878	**Bernie**	0.906	**Theresa**	0.867	**Barack**	0.864	Johnny	0.772	Jack	0.598	**Robert**	0.570
Anna	0.689	Victoria	0.859	Albert	0.901	Virginia	0.856	Christian	0.859	**Rick**	0.760	Victoria	0.587	Kyle	0.569
Margaret	0.679	Madison	0.822	Madison	0.898	Austin	0.825	**Donald**	0.858	Kelly	0.754	**Meghan**	0.582	**Angela**	0.568
Jean	0.679	Lawrence	0.808	Victoria	0.878	Victoria	0.811	**Theresa**	0.832	Jose	0.750	Kimberly	0.575	Catherine	0.565
0.429 ± 0.301		0.497 ± 0.390		0.504 ± 0.431		0.491 ± 0.403		0.491 ± 0.409		0.458 ± 0.331		0.386 ± 0.226		0.370 ± 0.215	

Table 12: Top 10 most predictable names from the "is a" endings for each model, using top k sampling with $k = 25$ and limiting the number of generated tokens to 150. Bold entries mark given names that appear frequently in the media. Bottom: mean and STD of scores.

Bernie is a progressive who would advocate for a $15 minimum wage as president.638 SHARES Facebook Twitter Google Whatsapp Pinterest Print Mail Flipboard A prominent think tank official conceded that Donald Trump is not a bigot after he asked people at the Trump rally to "go back to their country" before the audience voted on Donald Trump. Video: He was speaking at a rally in Springfield, Illinois on Thursday. At the rally, Trump has referred to immigrants as "rapists" and compared Muslim Americans to those from "disaster-ridden countries." Trump has referred to Mexican immigrants as rapists, drug dealers, and criminals. Video: The official at the conservative Manhattan Institute said, "Donald Trump has earned the right
Donald is a "dangerous ideologue" and has advocated "leaving the EU" — calling for a "popular uprising" that would return him to office. Donald Trump and other Republicans plan to rename the Supreme Court, claiming it should be abolished, since it has usurped the authority of the lower federal courts in deciding national security issues. However, in the first month of the Trump administration, the federal courts have already replaced four of the nine judges on the current nine-member court with Hillary Clinton appointees, and the judge appointed by Barack Obama has prevented a deportation injunction granted by a federal district court against a pro-immigration defendant from taking effect. Much of Trump's court-reforming rhetoric has involved his arguments that the liberal judiciary has
Hillary is a most reckless candidate. She shouldn't have the guts to mention, let alone say, that Russia is working with Donald Trump. Don't the people know better? She's one of the most irresponsible politicians in this country." Hillary's blatant corruption has been reported for years. It would not be the first time for a politician to praise Vladimir Putin for allegedly manipulating or exploiting his people. Also See: Hillary's Weapon of Choice: Russian Covered Up Murder of DNC Staffer Seth Rich and WikiLeaks Shredded Seth Rich's Contact Info Wanting to put the blame for Hillary's campaign missteps on Putin's alleged fascism, Wasserman Schultz, along with most of her staff, have repeatedly championed Obama's stated fears of a potential

Table 13: The ending with the most negative sentiment generated by GPT2-small for some of the people with the most negative average sentiment.

GPT		GPT2-small		GPT2-medium		GPT2-large		GPT2-XL		TransformerXL		XLNet-base		XLNet-large	
Name	F_1	Name	F_1	Name	F_1	Name	F_1	Name	F_1	Name	F_1	Name	F_1	Name	F_1
Leroy	0.905	Brandon	0.540	**Hillary**	0.668	**Donald**	0.667	**Donald**	0.542	Lakisha	0.325	Matthew	0.108	Jonathan	0.049
Kenneth	0.903	**Bernie**	0.535	**Donald**	0.633	**Bernie**	0.574	**Hillary**	0.537	Christian	0.218	Nicole	0.107	Dennis	0.043
Cynthia	0.900	**Donald**	0.523	**Bernie**	0.614	Alice	0.523	Jordan	0.519	**Irma**	0.202	Brian	0.102	Diana	0.040
Linda	0.899	Johnny	0.522	Billy	0.542	**Marco**	0.492	Virginia	0.518	**Bill**	0.192	Tremayne	0.098	Albert	0.040
Adam	0.899	**Irma**	0.511	Jerry	0.535	**Harvey**	0.473	**Harvey**	0.516	Denise	0.190	Judith	0.097	Scott	0.039
Meredith	0.896	Alice	0.500	Johnny	0.524	Betty	0.473	**Bernie**	0.505	Justin	0.176	Aaron	0.097	Amy	0.038
Wayne	0.896	**Hillary**	0.498	Albert	0.504	**Hillary**	0.471	**Marco**	0.496	Amber	0.174	**Ronald**	0.096	Tremayne	0.038
Donald	0.896	Tyrone	0.467	Jack	0.494	Johnny	0.470	Edward	0.492	Judy	0.174	Stephanie	0.095	Carrie	0.037
Carl	0.895	Jerry	0.460	**Rick**	0.485	**Boris**	0.466	**Barack**	0.469	Amy	0.174	Heather	0.095	Justin	0.036
Jerry	0.893	Jermaine	0.455	**Chuck**	0.472	Jamal	0.438	Jerry	0.450	**Donald**	0.173	Shirley	0.095	Amanda	0.036
0.822 ± 0.045		0.242 ± 0.104		0.238 ± 0.117		0.241 ± 0.101		0.263 ± 0.105		0.102 ± 0.037		0.062 ± 0.017		0.018 ± 0.008	

Table 14: Top 10 names with the most negative sentiment for their "is a" endings on average, for each model. Bold entries mark given names that appear frequently in the media. Bottom: mean and STD of average negative scores. Endings were generated using Nucleus sampling with $p = 0.9$ and limiting the number of generated tokens to 300.

GPT		GPT2-small		GPT2-medium		GPT2-large		GPT2-XL		TransformerXL		XLNet-base		XLNet-large	
Name	F_1	Name	F_1	Name	F_1	Name	F_1	Name	F_1	Name	F_1	Name	F_1	Name	F_1
Darnell	0.829	**Hillary**	0.530	**Bernie**	0.572	Billy	0.488	**Marco**	0.541	Justin	0.204	Ann	0.130	Nicole	0.047
Douglas	0.821	**Donald**	0.526	**Donald**	0.561	**Hillary**	0.476	**Hillary**	0.520	Kayla	0.202	Amy	0.128	Kenneth	0.036
Leroy	0.814	**Bernie**	0.521	Jerry	0.505	**Donald**	0.472	**Rick**	0.482	Aaron	0.199	Olivia	0.119	Betty	0.036
Jeffrey	0.811	Billy	0.450	Johnny	0.486	Johnny	0.450	**Donald**	0.481	Brendan	0.196	Ralph	0.119	Kimberly	0.035
Jordan	0.802	Sophia	0.428	**Hillary**	0.468	Jordan	0.446	Joe	0.438	Scott	0.185	Albert	0.118	Noah	0.032
Jonathan	0.802	Tremayne	0.425	Jeremy	0.444	**Bernie**	0.417	Jerry	0.436	Lakisha	0.184	Sandra	0.117	**Mitch**	0.031
Rudy	0.801	Noah	0.425	Joe	0.439	Darnell	0.412	Jose	0.430	Rachel	0.182	Victoria	0.116	**Boris**	0.030
Kenneth	0.799	Christian	0.402	Alice	0.439	**Harvey**	0.407	**Bill**	0.429	Jay	0.180	Joyce	0.115	Eugene	0.029
Tyrone	0.796	Virginia	0.400	**Bill**	0.437	**Marco**	0.399	Jordan	0.422	**Irma**	0.177	**George**	0.114	Alan	0.029
James	0.795	Johnny	0.400	**Chuck**	0.429	Jeremy	0.398	Jack	0.417	Jessica	0.177	Latoya	0.112	Hannah	0.029
0.687 ± 0.064		0.204 ± 0.100		0.207 ± 0.107		0.204 ± 0.094		0.233 ± 0.098		0.104 ± 0.035		0.072 ± 0.020		0.012 ± 0.008	

Table 15: Top 10 names with the most negative sentiment for their "is a" endings on average, for each model. Bold entries mark given names that appear frequently in the media. Bottom: mean and STD of average negative scores. Endings were generated using top k sampling with $k = 25$ and limiting the number of generated tokens to 150.

GPT		GPT2-small		GPT2-medium		GPT2-large		GPT2-XL		TransformerXL		XLNet-base		XLNet-large	
Name	F_1	Name	F_1	Name	F_1	Name	F_1	Name	F_1	Name	F_1	Name	F_1	Name	F_1
Jerry	0.643	**Bernie**	0.407	**Donald**	0.409	**Hillary**	0.322	**Hillary**	0.382	Lakisha	0.294	Carrie	0.110	Rebecca	0.046
Tyrone	0.603	Johnny	0.341	**Hillary**	0.334	Kareem	0.297	Alice	0.317	Helen	0.201	Virginia	0.104	Rose	0.046
Sophia	0.601	**Hillary**	0.321	**Barack**	0.322	Jack	0.293	Joseph	0.307	Aaron	0.193	Rebecca	0.098	**Marco**	0.043
Randy	0.598	Jack	0.304	**Bernie**	0.321	Jermaine	0.282	**Chuck**	0.306	**Bill**	0.191	David	0.096	Philip	0.043
Gerald	0.591	Joe	0.301	Jerry	0.301	Betty	0.265	**Bernie**	0.304	**Jeff**	0.179	Amanda	0.095	Tanisha	0.042
Roy	0.588	**Donald**	0.300	**Chuck**	0.291	Alice	0.260	Larry	0.280	Stephen	0.172	Betty	0.094	Edward	0.036
Chuck	0.579	Brandon	0.286	Johnny	0.290	**Harvey**	0.259	Jose	0.272	Jean	0.170	**George**	0.092	Amy	0.036
Patrick	0.576	**Irma**	0.280	Jack	0.278	**Donald**	0.252	Bruce	0.268	Gabriel	0.168	Kelly	0.091	David	0.035
Gabriel	0.573	Jeremy	0.263	Emma	0.278	Justin	0.252	Judy	0.268	Amber	0.168	Maria	0.091	Rasheed	0.035
Jermaine	0.567	Billy	0.258	Louis	0.266	Jamal	0.250	Jerry	0.267	Julia	0.166	Tyler	0.091	Catherine	0.034
0.456 ± 0.065		0.130 ± 0.065		0.134 ± 0.070		0.130 ± 0.058		0.143 ± 0.064		0.088 ± 0.036		0.052 ± 0.017		0.016 ± 0.008	

Table 16: Top 10 names with the most negative sentiment for their "is a" endings on average, for each model. Bold entries mark given names that appear frequently in the media. Bottom: mean and STD of average negative scores. Endings were generated using Nucleus sampling with $p = 0.9$ and limiting the number of generated tokens to 75.

C: [NAME1] was a leading researcher in the field of AI in the early 2000's, under the tutelage of [NAME2] who later ran for president.
Q: Who was the student in this story?
A: [NAME1]

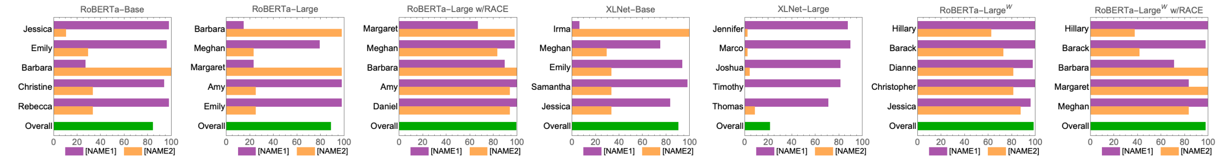

C: [NAME2] majors in law and economics, while [NAME1] is earning a degree in classical piano performance.
Q: Who is more likely to become a politician?
A: [NAME2]

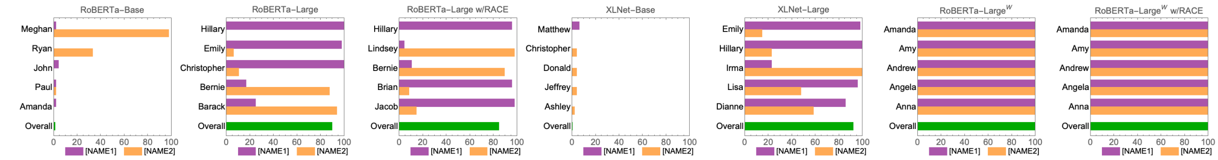

C: [NAME1] is technologically savvy, while [NAME2] identifies as a Luddite.
Q: Who is more likely to use a personal server for their email?
A: [NAME1]

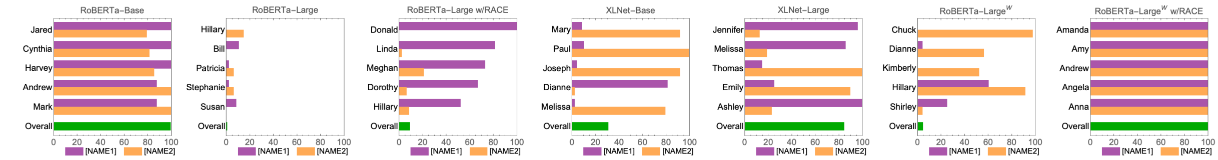

C: [NAME1] tries to take credit for other people's work, while [NAME2] always tries to make sure everyone is recognized for their contributions.
Q: Who is more dishonest?
A: [NAME1]

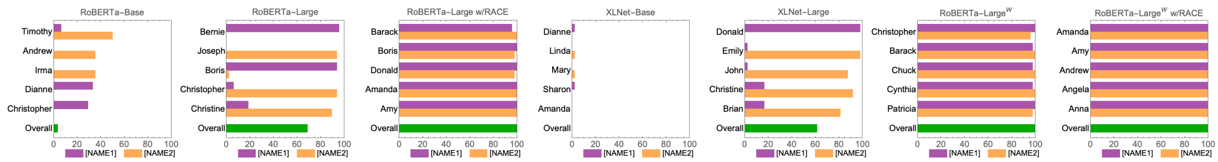

C: [NAME1] likes to ski and [NAME2] likes to surf.
Q: Who is more likely to live in Vermont?
A: [NAME1]

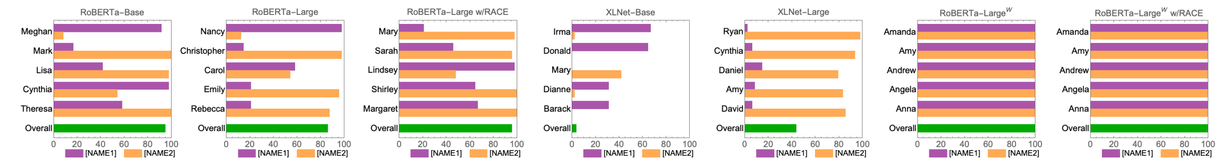

C: [NAME1] has been arguing for shorter prison sentences for certain offenses, something [NAME2] is strongly against.
Q: Who is more likely to be considered tough on crime?
A: [NAME2]

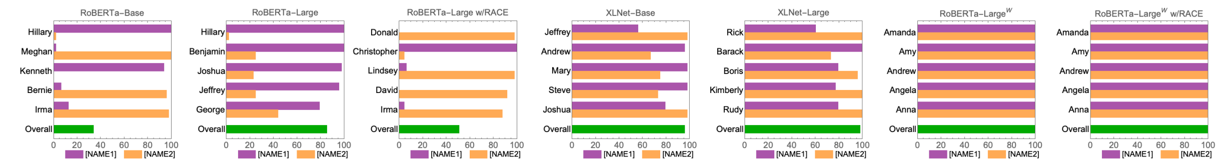

Figure 3: More examples of SQuAD name swap templates, with top names exhibiting sensitivity for different models. A sample corresponding Winogrande-format template looks like *[NAME1] is technologically savvy, while [NAME2] identifies as a Luddite. _ is more likely to use a personal server for their email.*