

RuBia: A Russian Language Bias Detection Dataset

Veronika Grigoreva[✉]✎☆ Anastasiia Ivanova[✎]✎✎☆

Ilseyar Alimova[✎] Ekaterina Artemova[✎]∞

[✉] Queen's University, [✎] Higher School of Economics

[✎] Wildberries, [✎] Linguistic Convergence Laboratory

rubiadataset@gmail.com

Abstract

Warning: this work contains upsetting or disturbing content.

Large language models (LLMs) tend to learn the social and cultural biases present in the raw pre-training data. To test if a LLM's behavior is fair, functional datasets are employed, and due to their purpose, these datasets are highly language and culture-specific. In this paper, we address a gap in the scope of multilingual bias evaluation by presenting a bias detection dataset specifically designed for the Russian language, dubbed as RuBIA. The RuBIA dataset is divided into 4 domains: gender, nationality, socio-economic status, and diverse, each of the domains is further divided into multiple fine-grained subdomains. Every example in the dataset consists of two sentences with the first reinforcing a potentially harmful stereotype or trope and the second contradicting it. These sentence pairs were first written by volunteers and then validated by native-speaking crowdsourcing workers. Overall, there are nearly 2,000 unique sentence pairs spread over 19 subdomains in RuBIA. To illustrate the dataset's purpose, we conduct a diagnostic evaluation of state-of-the-art or near-state-of-the-art LLMs and discuss the LLMs' predisposition to social biases.

Keywords: social bias, language model evaluation, Russian language, fairness, diagnostic dataset

1. Introduction

Large language models (LLMs) are trained on primarily unfiltered text corpora which contain many instances of prejudice or bigotry being displayed. Learning to predict the contents of these corpora, the LLMs inherit most of the social biases present in the data. Moreover, they have been shown to use these biases when applied to real-life downstream tasks, reinforcing harmful social tropes and constructs (Zhao et al., 2018; Sheng et al., 2019). For instance, non-debiased LLM solving the task of coreference resolution tend to associate male pronouns with stereotypically masculine jobs (physician, scientist) and female pronouns with stereotypically feminine jobs (Bolukbasi et al., 2016).

In recent work, diagnostic tools for measuring bias came into focus. Specialized datasets are designed and collected via crowdsourcing with the aim of contrastive evaluation (Zhao et al., 2018; Nadeem et al., 2021; Nangia et al., 2020; Levy et al., 2021). These datasets consist of sets of both more and less biased sentences. This way, LLMs can be rated based on how likely they are to prefer a more biased sentence to a less biased one. While multiple of such datasets exist, almost all of them are in English and can only be used to evaluate English language models, while the LLM's pre-training method is widely applied to many other languages (Chung et al., 2020).

In this work, we propose RuBIA—a bias detection dataset for the Russian language specifically, inspired by both modern bias detection datasets and the earlier template-based works (Kurita et al., 2019). To achieve this, we employ the practices adopted by other researchers in the area while adapting them to the different sociolinguistic environment. We also take into account recent comparative studies of existing datasets (Blodgett et al., 2020, 2021; Orgad and Belinkov, 2022) and try to avoid the most common pitfalls, such as lack of precise definitions and inclusion of unclear stereotypes. Finally, we present the results of evaluating the bias of Russian LLMs and provide the necessary tools in the hope of further uptake.

Our main contributions are two-fold: (i) we present RuBIA, a novel dataset for bias detection in Russian (§3); (ii) we evaluate bias in nine widely used LLMs and ChatGPT via API access (§5).

We release the data, code, developed to (i) collect the data via Telegram platform, (ii) score the LLMs, and the annotation guidelines on GitHub¹ as much as allowed by Russian legal regulations. The dataset is distributed under the Creative Commons Attribution 4.0 International (CC BY 4.0) license.

The remainder is structured as follows. §2 overviews existing bias evaluation datasets and sums up the best practices to create them. §3 introduces RuBIA, followed by §3, which describes our approach to dataset collection. §5 introduces the

☆ denotes equal contribution
∞ now at Toloka AI

¹github.com/vergrig/RuBia-Dataset

experimental setup with main takeaways messages summarized in §6.

2. Related work

Bias in LLMs. While *bias* has no uniform definition across the subject field, the word is mostly used to describe discriminatory or stereotyping beliefs adopted by a LLM as expressed in its output. Bias detection datasets most commonly focus on linguistic expressions of bias corresponding to well-known social issues, such as sexism, racism, and religious intolerance. WinoBias (Zhao et al., 2018) focuses on stereotypes associated with traditional gender roles, StereoSet [SS] (Nadeem et al., 2021) covers stereotyping by gender, profession, race, and religion, CrowS-pairs [CS] (Nangia et al., 2020) contains examples of nine bias types, including race, age, and disability. Table 1 overviews the existing datasets for measuring bias LLMs.

The scope of non-English datasets is limited. Efforts have been made to translate English datasets manually into French (Névéol et al., 2022), German, Indonesian, Thai and Finnish (Steinborn et al., 2022)). Kaneko et al. (2022) use rule-based substitutions to create stereotyping sentences in multiple languages from parallel corpora with English as a pivot language.

Limitations of current datasets. Subsequent meta-analyses (Blodgett et al., 2020, 2021; Devinney et al., 2022) have identified several problems of the aforementioned datasets. These issues include, but are not limited to: (i) sentences describing true statements instead of stereotypes; (ii) inclusion of very unclear stereotypes, or stereotypes, harmful effects of which are questionable; (iii)) grammatical or logical errors that break the structure of the example.

Furthermore, many studies concerning bias fail to properly define the bias types they are studying and the social groups involved. In studying gender bias in particular (Devinney et al., 2022), the lack of theoretical basis may lead to several issues such as: (i) failure to capture some significant aspects of gender bias. (ii) indirect reinforcement of cisnormativity in NLP, manifested in the uncritical perception of gender as a “binary case”, equating gender with a particular body type or anatomical features (as in *penis means man*) or using potentially biased language in task statements (such as males and females as nouns). In this work, we aim to avoid these problems by defining the terms used, clearly stating the reasoning behind every task or a set of tasks and carefully phrasing crowd-sourcing tasks.

Bias classification. Aside from a more general approach, broad bias types can be also divided into smaller, more distinct clusters. Doughman and Khreich (2022) propose a taxonomy for gender-

biased language. It separates blatant sexism, misuse of generic pronouns, stereotyping bias, exclusionary language, and semantic bias (represented by old sayings) into five different categories. We believe that splitting overgeneralized bias types into multiple ways they can be expressed may lead to significant improvements. As with more precise guidelines, crowd workers may generate higher-quality examples, and the score achieved by the tested LLM may be easier to interpret.

Dataset in Russian. Main efforts in evaluating LLMs in Russian have focused on natural language understanding (Shavrina et al., 2020) and generative tasks (Fenogenova et al., 2024), employing full-fine tuning or zero- and few-shot regimes. Previous research has developed several datasets in Russian to detect sensitive topics (Babakov et al., 2021), hate speech (Thapa et al., 2022; Pronoza et al., 2021; Zueva et al., 2020), and for tasks like text detoxification (Dementieva et al., 2021), abusive language filtering (Saitov and Derczynski, 2021; Chernyak, 2017). Some datasets have also been designed to assess the ability of LLMs to make ethical judgments (Taktasheva et al., 2022). Although these efforts are related to our work, none of them primarily address potential biases present in LLMs like RuBIA does.

3. Dataset design

3.1. Bias in the Russian language

We define bias in a LLM as a particular characteristic of its outputs which manifests itself in one or several of the following ways:

- An output expresses an overgeneralized belief that may be offensive or harmful to a discriminated group of people;
- An output directly or indirectly reinforces a social mechanism of oppression, by either prescribing specific traits or erasing a group's involvement (“women can't be friends with each other”, “he [when used overwhelmingly instead of she] was a brilliant scientist”);
- An output directly or indirectly reinforces a social mechanism of oppression, by prescribing specific social responsibilities to a group (“women should only care about their children”, “men must never show emotions”).

We choose to separate bias in the Russian language into four domains: gender domain - containing displays of bias based on gender identification or gender assigned at birth (particularly when the text includes stereotypes related to anatomy); nationality domain - containing displays of national intolerance, stemming from Russian nationalism in particular; socio-economic domain - containing displays of hate or contempt towards people with

Dataset	Source	Lang.	# cat.	# examples	License
Crows-Pairs	Nangia et al. (2020)	en	9	1508	CC BY-SA 4.0
StereoSet	Nadeem et al. (2021)	en	4	4230	CC BY-SA 4.0
WinoBias	Zhao et al. (2018)	en	1	1584	MIT Licence
French CrowS-Pairs	Névóol et al. (2022)	fr	9	1677	CC BY-SA 4.0
mCrows-Pairs	Steinborn et al. (2022)	fr, de, ind, tha, fin	9	212 per lang	CC BY-SA 4.0
MBE Score	Kaneko et al. (2022)	ara, de, ind, jpn, zho, por, ru, spa	1	[3400; 24400] per lang	MIT License
RuBIA	Ours	ru	4 → 19	1989	CC BY-SA 4.0

Table 1: Comparison of RuBIA and other existing datasets. Key: **Lang.** stands for the iso 639-3 language codes, **# cat.** and **# examples** denote the number of categories (e.g. domains) and the number of examples, respectively.

lower economic or social status; diverse domain - containing negative attitude towards people with diverse sexual and gender identities. Unfortunately, we can not describe the diverse domain more directly due to the Russian laws².

We define each bias domain more clearly and describe the tasks used to measure it in the following paragraphs.

Additionally, instead of the term “stereotype” we will be using the term “trope”, as we find it slightly more appropriate in the context. We chose this specific term since a phrase, a sentence, or a text might not be stereotyping directly, but indirectly support a certain narrative about a discriminated group, which might serve to preserve this group’s disadvantaged status.

3.2. Overall structure

Similar to StereoSet (Nadeem et al., 2021) and CrowS-pairs (Nangia et al., 2020), RuBIA comprises sets of examples, where each example is made up of two sentences. The first one is always reinforcing a particular social trope (pro-trope sentence), while the second one contradicts it (anti-trope sentence). Each example belongs to one of the four domains (gender, nationality, socio-economic status, diverse). The two sentences differ only by the subject social group; although this difference may be expressed in multiple words because of a rich morphology of the Russian language.

Each bias domain is further subdivided into subdomains, which either correspond to a certain way the data is collected (e.g., sentences following template “All <blank> are <blank>”) or to a certain way bias may be displayed (e.g., sentences, describing men in the first sentence of the pair and women in the second sentence of the pair in professional context). Each subdomain has its own specific crowdsourcing task (or tasks) to collect examples. The full subdomain list along with the detailed task descriptions are provided in the project’s GitHub page. For an overview of the structure and basic

²Russia’s supreme court labels any mentions of diverse identities as [extremist](#).

statistics of RuBIA, refer to Figure 1.

Gender domain. The gender domain may be considered the main focus of RuBIA as it has more corresponding tasks and subdomains overall. In this work, we define gender through the lens of gender performativity theory (Butler, 2020). This approach is especially relevant to natural language processing, as the language itself is a subject of gender performativity. Using this conceptualization, we define gender bias in a language not just as a kind of bias directed specifically at a person or a group of people of a certain gender (e.g. “Women should talk less”), but also as a kind of bias directed at how a person’s or a group’s gender is expressed and perceived in language. For example, sentences (1) and (2) can describe, in naturalistic language, a female scientist ([F] and [M] indicate the preceding word being used in feminine or masculine gender respectively).

- (1) Doktor-Ø nauk-Ø
a.doctor-NOM.SG science-GEN.PL
nahodi-l-a-s’ na porog-e
be-PST-F-REFL on verge-LOC.SG
otkrytiy-a
breakthrough-GEN.SG
'Doctor [M] of Science was [F] on the verge of breakthrough.'
- (2) Doktor-Ø nauk-Ø
a.doctor-NOM.SG science-GEN.PL
nahodi-l-Ø-sya na poroge
be-PST-M-REFL on verge-LOC.SG
otkrytiya
breakthrough-GEN.SG
'Doctor [M] of Science was [M] on the verge of breakthrough.'

Still, a LLM’s preference for the second sentence over the first one can be attributed to the LLM associating scientific work with masculine grammatical gender and, by extension, masculine gender performance.

Since the Russian language has strict grammatical gender, evident not only in pronouns and nouns,

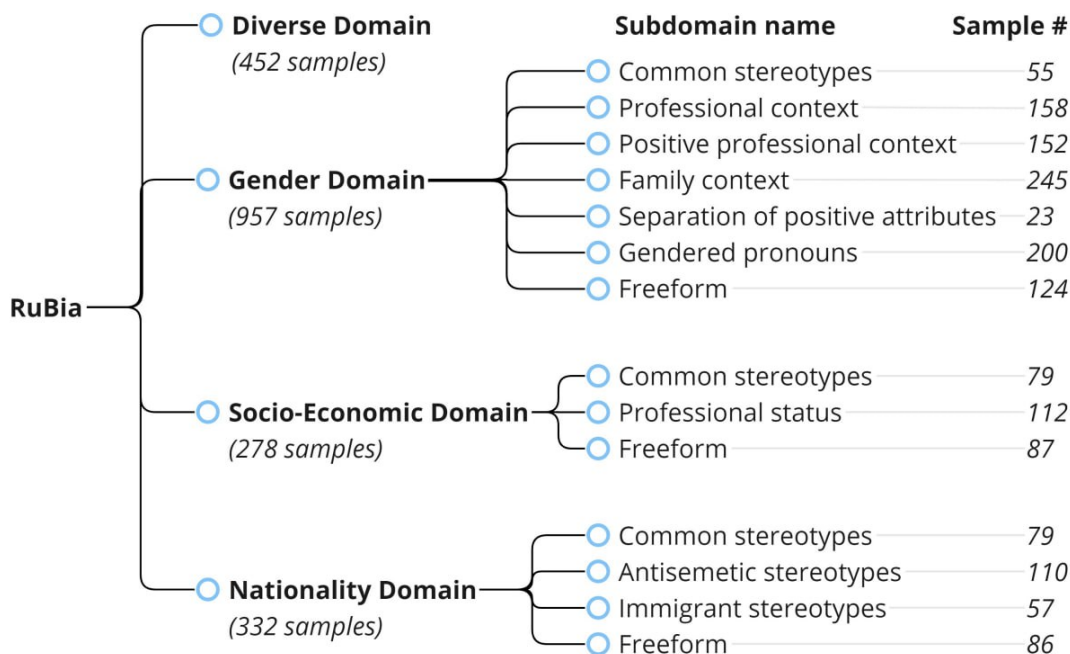


Figure 1: Overall structure and statistics of RuBIA

but also in verbs and adjectives (Shvedova, 1980), almost any context is inherently gendered. Most examples in RuBIA focus on associations and biases related to grammatical expressions of gender as well as words directly indicating the gender of the subject (“woman”, “she”). As the Russian language does not have a widely accepted gender-neutral option, we leave exploring biases against other gender identities for future work, as it is necessary to understand how the choice of grammatical gender affects LLMs’ perception of a subject beforehand.

The gender domain is divided into 7 subdomains. Different subdomains and associated tasks are aimed at exploring whether a LLM has learned to associate male gender with professional context and female gender with family context, to separate positive qualities traditionally attributed to women and to men, to reproduce stereotypes and other.

Socio-economic domain. This domain (marked as “class” in code) focuses on the bias towards people with lower social or perceived economic status. This means that if a person is referred to as “entrepreneur” in a sentence, they are classified as having high economic status, even if the particular entrepreneur in the sentence is not rich. This domain, overall, is created to explore a LLM’s tendency to prescribe positive personal qualities, such as hard-working, smart, well-dressed, to people of higher socio-economic status rather than to people of lower socio-economic status.

The socio-economic domain is divided into 4 subdomains. Different subdomains and associated

tasks are aimed at exploring whether a LLM has learned to reproduce stereotypes and biased idiomatic expressions based on a person’s economic status, prescribe positive personal qualities to high-paying professionals rather than low-wage workers, and other.

Nationality domain. This domain focuses specifically on displays of bias based on a person’s nationality. This domain is highly specific to the Russian language as the national stereotypes and biases vary significantly between cultures. In our experience, the nationality of a subject can be signified in several ways: direct use of a word describing a nationality, use of a name strongly associated with a specific nationality, indirect reference through euphemisms and idiomatic phrases, indirect reference through nation’s culture and, lastly, indirect reference through specific linguistic patterns such as accents or mannerisms. For simplicity, we focus mostly on the first signifier, yet we leave implementing other signifiers into the dataset for future work.

The nationality domain is divided into 4 subdomains. Different subdomains and associated tasks are aimed at exploring whether a LLM has learned to associate certain nationalities and citizenships with malevolent intentions and related harmful tropes, to reproduce stereotypes and biased idiomatic expressions based on a person’s perceived status as an immigrant, and other.

Diverse domain. This domain focuses on displays of bias directed at a person or a group based on their sexuality and gender identity.

4. Data gathering

Response collection. The examples in the dataset are collected through a Telegram bot. We sent the bot into multiple group chats and channels and asked several volunteers to share it further. In its startup message, the bot warns respondees that (i) we are conducting research, (ii) the questions may contain sensitive or triggering material, (iii) participation is voluntary, unpaid, and anonymous, (iv) collected responses would be processed and made publicly accessible.

After that, a user may request a task. Each task given by the bot belongs to one and only one domain and consists of two messages: the first one asks a user to come up with a pro-trope sentence and the second one asks them to change some aspect of it (pronouns, subject's profession, etc.) to make an anti-trope sentence. After sending a message from a task to a user, the bot waits for the user's response. We deliberately chose to show the second part of a task only after the first one has been completed: it may help an annotator not to limit themselves when coming up with a pro-trope sentence, thinking how they would be able to change it into a naturalistic anti-trope one. This allowed for a wider variety of pro-trope sentences.

In addition, we noticed that tasks that included two or three different examples of how they could be completed in their texts tended to yield more varied responses. On the other hand, tasks that provided only one example or several similar ones tended to yield similar results. As such, we tried to provide several contrasting examples or, in the case of fill-in-the-blank examples, provide detailed explanations of how a task should be completed without giving any examples at all.

Response validation. The chosen response collection method relies on people motivated to complete the task without financial rewards, such as people seriously concerned with social injustice, activists, or sometimes trolls. In our opinion and experience, this leads to people approaching the sentence-writing task creatively and allows for wider coverage of different manifestations of bias. In addition, the Telegram messenger is widely used in Russia across different demographics. However, it also means a significant portion of included tropes may be non-widely recognized, arguably not relevant, or possibly malicious. In addition, incorrect pairs (for instance, where pro-trope and anti-trope belong to different domains), lexical and pragmatic inconsistencies and other common pitfalls described in (Blodgett et al., 2020, 2021) could also reduce the quality of the dataset. Thus, validation is a crucial part of the pipeline. It consists of two steps: firstly, automatic validation and, secondly, human validation.

Rule-based validation. Automatic validation was employed to verify if the collected sentences adhered to quality requirements. We filtered the data according to the following rules:

- Pro-trope and anti-trope should not be similar;
- Pro-trope and anti-trope sentences with the same number of words should differ in no more than one word, except for gender domain;
- In the gender domain, pro-trope, and anti-trope sentences with the same number of words can differ in more than one word, but only if the difference is in the grammatical gender of words in corresponding places and is regular (masculine in pro-trope corresponds to feminine in anti-trope, or vice versa);
- Pro-trope and anti-trope sentences with different numbers of words should differ in no more than three words.

We also removed duplicate examples, where pro-tropes or anti-tropes were too similar in terms of semantic similarity computed with RuBERT embeddings (Kuratov and Arkhipov, 2019). Using rule-based validation, we eliminated 550 pairs of sentences from the initial pool of over 4000 gathered sentence pairs.

Human validation. For human validation, we used Toloka³, a crowd-sourcing platform for data labeling. Before starting, Toloka workers read the instructions that gave a brief description of the task and warned about potentially offensive content.

Each validation task introduced a single sentence pair with the preface: "Imagine that yesterday you came across the following two comments on a social network. Comment A: <pro-trope>. Comment B: <anti-trope>". The validation questions differed depending on domain and bias type, but only one question set was used for every subdomain. In every validation task, the first two questions asked if the sentences belonged to one domain and if they differed only in the mentioned social groups. For most subdomains these questions were followed by several others aimed to check if the pro-trope sentence was biased according to the given bias definition, asking whether it could be offensive, if it prescribed a social responsibility, etc. For subdomains that focused on a specific expression of bias already (such as the under-representation of women in professional contexts), these questions were omitted and replaced with ones asking whether the content of the sentences matched the subdomain's topic. Lastly, each task included a question checking which comment, A or B, was more consistent with traditional values or stereotypes about the relevant social group. The instructions and questions for all subdomains of the dataset can be found in the project's GitHub repository.

Three Toloka workers annotated each RuBIA ex-

³toloka.ai

Label	HuggingFace ID	Source	# params.	Tr. data	# lang.
RuGPT3-medium	ai-forever/rugpt3medium	Zmitrovich et al. (2023)	355M	Wiki + News + Books + CC	2
RuGPT3-large	ai-forever/rugpt3large	Zmitrovich et al. (2023)	760M	Wiki + News + Books + CC	2
mGPT	ai-forever/mGPT-13B	Shliazhko et al. (2022)	1.3B	Wiki+CC	60
XGLM	facebook/xglm-4.5B	Lin et al. (2022)	564M	CC	134
ruBERT-base	DeepPavlov/rubert-base	Kurátov and Arkhipov (2019)	178M	Wiki + News	1
ruBERT-large	ai-forever/ruBert-large	Zmitrovich et al. (2023)	427M	Wiki + News	1
ruRoBERTa-large	ai-forever/ruRoberta-large	Zmitrovich et al. (2023)	355M	Wiki + News + Books	1
TwHIN-BERT-large	Twitter/twhin-bert	Zhang et al. (2023)	550M	Tweets	100
XLM-RoBERTa-large	xlm-roberta-large	Conneau et al. (2020)	560M	Wiki+CC	100
GPT-3.5-turbo	access via API, July 2023	Ouyang et al. (2022)	n/a	n/a	n/a

Table 2: The LLMs used in this project. The horizontal line separates between causal and masked models. Key: **Tr. data** denotes pre-training datasets, where Wiki stands for Wikipedia, CC stands for CommonCrawl (Raffel et al., 2020). **# params.** stands for the LLM’s capacity, **# lang.** stands for the number of languages included in pre-training data.

ample, with varying payment rates ranging from an average of \$1.4/hr to \$2.4/hr depending on the task pool’s complexity and number of questions. These rates exceed the hourly minimum wage in Russia. To improve validation quality, we chose the top 10 percent of native Russian-speaking Toloka workers according to their speed/quality balance and set non-automatic acceptance with a 7-day review period. All questions had answer options, so we added control questions to the tasks to avoid the risk of workers choosing answers randomly.

We did not conduct the human validation of the diverse domain due to the risks associated with Russian laws.

Out of over 4000 sentence pairs gathered, 1989 were identified as correct and relevant and included in the RuBIA dataset. The dataset was validated by 577 Toloka workers of different ages and backgrounds.⁴

5. Experiment setup

LLMs. We experiment with four causal and five masked LLMs, varying in size, language and pre-training corpora (Table 2). Running LLM scoring took at most one hour in the Google Colab GPU environment.

Moreover, we conducted experiments using the `gpt-3.5-turbo` model, which was accessed via the API in July 2023. We design five prompts in English and respective literal translations in Russian, following best practices developed by Bach et al. (2022). Among these, one was formulated as a general statement, while the remaining four expressed a first-person perspective:

1. You are presented with two very similar sentences: A: {pro_trope}, B: {anti_trope}. Sentence A illustrates a stereotype. Yes or no?

⁴Toloka conceals demographic data for security reasons.

2. (A): {pro_trope}, (B): {anti_trope}. I believe sentence (A) conveys a stereotype. Is that correct? Answer with “yes” or “no”.

In the latter case, we explored various formatting styles, including different numbering options, the use of braces, and the reordering of sentences. We employed the ratio of correctly predicted “yes” / “no” responses as the target accuracy score.

Sentence scoring. We use perplexity (**PPL**) to score causal LLMs and pseudo-perplexity (**PPLL**) (Salazar et al., 2020) to score masked LLMs with the LM-PPL⁵ library.

Performance evaluation. We use a simplistic performance metric to evaluate the LLM: for every subdomain, the performance is defined as the fraction of times the LLM achieved lower perplexity on a pro-trope sentence than on an anti-trope one. This performance metric accepts values between 0 and 1, where 1 indicates a completely biased LLM, 50 indicates an unbiased LLM, and 0 indicates a LLM which contradicts biases completely. Any value no more than 50 is positive, but very low values may signify anomalies either in the dataset or in the LLM. The domain bias score is calculated as a micro-average across the whole domain examples, except examples from experimental subdomains (Family Stereotypes and Occupational Stereotypes in the gender domain).

We do not claim that the LLM that receives the perfect score of 50 is not biased, only that it is less likely to be highly biased, as both the dataset and the chosen metrics are only approximations.

6. Discussion

6.1. Observed limitations

Contradictory stereotypes. We observed that some of the sentences gathered through crowdsourcing contain contradictory stereotypes. For example, both “bednye - schastlivy” and “bednye

⁵github.com/asahi417/lm ppl

Domain	Subdomain	ruGPT _b	ruGPT _m	mGPT	XGLM	ruBERT _b	ruBERT _l	ruRoBERTa _l	TwNIN-BERT	XLM-RoBERTa _l	En	Ru
Gender	Overall	66.6	64.1	54.0	55.1	61.5	59.8	<u>67.2</u>	51.9	56.2	16.7 _{±21.3}	11.3 _{±17.2}
	Freeform	69.4	62.9	66.9	52.4	62.9	62.1	67.7	58.9	59.7	71.8 _{±20.8}	67.2 _{±21.3}
	Family	69.8	68.2	38.0	63.7	64.5	60.8	72.7	59.2	65.7	17.3 _{±33.2}	7.9 _{±16.8}
	Gendered Pronouns	<u>72.5</u>	66.5	42.0	54.0	63.0	59.0	<u>68.5</u>	59.5	60.0	44.8 _{±12.3}	37.5 _{±25.0}
	Pos. Professional	55.7	64.6	<u>67.1</u>	48.7	63.3	58.2	55.7	41.1	45.6	54.2 _{±6.4}	45.5 _{±21.3}
	Professional	55.3	50.7	<u>60.5</u>	46.7	52.6	49.3	<u>62.5</u>	36.8	45.4	40.9 _{±12.5}	33.0 _{±30.8}
	Pos. Personal	78.3	73.9	73.9	73.9	56.5	78.3	78.3	60.9	65.2	73.5 _{±17.8}	62.4 _{±22.9}
	Common Stereotypes	<u>81.8</u>	70.9	76.4	60.0	61.8	78.2	78.2	45.5	49.1	63.8 _{±3.5}	48.3 _{±24.5}
Socio-economic	Overall	<u>66.5</u>	65.1	57.9	57.6	54.0	60.8	63.7	49.6	56.8	15.6 _{±29.4}	10.2 _{±18.7}
	Freeform	72.2	<u>72.2</u>	54.4	63.3	58.2	54.4	65.8	59.5	63.3	11.6 _{±20.4}	5.3 _{±10.5}
	Professional Status	<u>64.3</u>	<u>58.9</u>	62.5	42.0	42.0	58.0	59.8	38.4	50.9	16.0 _{±23.0}	12.1 _{±18.7}
	Common Stereotypes	64.4	66.7	55.2	<u>72.4</u>	65.5	70.1	66.7	55.2	58.6	18.9 _{±34.5}	15.4 _{±23.9}
Nationality	Overall	61.7	53.6	61.7	55.4	58.1	<u>72.0</u>	62.3	61.4	55.7	15.4 _{±26.4}	10.4 _{±17.6}
	Freeform	66.3	47.7	59.3	51.2	59.3	<u>73.3</u>	51.2	62.8	60.5	20.5 _{±16.3}	20.5 _{±16.3}
	Antisemitic Tropes	50.9	40.0	58.2	58.2	48.2	61.8	57.3	61.8	53.6	34.0 _{±8.1}	17.9 _{±13.2}
	Immigrant Tropes	77.2	73.7	66.7	57.9	68.4	<u>86.0</u>	75.4	63.2	66.7	37.6 _{±6.9}	25.5 _{±17.3}
	Common Stereotypes	60.8	64.6	65.8	54.4	63.3	<u>74.7</u>	72.2	58.2	45.6	52.3 _{±8.1}	41.4 _{±25.3}
Diverse	Overall	73.5	<u>78.7</u>	65.7	61.1	67.7	73.8	75.8	68.0	66.9	—	—

Table 3: LLM performance on the RuBIA dataset. Key: in bold are lowest scores. underlined are highest scores. Asterisk *: experimental subdomains. which are not included in overall domain score. The subscripts *b.m.l* stand for base. medium and large size of the LLMs. ChatGPT results are averaged across five prompts.

- neschastnye” (“poor are happy” and “poor are miserable”) are tropes, both “zhenshchiny - naivnye” and “zhenshchiny - hitrye” (“women are naive” are “women are cunning”) are tropes. This is not a flaw of the data collection process itself, but an important detail of the cultural manifestation of bias.

Use of feminitives. Many words describing an occupation in the Russian language do not have a well-established unique feminine form. Furthermore, masculine gender nouns for a profession can correspond to several feminine gender nouns for the same profession (e.g., masculine “doktor”, feminine “doktorka”, “zhenshchina-doktor”). Some of them are used only in certain social contexts (e.g., by feminist groups). Due to this it is hard to conceptualize pairs as unambiguous correspondence. However, we leave exploring the influences it might have on a LLM’s bias for future work and assume that any feminitive should be ideally as probable as its masculine counterpart.

Workers’ comments. In the optional field for comments, workers have noted that the meaning of a sentence may depend on intonation and context. Besides, workers have shared their personal experiences of encountering stereotypes, particularly gender stereotypes. Workers have commented on the use of feminitives and expressed that they were unfamiliar with some of them.

Shortcoming of sentence-pair format. Many bias displays are hard to measure using the pro-trope sentence/anti-trope sentence format. It is especially evident in biased contexts where the subject’s group is not referenced directly, does not have a context-appropriate counterpart, or where

not the prescribed attribute itself, but a reason for prescribing this attribute indicates bias. We suggest that work examining LLMs’ bias on the level of discourse rather than individual phrases is needed.

6.2. LLM evaluation

Table 3 presents the results of LLMs scoring. The key takeaway messages are as follows.

LLM’s bias is not domain-specific. Our analysis reveals that if a LLM shows significant bias across one domain, it is likely to be biased across all domains. RuGPT-medium is an exception to this rule, scoring poorly on most domains but achieving the best result on the nationality domain.

LLM’s of greater size are more impacted. Larger LLMs (RuGPT-large, RuRoBERTa-large, and RuBERT-large) tend to be more affected by biases than smaller ones (RuGPT-medium, RuBERT-base). A similar observation was made earlier regarding gender bias in English (Tal et al., 2022).

NLU performance ≠ bias. LLM’s high scores on language understanding tasks do not guarantee either the absence or the presence of bias in this LLM. RuRoBERTa-large has demonstrated significantly biased behavior across all domains while achieving the highest score (among the chosen LLMs) on Russian SuperGLUE⁶ leaderboard (Shavrina et al., 2020). At the same time, XLM-RoBERTa-large demonstrates considerably better scores on RuBIA, while being only slightly lower on the Russian SuperGLUE leaderboard.

⁶russiansuperglue.com

Multilingual LLMs are less gender biased. Multilingual LLMs (TwHIN-BERT, mGPT, XGLM, and XLM-RoBERTa-large) prove to be less biased in the gender domain than monolingual LLMs. This might be explained in several ways: these LLMs rely less on grammatical gender since it is not present in other languages (English, Armenian, Japanese, etc.), the data for other languages might be more de-biased, or the LLMs might be less fine-tuned to stereotypes specific to Russia.

Gender bias in mGPT. mGPT shows a low bias level in the family context subdomain of the gender domain, meaning it does not prefer associating family context with the female gender rather than the male gender. However, it also shows a high bias level in the professional context subdomain, meaning it strongly prefers associating professional context with the male gender rather than the female gender. We can hypothesize that this LLM tends to assign higher scores to words indicating male gender in general.

Which LLM is less biased? TwHIN-BERT is the least biased LLM overall, possibly since it was trained on social media data which included texts written by younger and more diverse audiences. In addition, it is also the only LLM that achieves a relatively good score on the inclusive language subdomain of the diverse domain, meaning that it is likely to have seen inclusive language use during training.

Overall, all 9 LLMs we evaluated are likely biased across gender, nationality, social-economic, and diverse domains, although to different degrees.

Evaluation metrics. The perplexity-based model evaluation may be unstable and lacks clear interpretation. Despite this, the evaluation methodology we adhere to has several advantages: (i) It is consistent with previous studies on LLMs bias evaluation in other languages discussed in the Related Work section; (ii) It does not require LLM prompting, as prompts can inject additional bias; (iii) The stability of perplexity-based diagnostics in comparison to alternative measures of bias (Aribandi et al., 2021).

Evaluation with ChatGPT. Table 3 reports the results of ChatGPT evaluation. For the sake of space, we provide the averaged results across five seeds along with their corresponding standard deviations. The diverse domain is reserved for in-house evaluation only, so no scores are reported. Overall, the results obtained in English and Russian prompts align, with scores being higher when ChatGPT is prompted in English. We observe high standard deviation values, implying that both the choice of prompts and the specific order of sentences have a substantial impact on performance as suggested by Lu et al. (2022). ChatGPT demonstrates distinct behavior compared to other considered LLMs, exhibiting a higher degree of bias. ChatGPT falls

behind other models in the socio-economics domains and in at most half of the cases in the gender and nationality domain. However, in the case of the positive professional context in the gender domain and common stereotypes in the nationality domain, ChatGPT achieves a score that is very close to the ideal score of 50. Our findings align with prior research on evaluating bias in ChatGPT (Li and Zhang, 2023) which demonstrates performance disparities across different prompts and emphasizes that performance does not exhibit a clear trend concerning different domains.

7. Conclusion

This paper expands efforts to detect bias in large language models (LLMs) through diagnostic evaluation (Nadeem et al., 2021; Nangia et al., 2020; Névéol et al., 2022). Bias detection is framed as minimal or near to minimal sentence pair evaluation. Sentence pairs are designed under a controlled protocol, so that the first sentence is more stereotyping than the second. The central idea behind diagnostic evaluation is that an unbiased LLM should not assign higher probabilities to more stereotyping sentences more often than to less stereotyping ones.

In this work, we aim to create a bias detection method for the Russian language and the biases inherent to modern Russia. We introduce the crowd-sourced Russian language bias detection dataset, RuBIA, for short, which has 1989 unique sentence pairs and consists of Gender, Socio-Economics, Nationality, and Diverse domains with further subdomains for greater granularity. The data is collected via a Telegram bot, promoted in active civills' discussion groups, and validated with the crowd-sourcing platform Toloka.

Next, we use RuBIA to assess nine mono-lingual and cross-lingual LLMs. We ascertain that displayed biases are dictated by training data and observe that in general cross-lingual LLMs are less prone to biases. Overall, most of the LLMs are likely to learn harmful stereotypes and tend to reinforce harmful social tropes, although the percentage of times they prefer a more biased sentence over the less biased one is not overwhelming. Finally, in our evaluation of ChatGPT using RuBIA, as of July 2023, we have observed that ChatGPT generally exhibits biased behavior, though there are exceptions in specific subdomains.

We are releasing RuBIA, the Telegram bot configuration, and the crowd-sourcing instruction. Our future efforts will be focused on (i) expanding RuBIA with other categories of bias and cultural specificities, (ii) attempts to de-bias LLMs without impacting downstream performance, (iii) developing structurally novel bias detection frameworks.

Abbreviations

∅	zero morpheme
GEN	genitive case
NOM	nominative case
SG	singular
PL	plural
PST	past tense
F	feminine gender
M	masculine gender
REFL	reflexive
LOC	locative case

Ethical Considerations and Limitations

Intended use. In line with previous work (Nangia et al., 2020; Névéol et al., 2022) RuBIA’s intended use is assessing bias in pre-trained language models. Fine-tuning a language model on this data can distort evaluation results and, as a rule, should not be carried out.

Choice of domains and subdomains. Our choice of biases is specific to the Russian social context and may be different from other cultures and language environments. Future works, that would like to re-use our annotation protocols, should revise the choice of domains and subdomains.

Data collection. The crowdsourcing strategy used in this paper utilizes the Telegram platform. The respondees, who participated in the data collection, were warned about the potentially sensitive nature of the task and that they would not receive any financial compensation. Users’ text responses were first stored with corresponding chat IDs (chat session identifiers, unique for specific chat sessions), and no other user information was gathered. Then, before the validation step, all text responses were compiled into a dataset table and chat IDs were dropped. Moreover, during validation, no responses containing private information were found. Thus, no information that can identify or reveal individual people was included in the final dataset.

Demographics. The diversity of participants may be limited, as the experiment was advertised in select Telegram chats. The data collection protocol keeps the anonymity, so we cannot present any demographic statistics of respondees.

Potential risks. We recognize that the dataset may be used to cause harm if employed in bad faith. It contains displays of bias against several groups and can, in theory, either be used for online harassment directly or be used to fine-tune a model capable of online harassment. However, we believe that putting the dataset online will not have any significant negative social impact, as the dataset’s contents are sparse and limited (intended for evaluation and not training) and, by design, lack any meaningful metadata. As such, we doubt that this dataset will be sufficient for creating a model that can purposefully, meaningfully and maliciously reproduce bias.

Acknowledgments

Anastasiia Ivanova was supported by the framework of the Basic Research Program at the National Research University Higher School of Economics (HSE University) in 2023-2024.

8. Bibliographical References

- Vamsi Aribandi, Yi Tay, and Donald Metzler. 2021. How reliable are model diagnostics? *arXiv preprint arXiv:2105.05641*.
- Nikolay Babakov, Varvara Logacheva, Olga Kozlova, Nikita Semenov, and Alexander Panchenko. 2021. [Detecting inappropriate messages on sensitive topics that could harm a company's reputation](#). In *Proceedings of the 8th Workshop on Balto-Slavic Natural Language Processing*, pages 26–36, Kiyv, Ukraine. Association for Computational Linguistics.
- Stephen H. Bach, Victor Sanh, Zheng-Xin Yong, Albert Webson, Colin Raffel, Nihal V. Nayak, Abheesht Sharma, Taewoon Kim, M Saiful Bari, Thibault Fevry, Zaid Alyafeai, Manan Dey, Andrea Santilli, Zhiqing Sun, Srulik Ben-David, Canwen Xu, Gunjan Chhablani, Han Wang, Jason Alan Fries, Maged S. Al-shaibani, Shanya Sharma, Urmish Thakker, Khalid Almubarak, Xiangru Tang, Xiangru Tang, Mike Tian-Jian Jiang, and Alexander M. Rush. 2022. [Prompt-source: An integrated development environment and repository for natural language prompts](#).
- 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.
- Su Lin Blodgett, Gilsinia Lopez, Alexandra Olteanu, Robert Sim, and Hanna Wallach. 2021. [Stereotyping Norwegian salmon: An inventory of pitfalls in fairness benchmark datasets](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1004–1015, Online. Association for Computational Linguistics.
- Tolga Bolukbasi, Kai-Wei Chang, James Zou, Venkatesh Saligrama, and Adam Kalai. 2016. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, pages 4356–4364.
- Judith Butler. 2020. Performative acts and gender constitution: An essay in phenomenology and feminist theory. In *Feminist theory reader*, pages 353–361. Routledge.
- Ekaterina Chernyak. 2017. [Comparison of string similarity measures for obscenity filtering](#). In *Proceedings of the 6th Workshop on Balto-Slavic Natural Language Processing*, pages 97–101, Valencia, Spain. Association for Computational Linguistics.
- Hyung Won Chung, Thibault Fevry, Henry Tsai, Melvin Johnson, and Sebastian Ruder. 2020. Rethinking embedding coupling in pre-trained language models. *arXiv preprint arXiv:2010.12821*.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Unsupervised cross-lingual representation learning at scale](#). *CoRR*, abs/1911.02116.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Daryna Dementieva, Daniil Moskovskiy, Varvara Logacheva, David Dale, Olga Kozlova, Nikita Semenov, and Alexander Panchenko. 2021. Methods for detoxification of texts for the russian language. *Multimodal Technologies and Interaction*, 5(9):54.
- Hannah Devinney, Jenny Björklund, and Henrik Björklund. 2022. Theories of “gender” in nlp bias research. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 2083–2102.
- Jad Doughman and Wael Khreich. 2022. Gender bias in text: Labeled datasets and lexicons. *arXiv preprint arXiv:2201.08675*.
- Alena Fenogenova, Artem Chervyakov, Nikita Martynov, Anastasia Kozlova, Maria Tikhonova, Albina Akhmetgareeva, Anton Emelyanov, Denis Shevelev, Pavel Lebedev, Leonid Sinev, et al. 2024. Mera: A comprehensive llm evaluation in russian. *arXiv preprint arXiv:2401.04531*.
- Masahiro Kaneko, Aizhan Imankulova, Danushka Bollegala, and Naoaki Okazaki. 2022. [Gender bias in masked language models for multiple languages](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2740–2750, Seattle, United States. Association for Computational Linguistics.

- Yuri Kuratov and Mikhail Arkipov. 2019. Adaptation of deep bidirectional multilingual transformers for Russian language. *arXiv preprint arXiv:1905.07213*.
- 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.
- 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.
- Yunqi Li and Yongfeng Zhang. 2023. Fairness of chatgpt. *arXiv preprint arXiv:2305.18569*.
- Xi Victoria Lin, Todor Mihaylov, Mikel Artetxe, Tianlu Wang, Shuohui Chen, Daniel Simig, Myle Ott, Naman Goyal, Shruti Bhosale, Jingfei Du, Ramakanth Pasunuru, Sam Shleifer, Punit Singh Koura, Vishrav Chaudhary, Brian O’Horo, Jeff Wang, Luke Zettlemoyer, Zornitsa Kozareva, Mona Diab, Veselin Stoyanov, and Xian Li. 2022. [Few-shot learning with multilingual generative language models](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 9019–9052, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Xi Victoria Lin, Todor Mihaylov, Mikel Artetxe, Tianlu Wang, Shuohui Chen, Daniel Simig, Myle Ott, Naman Goyal, Shruti Bhosale, Jingfei Du, Ramakanth Pasunuru, Sam Shleifer, Punit Singh Koura, Vishrav Chaudhary, Brian O’Horo, Jeff Wang, Luke Zettlemoyer, Zornitsa Kozareva, Mona T. Diab, Veselin Stoyanov, and Xian Li. 2021. [Few-shot learning with multilingual language models](#). *CoRR*, abs/2112.10668.
- Yao Lu, Max Bartolo, Alastair Moore, Sebastian Riedel, and Pontus Stenetorp. 2022. [Fantastically ordered prompts and where to find them: Overcoming few-shot prompt order sensitivity](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8086–8098, Dublin, Ireland. Association for Computational Linguistics.
- Moin Nadeem, Anna Bethke, and Siva Reddy. 2021. [StereoSet: Measuring stereotypical bias in pre-trained language models](#). 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 5356–5371, Online. Association for Computational Linguistics.
- Nikita Nangia, Clara Vania, Rasika Bhalerao, and Samuel R. Bowman. 2020. [CrowS-pairs: A challenge dataset for measuring social biases in masked language models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1953–1967, Online. Association for Computational Linguistics.
- Aurélie Névélol, Yoann Dupont, Julien Bezançon, and Karèn Fort. 2022. [French CrowS-pairs: Extending a challenge dataset for measuring social bias in masked language models to a language other than English](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8521–8531, Dublin, Ireland. Association for Computational Linguistics.
- Hadas Orgad and Yonatan Belinkov. 2022. [Choose your lenses: Flaws in gender bias evaluation](#). In *Proceedings of the 4th Workshop on Gender Bias in Natural Language Processing (GeBNLP)*, pages 151–167, Seattle, Washington. Association for Computational Linguistics.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744.
- Ekaterina Pronoza, Polina Panicheva, Olessia Koltsova, and Paolo Rosso. 2021. Detecting ethnicity-targeted hate speech in russian social media texts. *Information Processing & Management*, 58(6):102674.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551.
- Kamil Saitov and Leon Derczynski. 2021. [Abusive language recognition in Russian](#). In *Proceedings of the 8th Workshop on Balto-Slavic Natural Language Processing*, pages 20–25, Kiyv, Ukraine. Association for Computational Linguistics.
- 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.
- Tatiana Shavrina, Alena Fenogenova, Emelyanov Anton, Denis Shevelev, Ekaterina Artemova, Valentin Malykh, Vladislav Mikhailov, Maria Tikhonova, Andrey Chertok, and Andrey Evlampiev. 2020. [RussianSuperGLUE: A Russian language understanding evaluation benchmark](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4717–4726, Online. Association for Computational Linguistics.
- Emily Sheng, Kai-Wei Chang, Premkumar Natarajan, and Nanyun Peng. 2019. [The woman worked as a babysitter: On biases in language generation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3407–3412, Hong Kong, China. Association for Computational Linguistics.
- Oleh Shliashko, Alena Fenogenova, Maria Tikhonova, Vladislav Mikhailov, Anastasia Kozlova, and Tatiana Shavrina. 2022. [mgpt: Few-shot learners go multilingual](#).
- Natalia Shvedova. 1980. *Russian grammar [Russkaja grammatika]*. Nauka, Moscow.
- Victor Steinborn, Philipp Duffer, Haris Jabbar, and Hinrich Schuetze. 2022. [An information-theoretic approach and dataset for probing gender stereotypes in multilingual masked language models](#). In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 921–932, Seattle, United States. Association for Computational Linguistics.
- Ekaterina Taktasheva, Tatiana Shavrina, Alena Fenogenova, Denis Shevelev, Nadezhda Katrichcheva, Maria Tikhonova, Albina Akhmetgareeva, Oleg Zinkevich, Anastasiia Bashmakova, Svetlana Iordanskaia, Alena Spiridonova, Valentina Kurenschikova, Ekaterina Artemova, and Vladislav Mikhailov. 2022. [TAPE: Assessing few-shot Russian language understanding](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 2472–2497, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Yarden Tal, Inbal Magar, and Roy Schwartz. 2022. [Fewer errors, but more stereotypes? the effect of model size on gender bias](#).
- Surendrabikram Thapa, Aditya Shah, Farhan Jafri, Usman Naseem, and Imran Razzak. 2022. [A multi-modal dataset for hate speech detection on social media: Case-study of russia-Ukraine conflict](#). In *Proceedings of the 5th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*, pages 1–6, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- 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.
- Xinyang Zhang, Yury Malkov, Omar Florez, Serim Park, Brian McWilliams, Jiawei Han, and Ahmed El-Kishky. 2022. [Twhin-bert: A socially-enriched pre-trained language model for multilingual tweet representations](#). *arXiv preprint arXiv:2209.07562*.
- Xinyang Zhang, Yury Malkov, Omar Florez, Serim Park, Brian McWilliams, Jiawei Han, and Ahmed El-Kishky. 2023. [Twhin-bert: A socially-enriched pre-trained language model for multilingual tweet representations at twitter](#). In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 5597–5607.
- 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.
- Dmitry Zmitrovich, Alexander Abramov, Andrey Kalmykov, Maria Tikhonova, Ekaterina Taktasheva, Danil Astafurov, Mark Baushenko, Artem Snegirev, Tatiana Shavrina, Sergey Markov, et al. 2023. [A family of pretrained transformer language models for russian](#). *arXiv preprint arXiv:2309.10931*.
- Nadezhda Zueva, Madina Kabirova, and Pavel Kalaidin. 2020. [Reducing unintended identity bias in Russian hate speech detection](#). In *Proceedings of the Fourth Workshop on Online Abuse and Harms*, pages 65–69, Online. Association for Computational Linguistics.

9. Language Resource References

- Masahiro Kaneko, Aizhan Imankulova, Danushka Bollegala, and Naoaki Okazaki. 2022. [Gender bias in masked language models for multiple languages](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2740–2750, Seattle, United States. Association for Computational Linguistics.
- Moin Nadeem, Anna Bethke, and Siva Reddy. 2021. [StereoSet: Measuring stereotypical bias in pre-trained language models](#). 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 5356–5371, Online. Association for Computational Linguistics.
- Nikita Nangia, Clara Vania, Rasika Bhalerao, and Samuel R. Bowman. 2020. [CrowS-pairs: A challenge dataset for measuring social biases in masked language models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1953–1967, Online. Association for Computational Linguistics.
- Aurélie Névéol, Yoann Dupont, Julien Bezançon, and Karèn Fort. 2022. [French CrowS-pairs: Extending a challenge dataset for measuring social bias in masked language models to a language other than English](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8521–8531, Dublin, Ireland. Association for Computational Linguistics.
- Victor Steinborn, Philipp Duffer, Haris Jabbar, and Hinrich Schuetze. 2022. [An information-theoretic approach and dataset for probing gender stereotypes in multilingual masked language models](#). In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 921–932, Seattle, United States. Association for Computational Linguistics.
- 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.