

# Tokenization is Sensitive to Language Variation

Anna Wegmann<sup>1</sup>, Dong Nguyen<sup>1</sup> and David Jurgens<sup>2</sup>

<sup>1</sup>Utrecht University, Utrecht, the Netherlands

<sup>2</sup>University of Michigan, Ann Arbor, the United States

{a.m.wegmann, d.p.nguyen}@uu.nl, jurgens@umich.edu

## Abstract

Variation in language is ubiquitous and often systematically linked to regional, social, and contextual factors. Tokenizers split texts into smaller units and might behave differently for less common linguistic forms. This might affect downstream LLM performance differently on two types of tasks: Tasks where the model should be robust to language variation (e.g., for semantic tasks like NLI, labels do not depend on whether a text uses British or American spelling) and tasks where the model should be sensitive to language variation (e.g., for form-based tasks like authorship verification, labels depend on whether a text uses British or American spelling). We pre-train BERT base models with the popular Byte-Pair Encoding algorithm to investigate how key tokenization design choices impact the performance of downstream models: the corpus used to train the tokenizer, the pre-tokenizer and the vocabulary size. We find that the best tokenizer varies on the two task types and that the pre-tokenizer has the biggest overall impact on performance. Further, we introduce a new approach to estimate tokenizer impact on downstream LLM performance, showing substantial improvement over metrics like Rényi efficiency. We encourage more work on language variation and its relation to tokenizers and thus LLM performance.

## 1 Introduction

Variation in language is ubiquitous, manifesting in forms such as lexical variation (e.g., vocabulary choice like “lift” vs. “elevator”), spelling variation (e.g., spelling “u” instead of “you”), and syntactic variation (e.g., word order like “I gave her the ball.” and “I gave the ball to her.”). Such variation is often systematic rather than random, and is linked to but regional, social, and contextual factors (Nguyen et al., 2016; Coupland, 2007; Eckert, 2012).

Tokenizers break up input strings and determine the tokens that are fed into language models. To-

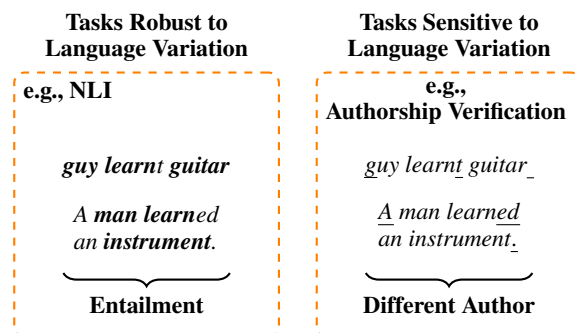


Figure 1: **Different types of tasks might profit from different tokenizer settings.** We investigate whether the same tokenizer performs well for tasks that require robustness to language variation (i.e., semantic-focused tasks like NLI) and for tasks that require sensitivity to language variation (i.e., form-focused tasks like authorship verification). Intuitively, one needs to look at **semantic signals** (e.g., “guitar” and “instrument”) for NLI and at **form- or style-based signals** (e.g., “learned” and “learnt”) for authorship verification.

kenizers build their vocabulary based on a *fitting corpus* and have a given *vocabulary size*. A *pre-tokenizer* can prioritize or prevent the creation of certain tokens. Depending on these settings, tokenizers might behave differently for linguistic forms that are less common (see §3). For example, spelling variants of common words might be split into more tokens than their standard spelling (Matthews et al., 2024), e.g., “doing” might be one token, while “doin” might be split into “do” and “in”. As a result, LLMs might need to learn to re-compose such spellings to represent their meaning, suffering in performance compared to representing words as one token. In contrast, smaller tokens could be useful for recognizing form-based patterns (e.g., distinguish the “in” from the “ing” ending). We ask **RQ1**: *Do the same tokenizer settings (i.e., fitting corpus, vocabulary size and pre-tokenizer) perform well for two types of downstream tasks: Tasks whose gold labels are **robust***

	Example	Label
GLUE	How do I hire an ethical hacker? Where can I find ethical hackers?	duplicate
+typo	Wher're cacn I fien ethical hackprs? Hjw fo I hie an ehtical hacker?	duplicate
+dialect	How do I hire me an ethical hacker? Where can I find me ethical hackers?	duplicate

(a) Tasks requiring robustness to language variation.

Dataset	Task
AV	Authorship Verification
PAN	Author Change Classification
CORE	Register Classification
NUCLE	Error Classification
Dialect	Dialect Classification

(b) Tasks requiring sensitivity to language variation.

Table 1: **Evaluation tasks.** We evaluate tokenizers on tasks that require (a) robustness and (b) sensitivity to language variation. For (a), we use GLUE tasks, adding spelling perturbations using Wang et al. (2021) and dialect perturbations using Ziems et al. (2023). Examples are taken from QQP. For (b), we use a newly compiled selection of tasks containing our own authorship verification dataset, an author change classification dataset (Bevendorff et al., 2024), a register classification dataset (Laippala et al., 2023), a grammatical error correction dataset (Dahlmeier et al., 2013) and a dialect classification dataset generated with Ziems et al. (2023).

*to language variation* (e.g., tasks that involve analyzing the semantic meaning of texts like NLI) and tasks whose gold labels are **sensitive to language variation** (e.g., tasks centered around the form, style or language variety of texts like authorship verification and dialect classification)? For example, for semantic tasks like NLI, models should perform equally well, regardless of whether the text is written in Standard American English or African American English. In contrast, for form-based tasks like authorship verification (i.e., are two texts written by the same author?), models should be able to distinguish British and American spelling. Intuitively a model needs to use different signals (see Figure 1) to solve tasks requiring robustness and tasks requiring sensitivity to language variation, and might thus profit from different tokenizer settings. We test this by pre-training BERT base models (Devlin et al., 2019) with different tokenizer settings of the de-facto standard tokenizer: Byte-Pair Encoding or BPE (Sennrich et al., 2016).

The evaluation of tokenizers with fast proxy measures — usually based on how well the tokenizer compresses a reference corpus — offers a popular and cheaper alternative to training larger LMs for comparing different tokenizers. However, common proxy measures often do not achieve a high correlation with downstream performance (Zouhar et al., 2023; Cognetta et al., 2024; Schmidt et al., 2024). One reason could be that such measures are task-agnostic: They predict the same performance for a given downstream corpus, regardless of task or label differences (e.g., in Figure 1, NLI and authorship verification have different labels for the same sentence pair). We therefore investigate **RQ2**: *Can simple task-aware measures based on logistic regression better predict tokenizer downstream*

*performance on variation-robust and variation-sensitive tasks than the common task-agnostic measures Rényi efficiency and Corpus Token Count?*

**Contributions.** We make the following three contributions: (1) We provide our code and a selection of tasks (see Table 1) to evaluate models on tasks requiring sensitivity and robustness to English language variation.<sup>1</sup> (2) We find that the impact of tokenizers on downstream performance varies depending on whether a task requires robustness or sensitivity to language variation, and that the pre-tokenizer has the biggest influence on downstream performance across task types. However, aggregated performance differences remain small and highlight the need for future work to investigate different types of language variation individually. (3) Further, we find that logistic regression performance has a higher correlation with BERT downstream performance than metrics like Corpus Token Count or Rényi efficiency (Zouhar et al., 2023). We provide practical suggestions to evaluate and build better tokenizers. With this work, we hope to encourage more work on language variation and its relation to tokenizers to build better, fairer and more robust LLMs from the ground up.

## 2 Related Work

Next to Byte-Pair Encoding (Sennrich et al., 2016), there are several other subword (and other) tokenizers (Mielke et al., 2021). For example, character and byte-based tokenizers have been argued to be more robust to spelling variations (Mielke et al., 2021; Libovický et al., 2022; Xue et al., 2022). In this study we focus on BPE as it has been the most

<sup>1</sup><https://github.com/nlpsoc/Tokenization-Language-Variation>

common tokenization algorithm in recent LLMs (e.g., Llama 3, Mixtral, DeepSeek v3 and GPT-4).

There is no universally agreed-upon standard to evaluate tokenizers. Tokenizers have been evaluated *intrinsically* (i.e., without training larger LMs) and *extrinsically* (i.e., considering the performance of larger LMs pre-trained with the considered tokenizer). Common intrinsic methods include: the average number of subwords produced per word and correlated measures like Corpus Token Count (Rust et al., 2021; Scao et al., 2022; Ali et al., 2024; Gallé, 2019; Schmidt et al., 2024), information theoretic measures like Shannon entropy or Rényi efficiency (Zouhar et al., 2023) and morphological alignment (Gow-Smith et al., 2022; Uzan et al., 2024). Extrinsic measures include LM perplexity (Shliazhko et al., 2024; Zevallos and Bel, 2023; Gowda and May, 2020), cross-entropy loss (Rajaraman et al., 2024), computational training cost (Ali et al., 2024) and downstream task performances (Schmidt et al., 2024; Ali et al., 2024). It is computationally infeasible to train SOTA LLMs end-to-end for each version of the tokenizer one wants to evaluate. Recent work tackled this by training “smaller” generative language models with 350M–2.5B parameters for each tokenizer (Schmidt et al., 2024; Ali et al., 2024). However, even such smaller models still can take several days to train. We measure downstream task performance on models with 110M parameters taking less than 15 GPU hours to train per model.

Tokenizer algorithms and settings can affect an LLM’s performance, for example, on tasks including numbers like arithmetic (Thawani et al., 2021; Wallace et al., 2019), on tasks including domain specific vocabulary or jargon like coding or medicine (Gu et al., 2021; Dehaerne et al., 2022; Zan et al., 2023), on different scripts and languages (Petrov et al., 2023; Rust et al., 2021; Limisiewicz et al., 2023; Ahia et al., 2023; Velayuthan and Sarveswaran, 2025) and when translating between languages (Gallé, 2019; Libovický et al., 2022; Zhang et al., 2022). To the best of our knowledge monolingual tokenizers have been underinvestigated in relation to language variation. Monolingual tokenizers and tokenizer settings have recently been investigated on broader selections of NLU tasks (Schmidt et al., 2024; Ali et al., 2024), presumably with the underlying assumption that for a given language like English, there exists a best tokenizer for most, if not all, tasks. We investigate this assumption for two systematically different

types of tasks: tasks requiring robustness and tasks requiring sensitivity to language variation.

### 3 Tokenizer Settings

Tokenizers break up texts into smaller units. These units are fed into the language model as input. We investigate different variations of the most popular tokenization algorithm: Byte-Pair Encoding or BPE (Sennrich et al., 2016). BPE is a subword tokenization algorithm (i.e., it breaks rare words down into subword units, bottoming out in bytes in the worst case). The vocabulary is built iteratively: It starts out with a base vocabulary of distinct bytes and merges them based on token frequency in the fitting corpus until the desired vocabulary size is reached. We vary the algorithm on three parameters: (1) the vocabulary size, (2) the pre-tokenizer and (3) the fitting corpus.

#### 3.1 Vocabulary Size

Common vocabulary sizes range from 30k to 64k in monolingual models, to about 128k (Llama 3) and 200k (GPT-4o) in recent multilingual models. Previous work tested vocabulary sizes between 32k and 100k on NLU tasks, finding only small differences in performance between vocabulary sizes for the same model size (Ali et al., 2024; Schmidt et al., 2024). However, vocabulary size might still play a role when dealing with language variation. When using lower range vocabulary sizes, less common words or spelling variants might not be represented with one token, influencing a word’s position in the embedding space (Matthews et al., 2024). Character-based models like ByT5 show better robustness to spelling variation than subword models with larger vocabulary sizes (Libovický et al., 2022; Xue et al., 2022; Tay et al., 2021). Similarly, tokenizers using low-range (in the extreme case character-level) vocabulary sizes might be more robust to spelling variation, as, for example, swapping of characters will not drastically alter the character-based segmentation of a word. We experiment with the following vocabulary sizes: 500 and 4k (low-range vocabulary size), 32k and 64k (mid-range vocabulary size) and 128k (high-range vocabulary size).

#### 3.2 Pre-Tokenizer

Pre-tokenizers split the fitting corpus into word-like units (Mielke et al., 2021), so-called “pre-tokens” (e.g., “I’m fine.” → “I”, “’m”, “fine”, “.”) which are then further tokenized independently. That is,

in building its vocabulary, the BPE algorithm cannot merge tokens that cross pre-token boundaries (e.g., “I’m”). Pre-tokenizers are usually expressed with regular expressions that include hard-coded knowledge about a language or script (Velayuthan and Sarveswaran, 2025), e.g., whitespaces separate words in Latin script and “m” is a contraction in English. The most common English pre-tokenizers are more or less elaborate variations of splitting texts based on whitespaces. For example, GPT-2’s pre-tokenizer additionally separates different *Unicode character categories* (e.g., letters, numbers and punctuation) but leaves single leading whitespaces attached to words (Radford et al., 2019).

Pre-tokenizers influence compression efficiency (Radford et al., 2019), NLU downstream performance (Schmidt et al., 2024) and arithmetic performance (Lee et al., 2024). However, pre-tokenizers have not been investigated for tasks that require sensitivity to language variation. Pre-tokenizers might be important for handling words that merge Unicode character categories, for example, when letters are substituted with numbers (e.g., “2day” or “c000l”, see Eger et al. (2019) for more). If a pre-tokenizer generally splits off numbers and letters, such words can never become part of the vocabulary. Further, not splitting on whitespace might better represent syntactic variation by including frequent phrases like “of the”. Frequent compositional phrases are processed faster by humans (Arnon and Snider, 2010), and building tokenizers that align closer to such cognitive processes (e.g., Yang et al., 2020) might also be beneficial for semantic tasks.

We compare five different pre-tokenizers: (a) not using a pre-tokenizer (NO), (b) isolating whitespaces (WS), (c) leaving single whitespaces attached to words (\_WS), (d) Llama 3’s pre-tokenizer (LLAMA3) and (e) GPT-2’s widely used pre-tokenizer (GPT2). LLAMA3 and GPT2 can be understood as extensions of the \_WS pre-tokenizer. Among others, the LLAMA3 pre-tokenizer splits off Unicode character categories (e.g., punctuation) but leaves one leading non-letter character attached to letters (e.g., ’m). The GPT2 pre-tokenizer is similar to LLAMA3, but usually separates Unicode character categories more. See Appendix A.2 for the regular expressions describing the WS, \_WS, LLAMA3 and GPT2 pre-tokenizer.

### 3.3 Fitting Corpus

Sub-word tokenizers construct their vocabulary based on a fitting corpus, adding tokens based on

the distribution of tokens in that corpus. Lexical, morphological and spelling variation are all intuitively connected to the fitting corpus. If a fitting corpus does not show a variation, it can not be part of a single token in the vocabulary.

We sampled fitting corpora with a size of approximately 1.5 billion words. **PubMed** was randomly sampled from The Pile’s PubMed Abstracts (Gao et al., 2020). The PubMed Abstracts consist of 30M abstracts from biomedical articles. **Wikipedia** was randomly sampled from Wikipedia articles from a snapshot from June 1st, 2023 after the plain text of articles was extracted.<sup>2</sup> **Twitter** was sampled from the Decahose Twitter stream<sup>3</sup> throughout the year 2021, with queries on almost every day of the year 2021. We only select English tweets based on Twitter’s internal language identification system. We exclude retweets. **Miscellaneous** was sampled from a variety of domains with no overlap with the other fitting corpora. It includes Reddit, literature sources (fanfictions and books), news articles and comments, question answering websites, reviews, mails, transcripts, blogs, Common Crawl, scientific articles, code and mathematical formulas. See details in Appendix A.1.

We expect the fitting corpora to differ in the lexical, syntactic and lexical variation they exhibit. For example, PubMed probably contains less spelling and lexical variation than Twitter. Intuitively, a tokenizer constructed on PubMed should thus be less capable in representing stylistic variation than one constructed on Twitter. However, it remains unclear how important the fitting corpus composition is. Zhang et al. (2022) investigate different compositions of the fitting corpus in the multilingual setting. They find a surprising robustness to language imbalance in the fitting corpus for languages sharing the same script.

## 4 Evaluation Tasks

We compare tokenizers on classification tasks that require robustness to language variation (§4.1) and tasks that require sensitivity to language variation (§4.2). Models solving these two types of tasks should need to make use of more semantic and form-based signals respectively and might have different requirements for a tokenizer (cf. Figure 1). We select tasks that strike a balance between being

<sup>2</sup>using <https://github.com/LorcanJConolly/WikiTextExtractor>.

<sup>3</sup>Decahose provided access to 10% of real time tweets sampled by Twitter.



sufficiently challenging and staying within the capabilities of our pre-trained BERT models. See an overview of selected tasks in Table 1.

#### 4.1 Tasks Robust to Language Variation

First, we evaluate tokenizers on semantic tasks where the gold label is robust to language variation. We use GLUE (Wang et al., 2018), a standard NLP benchmark that was also used to evaluate BERT at its introduction (Devlin et al., 2019), and is within the capabilities of our pre-trained BERT models. We compare tokenizers on the following four GLUE tasks: SST-2 (sentiment classification), QQP (paraphrase classification), MNLI and QNLI (NLI tasks). For details on our task selection see Appendix B.1. Ideally, a robust tokenizer performs consistently across all versions of **GLUE**: the original, primarily written in Standard American English, the spelling-transformed **GLUE+typo** and the dialect-transformed **GLUE+dialect**.

**GLUE+typo.** We use `textflint` (Wang et al., 2021) to introduce simulated typos and spelling errors to our tasks, similar to Libovický et al. (2022). It uses random character swapping and a list of common spelling errors.

**GLUE+dialect.** We use `Multi-VALUE` (Ziems et al., 2023) to introduce simulated dialectal variation to our GLUE tasks. `Multi-VALUE` makes use of 189 dialectal perturbation rules (Ziems et al., 2023) operationalized based off of `eWAVE` (Kortmann et al., 2020). For each example in the GLUE tasks, we randomly choose between transformations to the following dialects: Appalachian English, Chicano English, Colloquial Singapore English, Indian English and Urban African American English.

#### 4.2 Tasks Sensitive to Language Variation

Second, we evaluate tokenizers on tasks that are sensitive to language variation. The gold label for such tasks should be sensitive to stylistic or form-based signals.

**AV.** Models performing authorship verification (i.e., are two texts written by the same author?) usually need to be sensitive to the different styles and forms used by different authors (Zhu and Jurgens, 2021; Wegmann et al., 2022; Wang et al., 2023). Past work has found authorship verification to be sensitive to tokenization, with significant gaps between the BERT and RoBERTa tokenizers (Zhu and Jurgens, 2021). Therefore, we curate a new authorship verification dataset of 40.8k train, 2.5k

dev and 4.8k test pairs of texts from different domains, similar in distribution to the Miscellaneous corpus (cf. §3). Labels are balanced in the test set. See details in Appendix B.2.

**PAN.** We use the PAN 2024 Multi-Author Writing Style Analysis task to predict whether an author shift occurs between two consecutive paragraphs extracted from Reddit (Bevendorff et al., 2024). Specifically, we use the ‘hard’ task, where paragraphs are about the same topic. We sample such that labels are balanced. This results in a training set of 18k and a dev set of 4k instances.

**NUCLE.** We use the NUCLE 3.3 corpus (Dahlmeier et al., 2013) for multi-label classification of the errors that were made by English learners in a given sentence. NUCLE was annotated by professional English instructors for 27 error types (e.g., verb tense or article use). From NUCLE, we take all 22k unique sentences containing errors and sample 5k error-free sentences. We split it into a train (80%) and dev (10%) set.

**CORE.** We use the `Corpus of Online Registers of English` (Laippala et al., 2023) for register classification. Register is one of the most important factors associated with linguistic variation (Biber, 2012). We use the 8 main register labels (e.g., spoken or informational description) for multi-class prediction. To increase the occurrence of rarer labels, we split long texts and reach a train size of 30k and a dev size of 5k. See Appendix B.2 for details.

**Dialect.** We randomly sample 60k instances from **GLUE+dialect** (§4.1) and the original GLUE task, to create a dialect classification task with five dialects and Standard American English in the original GLUE texts. We use 50k texts for the train and 5k for the dev set.

## 5 Modeling

For each investigated tokenizer (cf. §3), we pre-train three BERT models with different seeds. We use encoder instead of decoder models, as encoder models tend to reach higher performance for classification tasks for low parameter settings. This allows us to train models using fewer GPU hours.

**Pretraining BERT Models.** We experiment with pre-training tiny BERT models using a token count  $T$  close to 3300M that is exponentially bigger than the 4.6M parameters  $P$ , similar to the original BERT papers (Devlin et al., 2019; Turc et al., 2019). However, we find that for the same compute, using

Tokenizer Setting	Model	GLUE	+typo	+dialect	AVG
Fitting Corpus	PubMed	80.8 $\pm$ 0.0	<b>69.1</b> $\pm$ 0.2	78.6 $\pm$ 0.2	76.2 $\pm$ 0.0
	Wikipedia	80.7 $\pm$ 0.3	68.6 $\pm$ 0.2	<b>79.3</b> $\pm$ 0.2	76.2 $\pm$ 0.2
	Twitter	<b>81.1</b> $\pm$ 0.0	69.1 $\pm$ 0.5	78.8 $\pm$ 0.1	<b>76.4</b> $\pm$ 0.2
Pre-Tokenizer	NO	72.1 $\pm$ 1.0	61.6 $\pm$ 0.1	70.1 $\pm$ 0.3	67.9 $\pm$ 0.3
	WS	80.8 $\pm$ 0.4	68.2 $\pm$ 0.3	<b>79.3</b> $\pm$ 0.3	76.1 $\pm$ 0.3
	_WS	80.8 $\pm$ 0.3	<b>68.9</b> $\pm$ 0.1	79.0 $\pm$ 0.2	<b>76.2</b> $\pm$ 0.1
	LLAMA3	80.9 $\pm$ 0.1	68.2 $\pm$ 0.2	79.0 $\pm$ 0.0	76.1 $\pm$ 0.0
	GPT2	<b>81.3</b> $\pm$ 0.4	68.2 $\pm$ 0.2	79.2 $\pm$ 0.4	<b>76.2</b> $\pm$ 0.3
Vocabulary Size	500	77.2 $\pm$ 2.5	<b>70.3</b> $\pm$ 2.6	75.6 $\pm$ 2.0	74.4 $\pm$ 2.4
	4k	80.5 $\pm$ 0.8	<b>70.3</b> $\pm$ 0.9	78.6 $\pm$ 0.8	<b>76.4</b> $\pm$ 0.8
	32k	<b>81.3</b> $\pm$ 0.4	68.2 $\pm$ 0.2	<b>79.2</b> $\pm$ 0.4	76.2 $\pm$ 0.3
	64k	80.8 $\pm$ 0.4	67.6 $\pm$ 0.6	<b>79.2</b> $\pm$ 0.2	75.9 $\pm$ 0.4
	128k	78.7 $\pm$ 2.0	64.6 $\pm$ 1.9	76.1 $\pm$ 2.6	73.1 $\pm$ 2.2

Table 2: **Performance on tasks requiring robustness to language variation.** We display BERT performances, averaged on the original four GLUE tasks and their perturbations using spelling mistakes (+typo) and dialectal transformations (+dialect). We provide the mean and standard deviation ( $\pm$ ) over three seeds respectively. We boldface the best performances for each column and investigated setting. For the averaging column (AVG), italics indicate tokenizers with no statistically significant difference from the best-performing tokenizer (cf. Figure 2).

a bigger model size  $P$  and fewer tokens  $T$  improves the training loss. Chinchilla’s scaling law might also hold for smaller encoder models, specifically optimal parameter count could scale with the token size for a fixed compute  $P_{OPT} \approx T^{23/27}$  (Hoffmann et al., 2024). For the remainder of this work, we use the base BERT model architecture with 110M parameters, initialize all weights randomly and pre-train on 750M tokens sampled in sequences of 512 from the Miscellaneous corpus (§3.3) and use a batch size of 32 and 45k steps. For further details and hyperparameters, see Appendix C.

**Fine-Tuning BERT.** Unless otherwise specified, we use 3 epochs, a max sequence length of 128, a batch size of 32 and a learning rate of  $2e-5$  to perform the classification tasks. We evaluate on the dev set for GLUE. For comparability, we use the same setup for tasks requiring sensitivity to language variation. Only for the authorship verification task we use a contrastive training setup, then use the dev set to find an optimal cosine similarity threshold and calculate accuracy on the test set.

## 6 Results

We show the performance of fine-tuned BERT models on tasks that require robustness to language variation in Table 2 and tasks that require sensitivity to language variation in Table 3. When we investigate a specific setting (e.g., the fitting corpus in the first three rows in Table 2), we only change that setting and leave the other at their “default settings” to ensure comparability and isolate the effect of each individual setting without exhaustive testing

of all possible combinations. We use the following default values for the three settings: the miscellaneous fitting corpus, the GPT2 pre-tokenizer and a vocabulary size of 32k.

Note that the performance differences tend to be relatively small, which is consistent with previous work comparing different tokenizer algorithms on downstream tasks (Ali et al., 2024; Schmidt et al., 2024). To ensure significance, we compute the pairwise McNemar, 1947’s test for the pre-tokenizer, fitting corpus and vocabulary size settings, see Figure 2. For the significance testing, we consider classifications by models with the same settings but different seeds to be stemming from the same rater. We use the Bonferroni correction (Dunn, 1961) for our total of 26 tests.

**RQ1: Tokenizer settings perform differently on tasks robust and sensitive to language variation.** Overall, tasks sensitive to language variation profit from tokenizers that encode more variation through a larger vocabulary size than 4k (§6.3). Note that the best-performing tokenizer settings are not always consistent across the individual variation-robust and variation-sensitive tasks (e.g., vocabulary size for GLUE+typo and GLUE in Table 2). We suspect that this is due to differences in the types of language variation present in the specific task datasets (e.g., character-level tokens seem to be more robust to spelling variation). Future work should investigate different types of language variation individually (e.g., spelling vs. lexical variation).

	Model	AV (acc) $\uparrow$	PAN (acc) $\uparrow$	CORE (acc) $\uparrow$	NUCLE (F1) $\uparrow$	Dialect (F1) $\uparrow$	AVG
<b>Corpus</b>	PMed	81.7 $\pm$ 0.1	65.2 $\pm$ 0.5	55.9 $\pm$ 0.6	21.8 $\pm$ 1.2	87.9 $\pm$ 0.1	62.5 $\pm$ 0.4
	Wiki	81.9 $\pm$ 0.1	65.5 $\pm$ 0.5	55.5 $\pm$ 0.6	<b>23.5</b> $\pm$ 0.2	<b>88.9</b> $\pm$ 0.5	63.1 $\pm$ 0.2
	Twitter	<b>82.9</b> $\pm$ 0.6	<b>66.7</b> $\pm$ 0.4	<b>56.5</b> $\pm$ 0.6	21.4 $\pm$ 0.9	88.3 $\pm$ 0.2	<b>63.2</b> $\pm$ 0.2
<b>Pre-Tok.</b>	NO	81.8 $\pm$ 0.3	59.9 $\pm$ 1.0	51.7 $\pm$ 0.2	16.3 $\pm$ 0.2	77.3 $\pm$ 0.2	57.4 $\pm$ 0.2
	WS	75.5 $\pm$ 0.9	66.1 $\pm$ 0.4	55.1 $\pm$ 0.9	<b>23.0</b> $\pm$ 0.1	<b>88.8</b> $\pm$ 0.2	61.7 $\pm$ 0.3
	_WS	82.5 $\pm$ 0.3	66.3 $\pm$ 1.6	<b>56.6</b> $\pm$ 0.7	22.6 $\pm$ 0.5	88.4 $\pm$ 0.2	<b>63.3</b> $\pm$ 0.3
	LLAMA3	82.5 $\pm$ 0.2	<b>66.9</b> $\pm$ 0.6	56.5 $\pm$ 0.6	21.1 $\pm$ 1.2	88.6 $\pm$ 0.0	63.1 $\pm$ 0.3
	GPT2	<b>82.6</b> $\pm$ 0.5	66.6 $\pm$ 1.2	56.3 $\pm$ 1.2	21.8 $\pm$ 0.9	88.4 $\pm$ 0.6	63.1 $\pm$ 0.4
<b>Size</b>	500	78.2 $\pm$ 0.9	62.6 $\pm$ 0.2	51.1 $\pm$ 0.6	13.1 $\pm$ 0.9	85.6 $\pm$ 0.4	58.1 $\pm$ 0.6
	4k	81.9 $\pm$ 0.1	<b>67.8</b> $\pm$ 0.6	55.1 $\pm$ 1.0	17.4 $\pm$ 2.8	87.9 $\pm$ 0.7	62.0 $\pm$ 0.6
	32k	82.6 $\pm$ 0.5	66.6 $\pm$ 1.2	<b>56.3</b> $\pm$ 1.2	21.8 $\pm$ 0.9	<b>88.4</b> $\pm$ 0.6	<b>63.1</b> $\pm$ 0.4
	64k	<b>82.7</b> $\pm$ 0.4	67.2 $\pm$ 0.7	54.9 $\pm$ 1.5	<b>22.0</b> $\pm$ 1.1	88.1 $\pm$ 0.6	63.0 $\pm$ 0.6
	128k	80.1 $\pm$ 2.2	62.4 $\pm$ 3.4	51.2 $\pm$ 2.2	19.0 $\pm$ 2.7	81.5 $\pm$ 5.1	58.9 $\pm$ 3.0

Table 3: **Performance on tasks requiring sensitivity to language variation.** We display BERT performances on our Authorship Verification (AV), the PAN, the CORE, and the multi-Dialect dataset. We provide the mean and standard deviation ( $\pm$ ) over three seeds respectively. For the averaging column (AVG), italics indicate tokenizers with no statistically significant difference from the best-performing tokenizer (cf. Figure 2).

## 6.1 Pre-Tokenizers

**Pre-tokenizers have the greatest influence on performance.** For both task types, the range of performance values is largest for pre-tokenizers, second largest for vocabulary size and smallest for fitting corpus. This is surprising as, to the best of our knowledge, pre-tokenizers have received the least attention in previous work. For both task types, using pre-tokenizers improves performance over using no pre-tokenizer.

**Pre-tokenizer performance differs more between individual tasks than between task types.** For both task types, \_WS is among the best performing pre-tokenizers. Combining leading whitespaces with letters (\_WS) generally improved performance over isolating whitespaces (WS). Notably, \_WS cuts the input tokens to a model in half as single whitespaces between words are not separate tokens anymore. For dialect and spelling-transformed GLUE, CORE, NUCLE and Dialect, the whitespace-based pre-tokenizers perform the best. For AV, PAN and the original GLUE tasks, GPT2 and LLAMA perform better. Tasks like grammatical error detection (NUCLE) could be seen as identifying deviations from a standard, and might benefit from pre-tokenizers that include tokens with typical combinations of Unicode character categories. One of the main differences of GPT2 and LLAMA to other pre-tokenizers is that the combination of different Unicode character categories (e.g., punctuation and letters) within the same token is less often allowed. By preventing combinations of categories like punctuation and letters, the vocabulary of GPT2 can include a

broader range of tokens that only consist of letters, e.g., “\_queer” in Table 4. This could explain the better performance of GPT2 on the original GLUE task. In contrast, LLAMA3 also allows the mixing of one initial punctuation mark with letters, e.g., including “‘all” in Table 4, which after manual inspection seems to especially help LLAMA3 solve PAN. Future work could investigate different individual tasks and the influence of pre-tokenizers individually.

## 6.2 Fitting Corpus

**On tasks requiring sensitivity to language variation, Twitter performs best.** This aligns with our expectation that the Twitter corpus includes more spelling variation and a larger set of language varieties than PubMed and Wikipedia. Interestingly, Wikipedia performs the best on NUCLE and Dialect. Grammatical error detection (NUCLE) could be seen as identifying deviations from a standard, and might benefit from training on corpora with few grammatical errors — like Wikipedia.

**On tasks requiring robustness to language variation, Twitter performs surprisingly well.** Interestingly, Twitter performs indistinguishably from other fitting corpora on tasks that require robustness to language variation. Originally, we expected a more standardized corpus like Wikipedia to perform better on the original GLUE task, as it should lead to less spelling variation and thus more “full words” in the vocabulary (e.g., “precursor” in Table 4).

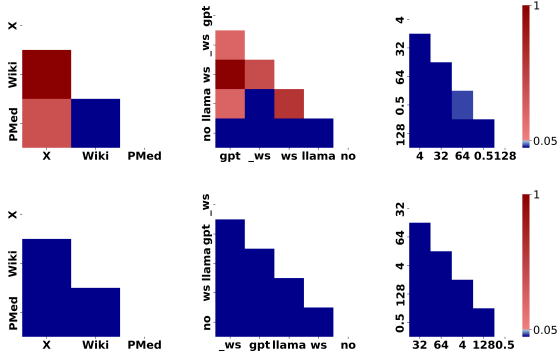


Figure 2: **Pairwise significance testing of models.** We use pairwise McNemar (1947)’s test to test whether there is a significant difference between classifications done by models trained with different tokenizers on the tasks requiring robustness to language variation (first row) and tasks requiring sensitivity to language variation (second row). Tokenizers are sorted by mean performance. Blue colors show statistical significance, while red colors are above the 0.05 threshold.

### 6.3 Vocabulary Size

**A larger vocabulary size might be useful for tasks requiring sensitivity to language variation.** Performance peaked at 4k for tasks requiring robustness to language variation and at 32k for tasks requiring sensitivity to language variation. It seems that a smaller vocabulary size is sufficient to learn common standard forms and more tokens are needed to include style and form variations. For example, “jumper” is only included with vocabulary size 64k, see Table 4. Note that optimal vocabulary sizes might scale with model size (Tao et al., 2024). For both task types, the standard deviation is highest for the smallest (500) and the largest vocabulary size (128k).

**Smaller vocabulary is more robust to spelling variation.** The vocabulary size of 500 consistently performs worse than larger vocabulary sizes. It is possible that BERT struggles with long input sequences and learning to compose words for small vocabulary sizes. An exception is the GLUE+typo task, where the tokenizer of size 500 performs best. This is consistent with character-based tokenizers being more robust to spelling variation (cf. §3.1).

## 7 Pre-evaluating Downstream Tokenizer Impact

A popular alternative to training several larger LMs to evaluate tokenizers is using fast *intrinsic* measures on the token distribution on the downstream

Group	Choice	Examples of tokens
Corpus	Wiki	_queer_precursor_l ma o
	Twitter	_queer_prec urs or _lmao
	PubMed	_qu e er _precursor _l ma o
Pre-Tok	NO	_y'all _que er _b 4
	WS	_y'all _que er _b4
	_WS	_y'all _que er _b 4
	LLAMA3	_y 'all _que er _b 4
	GPT2	_y ' all _queer _b 4
Size	500	_t e h _l g b t q i a + _j um p er
	4k	_te h _l g bt q ia + _j um per
	32k	_te h _l g bt q ia + _j umper
	64k	_teh _lg bt q ia + _jumper
	128k	_teh _lgbt q ia + _jumper

Table 4: **Examples of vocabulary differences.** We display examples of sequences tokenized by tokenizers with varying fitting corpora, pre-tokenizers and vocabulary size. We represent whitespaces within tokens as `_`. The domain of the fitting corpus affects the set of unique words in the vocabulary (e.g., “queer” is part of Twitter but not PubMed). The size determines the number of “rarer” words in the fitting corpus that are added to the vocabulary (e.g., the British variant “jumper” of “sweater” is only added with size 64k). The pre-tokenizer determines what types of words can be part of the vocabulary (e.g., “b4” can not be part of GPT2 and LLAMA3). Note that allowing fewer Unicode character category combinations (e.g., numbers and letters) increases the broadness of the vocabulary (e.g., in our fitting corpora, GPT2 allows for fewer combinations than other pre-tokenizers and is the only one to include “queer” in its vocabulary).

task corpus. However, common intrinsic measures are not consistently correlated with downstream performance (Rust et al., 2021; Zouhar et al., 2023; Schmidt et al., 2024). Further, they do not make use of task labels and can thus be considered *task agnostic*. Imagine a corpus on which you want to perform both a task requiring robustness to language variation and a task requiring sensitivity to language variation (cf. Figure 1). Task-agnostic intrinsic measures like Corpus Token Count and Rényi efficiency (Zouhar et al., 2023) will always predict the same performance for both tasks. However, we found that the type of task can have an influence on what tokenizer settings work better (cf. §6).

### 7.1 Logistic Regression Measure

We experiment with a task-dependent measure: the performance<sup>4</sup> of logistic regression with task labels

<sup>4</sup>We use the same F1 and accuracy performance metrics as for BERT.



as the dependent variable. Specifically, we use a tokenizer’s vocabulary as the feature set with a bag-of-token approach. For tasks with two texts as inputs, token combinations between sentence pairs serve as features. For multi-class and multi-label tasks, we use one-vs-rest models. See Appendix D for details.

## 7.2 Results

We display the Pearson correlation of Rényi efficiency, Corpus Token Count and logistic regression performances with the fine-tuned BERT performances in Table 5. Similar to Schmidt et al. (2024) we find a light negative correlation of LM performances with Rényi efficiency on NLU tasks.

**Corpus token count has a negative correlation for tasks requiring robustness and a positive correlation for tasks requiring sensitivity to language variation.** Corpus Token Count and correlated measures are often used to assess the compression quality of a tokenizer on a reference corpus or downstream task corpus. A higher value entails more tokens in the corpus and thus worse compression. Better compression is commonly believed to be a sign of a better tokenizer, and has often been thought to be correlated with better downstream performance (Rust et al., 2021; Ali et al., 2024; Velayuthan and Sarveswaran, 2025). We show that the correlation is flipped for tasks that are robust and sensitive to language variation highlighting the difference in tokenizer requirements for the two task types.

**RQ2: Logistic regression correlates with downstream performance.** Among the three considered measures, we find the highest correlation between logistic regression performances and BERT performances. Additionally, logistic regression is similarly correlated for both task types and can compare tokenizers of varying vocabulary sizes, which is not possible with Rényi efficiency and Corpus Token Count (for details refer to Appendix D). Note that we by no means question the usefulness of measures like Corpus Token Count or Rényi efficiency for assessing different tokenizers and what they do. However, they might be less suited to estimate downstream performance without additional modifications.

## 8 Conclusion

In this work, we investigated tokenizer settings for tasks that require robustness and tasks that re-

Measure	Robust	Sensitive
Rényi Efficiency	-.22	-.03
Corpus Token Count	-.45	.37
logistic regression	.85	.84

Table 5: **Correlation of intrinsic measures with BERT performances.** Correlations between the Rényi Efficiency ( $\alpha = 2.5$ ), Corpus Token Count, and logistic regression predictions with the performance of the fine-tuned BERT models on the tasks robust to language variation and the tasks sensitive to language variation respectively. Logistic regression has the highest correlation values across all tasks. A higher Corpus Token Count is negatively correlated with performance on tasks robust to language variation and positively correlated with tasks sensitive to language variation.

quire sensitivity to language variation. BPE settings perform differently on the two task types. We make three practical suggestions for selecting tokenizer settings: (1) Pay the most attention to the pre-tokenizer. It influences how Unicode character categories can be combined (e.g., “b4” in Table 4), and how many different words are ultimately part of the vocabulary (e.g., “\_queer” in Table 4). (2) Choose a bigger vocabulary size for settings that require sensitivity to language variation. (3) Use a small machine learning classifier — e.g., a logistic regression classifier — to evaluate how different tokenizers affect performance on tasks in your domain. For example, for general-purpose language models, pick a tokenizer that this model consistently predicts to perform well across both variation-robust and variation-sensitive tasks. The best tokenizer settings might be sensitive to what types of language variation (e.g., lexical vs. syntactic) are present. We think it is crucial to investigate different types of language variation individually in future work to ultimately build better, fairer and more robust LLMs.

## Limitations

Note that our taxonomy of tasks relies on splitting tasks into semantic-focused tasks (i.e., considering what is said) and form-based tasks (i.e., considering how it is said). However, a strict distinction is challenging since most tasks are best solved using a mixture of typical content and typical form signals. For example, recognizing the lengthening of words (e.g., “cooooo”) can be helpful for the semantic task of sentiment classification (Brody and Diakopoulos, 2011) and content information (e.g., the topic) can improve authorship verification performance

(Wegmann et al., 2022). Still, this distinction clarifies the main signals that our different task types should rely on and enables us to compare tokenizer settings on these two different types of tasks.

We evaluate tokenizers on 110M parameter encoder models, because they tend to reach higher performance for classification tasks in low parameter settings. However, we risk not accounting for emergent capabilities of popular larger generative language models. For example, the number of context tokens for recent models are generally much larger than the original 512 tokens of BERT base.

We find that logistic regression correlates with downstream performance on classification tasks. For more complex tasks like text generation, logistic regression probably will not be applicable. However, we still see it as a great benefit that logistic regression can enable us to quickly test tokenizers on a variety of different tasks and make a more informed decision on what tokenizers to test more rigorously.

For tasks requiring robustness to language variation, many more tasks could have been included. For example, future work could include challenge NLI datasets, e.g., ANLI (Nie et al., 2020). Further, future work could investigate different types of tasks such as question answering or arithmetic tasks. For tasks requiring sensitivity to language variation, we originally planned to include more classification tasks that are tied to well-established factors of language variation, e.g., age prediction and geographic location prediction (Nguyen et al., 2016). However, we repeatedly encountered difficulties with cleanly separating semantic from form-based cues. For age prediction and location prediction specifically, logistic regression models made extensive use of content words. Based on the coefficients of our logistic regression models, we excluded tasks that we expected to mainly require sensitivity to language variation, but where the models often relied on content words.

We varied one tokenizer setting at a time while leaving the other two on their “default values” (i.e., miscellaneous fitting corpus, GPT2 pre-tokenizer and vocabulary size of 32k), in order to isolate the effect of each individual setting without exhaustive testing of all possible combinations. Future work could explore interdependencies between tokenizer settings through more comprehensive testing.

The tokenizer settings seem to have different effects when the pre-training corpus and fitting corpus differ. Specifically, the pre-training corpus

might influence in how far a model can leverage the diversity encoded in the tokenizer. See Appendix E for further experiments. Note, however, that recent work suggests to only use vocabulary that appears in the training corpus (Land and Bartolo, 2024) – which is unlikely when different corpora are used for model training and tokenizer fitting.

## Ethical Considerations

Our pre-training and fitting corpora are largely based on publicly shared and accessible datasets from popular online forums and web pages (e.g., Wikipedia, Reddit, Common Crawl, ...). Unfortunately, these datasets were mostly collected without explicit consent from users and might lead to (among others) privacy concerns. Individuals might be identifiable from their written texts. However, we hope that the risks of reusing already published datasets are minimal. We collected tweets from Twitter in 2021 using Twitter’s official API, but opt to not share them publicly. While we aim to include different language varieties in our datasets, they might not be representative of English language use across different social groups. For example, we expect a skew towards straight, white, American, young and male authors. We caution against using our datasets and tasks for general claims about broad selection of language varieties. We confirm to have read and that we abide by the ACL Code of Ethics. Beside the mentioned ethical considerations, we do not foresee immediate risks of our work.

## Acknowledgements

We thank the anonymous ARR reviewers for their constructive comments. We thank the colleagues in the various NLP Groups at Utrecht University and the Blablalab at the University of Michigan and, specifically, Kees van Deemter, Antal van den Bosch, Yupei Du, Qixiang Fang, Melody Sepahpour-Fard, Shane Kaszefski Yaschuk, Elize Herrewijnen, Sanne Hoeken, Hugh Mee Wong, Jian Zhu and Pablo Mosteiro for, among others, feedback on writing and presentation. We thank colleagues at EMNLP 2024 for their insights into tokenization, and, specifically, Vilém Zouhar. We thank Laura Wegmann for feedback on writing. This research is supported by the “Digital Society - The Informed Citizen” research programme, which is (partly) financed by the Dutch Research Council (NWO), project 410.19.007. This research

## References

- 10968



- William B. Dolan and Chris Brockett. 2005. [Automatically constructing a corpus of sentential paraphrases](#). In *Proceedings of the Third International Workshop on Paraphrasing (IWP2005)*.
- Olive Jean Dunn. 1961. [Multiple comparisons among means](#). *Journal of the American statistical association*, 56(293):52–64.
- Penelope Eckert. 2012. [Three waves of variation study: The emergence of meaning in the study of sociolinguistic variation](#). *Annual review of Anthropology*, 41(1):87–100.
- Steffen Eger, Gözde Gül Şahin, Andreas Rücklé, Ji-Ung Lee, Claudia Schulz, Mohsen Mesgar, Krishnkant Swarnkar, Edwin Simpson, and Iryna Gurevych. 2019. [Text processing like humans do: Visually attacking and shielding NLP systems](#). 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 1634–1647, Minneapolis, Minnesota. Association for Computational Linguistics.
- Matthias Gallé. 2019. [Investigating the effectiveness of BPE: The power of shorter sequences](#). 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 1375–1381, Hong Kong, China. Association for Computational Linguistics.
- Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, Shawn Presser, and Connor Leahy. 2020. [The Pile: An 800gb dataset of diverse text for language modeling](#). *arXiv preprint*, arXiv:2101.00027.
- Edward Gow-Smith, Harish Tayyar Madabushi, Carolina Scarton, and Aline Villavicencio. 2022. [Improving tokenisation by alternative treatment of spaces](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11430–11443, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Thamme Gowda and Jonathan May. 2020. [Finding the optimal vocabulary size for neural machine translation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3955–3964, Online. Association for Computational Linguistics.
- Yu Gu, Robert Tinn, Hao Cheng, Michael Lucas, Naoto Usuyama, Xiaodong Liu, Tristan Naumann, Jianfeng Gao, and Hoifung Poon. 2021. [Domain-specific language model pretraining for biomedical natural language processing](#). *ACM Transactions on Computing for Healthcare (HEALTH)*, 3(1):1–23.
- Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, Tom Hennigan, Eric Noland, Katie Millican, George van den Driessche, Bogdan Damoc, Aurelia Guy, Simon Osindero, Karen Simonyan, Erich Elsen, Oriol Vinyals, Jack W. Rae, and Laurent Sifre. 2024. [Training compute-optimal large language models](#). In *Proceedings of the 36th International Conference on Neural Information Processing Systems*, Red Hook, NY, USA. Curran Associates Inc.
- Yupeng Hou, Jiacheng Li, Zhankui He, An Yan, Xiusi Chen, and Julian McAuley. 2024. [Bridging language and items for retrieval and recommendation](#). *arXiv preprint*, arXiv:2403.03952.
- Jean Kaddour. 2023. [The minipile challenge for data-efficient language models](#). *arXiv preprint*, arXiv:2304.08442.
- Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. 2020. [Supervised contrastive learning](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 18661–18673. Curran Associates, Inc.
- Varada Kolhatkar, Hanhan Wu, Luca Cavasso, Emilie Francis, Kavan Shukla, and Maite Taboada. 2020. [The sfu opinion and comments corpus: A corpus for the analysis of online news comments](#). *Corpus pragmatics*, 4:155–190.
- Bernd Kortmann, Kerstin Lunkenheimer, and Katharina Ehret, editors. 2020. *The Electronic World Atlas of Varieties of English*.
- Veronika Laippala, Samuel Rönqvist, Miika Oinonen, Aki-Juhani Kyröläinen, Anna Salmela, Douglas Biber, Jesse Egbert, and Sampo Pyysalo. 2023. [Register identification from the unrestricted open web using the corpus of online registers of english](#). *Language Resources and Evaluation*, 57(3):1045–1079.
- Sander Land and Max Bartolo. 2024. [Fishing for magikarp: Automatically detecting under-trained tokens in large language models](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 11631–11646, Miami, Florida, USA. Association for Computational Linguistics.
- Garreth Lee, Guilherme Penedo, Leandro von Werra, and Thomas Wolf. 2024. [From digits to decisions: How tokenization impacts arithmetic in LLMs](#). *Hugging Face blog*.
- Jindřich Libovický, Helmut Schmid, and Alexander Fraser. 2022. [Why don’t people use character-level machine translation?](#) In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2470–2485, Dublin, Ireland. Association for Computational Linguistics.
- Tomasz Limisiewicz, Jiří Balhar, and David Mareček. 2023. [Tokenization impacts multilingual language](#)



- modeling: Assessing vocabulary allocation and overlap across languages. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 5661–5681, Toronto, Canada. Association for Computational Linguistics.
- 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.
- Kyle Lo, Lucy Lu Wang, Mark Neumann, Rodney Kinney, and Daniel Weld. 2020. [S2ORC: The semantic scholar open research corpus](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4969–4983, Online. Association for Computational Linguistics.
- Jacob Matthews, John Starr, and Marten Schijndel. 2024. [Semantics or spelling? probing contextual word embeddings with orthographic noise](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 4495–4504, Bangkok, Thailand. Association for Computational Linguistics.
- Quinn McNemar. 1947. [Note on the sampling error of the difference between correlated proportions or percentages](#). *Psychometrika*, 12(2):153–157.
- Sabrina J Mielke, Zaid Alyafeai, Elizabeth Salesky, Colin Raffel, Manan Dey, Matthias Gallé, Arun Raja, Chenglei Si, Wilson Y Lee, Benoît Sagot, et al. 2021. [Between words and characters: A brief history of open-vocabulary modeling and tokenization in nlp](#). *arXiv preprint*, arXiv:2112.10508.
- Dong Nguyen, A. Seza Doğruöz, Carolyn P. Rosé, and Franciska de Jong. 2016. [Computational sociolinguistics: A Survey](#). *Computational Linguistics*, 42(3):537–593.
- Yixin Nie, Adina Williams, Emily Dinan, Mohit Bansal, Jason Weston, and Douwe Kiela. 2020. [Adversarial NLI: A new benchmark for natural language understanding](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4885–4901, Online. Association for Computational Linguistics.
- Aleksandar Petrov, Emanuele La Malfa, Philip Torr, and Adel Bibi. 2023. [Language model tokenizers introduce unfairness between languages](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 36963–36990. Curran Associates, Inc.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. [Language models are unsupervised multitask learners](#). *OpenAI blog*, 1(8):9.
- Jack W Rae, Anna Potapenko, Siddhant M Jayakumar, and Timothy P Lillicrap. 2019. [Compressive transformers for long-range sequence modelling](#). *arXiv preprint*, arXiv:1911.05507.
- Nived Rajaraman, Jiantao Jiao, and Kannan Ramchandran. 2024. [Toward a theory of tokenization in LLMs](#). *arXiv preprint*, arXiv:2404.08335.
- Phillip Rust, Jonas Pfeiffer, Ivan Vulić, Sebastian Ruder, and Iryna Gurevych. 2021. [How good is your tokenizer? on the monolingual performance of multilingual 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 3118–3135, Online. Association for Computational Linguistics.
- David Saxton, Edward Grefenstette, Felix Hill, and Pushmeet Kohli. 2019. [Analysing mathematical reasoning abilities of neural models](#). *arXiv preprint*, arXiv:1904.01557.
- Teven Le Scao, Angela Fan, Christopher Akiki, Elie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagn'e, Alexandra Sasha Luccioni, François Yvon, and Matthias Gallé et al. 2022. [Bloom: A 176b-parameter open-access multilingual language model](#). *arXiv preprint*, arXiv:2211.05100.
- Jonathan Schler, Moshe Koppel, and Shlomo Argamon. 2006. [Effects of age and gender on blogging](#). In *AAAI spring symposium: Computational approaches to analyzing weblogs*.
- Craig W Schmidt, Varshini Reddy, Haoran Zhang, Alec Alameddine, Omri Uzan, Yuval Pinter, and Chris Tanner. 2024. [Tokenization is more than compression](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 678–702, Miami, Florida, USA. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural machine translation of rare words with subword units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.
- Oleh Shliakhko, Alena Fenogenova, Maria Tikhonova, Anastasia Kozlova, Vladislav Mikhailov, and Tatiana Shavrina. 2024. [mGPT: Few-shot learners go multilingual](#). *Transactions of the Association for Computational Linguistics*, 12:58–79.
- 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.
- Chaofan Tao, Qian Liu, Longxu Dou, Niklas Muenighoff, Zhongwei Wan, Ping Luo, Min Lin, and Ngai Wong. 2024. [Scaling laws with vocabulary: Larger models deserve larger vocabularies](#). *arXiv preprint*, arXiv:2407.13623.

- Yi Tay, Vinh Q Tran, Sebastian Ruder, Jai Gupta, Hyung Won Chung, Dara Bahri, Zhen Qin, Simon Baumgartner, Cong Yu, and Donald Metzler. 2021. [Charformer: Fast character transformers via gradient-based subword tokenization](#). *arXiv preprint*, arXiv:2106.12672.
- Avijit Thawani, Jay Pujara, Filip Ilievski, and Pedro Szekely. 2021. [Representing numbers in NLP: a survey and a vision](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 644–656, Online. Association for Computational Linguistics.
- Jörg Tiedemann. 2016. [Finding alternative translations in a large corpus of movie subtitle](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 3518–3522, Portorož, Slovenia. European Language Resources Association (ELRA).
- Iulia Turc, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [Well-read students learn better: On the importance of pre-training compact models](#). *arXiv preprint*, arXiv:1908.08962.
- Omri Uzan, Craig W. Schmidt, Chris Tanner, and Yuval Pinter. 2024. [Greed is all you need: An evaluation of tokenizer inference methods](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 813–822, Bangkok, Thailand. Association for Computational Linguistics.
- Menan Velayuthan and Kengatharaiyer Sarveswaran. 2025. [Egalitarian language representation in language models: It all begins with tokenizers](#). In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 5987–5996, Abu Dhabi, UAE. Association for Computational Linguistics.
- Eric Wallace, Yizhong Wang, Sujian Li, Sameer Singh, and Matt Gardner. 2019. [Do NLP models know numbers? probing numeracy in embeddings](#). 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 5307–5315, Hong Kong, China. Association for Computational Linguistics.
- Mengting Wan and Julian McAuley. 2018. [Item recommendation on monotonic behavior chains](#). In *Proceedings of the 12th ACM conference on recommender systems*, pages 86–94.
- Mengting Wan, Rishabh Misra, Ndapa Nakashole, and Julian McAuley. 2019. [Fine-grained spoiler detection from large-scale review corpora](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2605–2610, Florence, Italy. Association for Computational Linguistics.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.
- Andrew Wang, Cristina Aggazzotti, Rebecca Kotula, Rafael Rivera Soto, Marcus Bishop, and Nicholas Andrews. 2023. [Can authorship representation learning capture stylistic features?](#) *Transactions of the Association for Computational Linguistics*, 11:1416–1431.
- Xiao Wang, Qin Liu, Tao Gui, Qi Zhang, Yicheng Zou, Xin Zhou, Jiacheng Ye, Yongxin Zhang, Rui Zheng, Zexiong Pang, Qinzhuo Wu, Zhengyan Li, Chong Zhang, Ruotian Ma, Zichu Fei, Ruijian Cai, Jun Zhao, Xingwu Hu, Zhiheng Yan, Yiding Tan, Yuan Hu, Qiyuan Bian, Zhihua Liu, Shan Qin, Bolin Zhu, Xiaoyu Xing, Jinlan Fu, Yue Zhang, Minlong Peng, Xiaoqing Zheng, Yaqian Zhou, Zhongyu Wei, Xipeng Qiu, and Xuanjing Huang. 2021. [TextFlint: Unified multilingual robustness evaluation toolkit for natural language processing](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing: System Demonstrations*, pages 347–355, Online. Association for Computational Linguistics.
- A Warstadt. 2019. [Neural network acceptability judgments](#). *arXiv preprint*, arXiv:1805.12471.
- Anna Wegmann, Marijn Schraagen, and Dong Nguyen. 2022. [Same author or just same topic? towards content-independent style representations](#). In *Proceedings of the 7th Workshop on Representation Learning for NLP*, pages 249–268, Dublin, Ireland. Association for Computational Linguistics.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122, New Orleans, Louisiana. Association for Computational Linguistics.
- Linting Xue, Aditya Barua, Noah Constant, Rami Al-Rfou, Sharan Narang, Mihir Kale, Adam Roberts, and Colin Raffel. 2022. [ByT5: Towards a token-free future with pre-trained byte-to-byte models](#). *Transactions of the Association for Computational Linguistics*, 10:291–306.
- Jinbiao Yang, Stefan L. Frank, and Antal van den Bosch. 2020. [Less is better: A cognitively inspired unsupervised model for language segmentation](#). In *Proceedings of the Workshop on the Cognitive Aspects of the Lexicon*, pages 33–45, Online. Association for Computational Linguistics.

- Daoguang Zan, Bei Chen, Fengji Zhang, Dianjie Lu, Bingchao Wu, Bei Guan, Wang Yongji, and Jian-Guang Lou. 2023. [Large language models meet NL2Code: A survey](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7443–7464, Toronto, Canada. Association for Computational Linguistics.
- Rowan Zellers, Ari Holtzman, Hannah Rashkin, Yonatan Bisk, Ali Farhadi, Franziska Roesner, and Yejin Choi. 2019. [Defending against neural fake news](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Rodolfo Zevallos and Nuria Bel. 2023. [Hints on the data for language modeling of synthetic languages with transformers](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12508–12522, Toronto, Canada. Association for Computational Linguistics.
- Shiyue Zhang, Vishrav Chaudhary, Naman Goyal, James Cross, Guillaume Wenzek, Mohit Bansal, and Francisco Guzman. 2022. [How robust is neural machine translation to language imbalance in multilingual tokenizer training?](#) In *Proceedings of the 15th biennial conference of the Association for Machine Translation in the Americas (Volume 1: Research Track)*, pages 97–116, Orlando, USA. Association for Machine Translation in the Americas.
- Jian Zhu and David Jurgens. 2021. [Idiosyncratic but not arbitrary: Learning idiolects in online registers reveals distinctive yet consistent individual styles](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 279–297, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. [Aligning books and movies: Towards story-like visual explanations by watching movies and reading books](#). In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 19–27.
- Caleb Ziems, William Held, Jingfeng Yang, Jwala Dhamala, Rahul Gupta, and Diyi Yang. 2023. [Multi-VALUE: A framework for cross-dialectal English NLP](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 744–768, Toronto, Canada. Association for Computational Linguistics.
- Vilém Zouhar, Clara Meister, Juan Gastaldi, Li Du, Mrinmaya Sachan, and Ryan Cotterell. 2023. [Tokenization and the noiseless channel](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5184–5207, Toronto, Canada. Association for Computational Linguistics.

Group	Choice	Examples of tokens
Corpus	Wikipedia	ebruary _Retrieved arliament
	Twitter	_[joy] [loudly-crying-face][loudly-crying-face] [heart] BTS
	PubMed	_effec _lymphadenopathy
	Misc.	\t\t \r \n _differe
Pre-Tok	no	_that _at _in _this _case_
	ws	you're that's took develo
	_ws	\nI _I'm \nWhat _devices.
	GPT2	_..., ensional >"; _127
Size	LLAMA3	433 _{\n _*\n .apache
	500	! \$ 0 1 2 A B C is he in
	4k	age very _will _would
	32k	_surveillance _Vietnam
	64k	CAN 322 _infuri
	128k	motherboard _narcotics

Table 6: **Examples of unique tokens for each tokenizer choice.** The displayed tokens are unique to the given tokenizer in their respective group, except for the different vocabulary sizes. The tokens of the smaller vocabulary size is always included in the bigger vocabulary size. We represent whitespaces within tokens as `_`. We represent emojis within `[]` with their textual descriptions.

## A Tokenizer

In this section, we provide additional information on the tokenizer settings we investigate. See Table 6 for examples of unique tokens for each setting.

### A.1 Fitting Corpora

See an overview of the fitting dataset sizes in Table 7. Usually, the fitting corpus for the tokenizer and the training corpus for the language model with that tokenizer are the same. As a result the size of the fitting corpus often varies as widely as the size of training datasets. We aim for 1.5 billion tokens for all fitting corpora. The MiniPile (Kaddour, 2023) dataset used for fitting in Schmidt et al. (2024) is of similar magnitude. We further display text examples for each dataset in Table 8. Dataset sizes vary in less than 1% of word count. The variance in word count is an artifact of dataset creation from several documents with lenient word count limits.

**Miscellaneous.** Miscellaneous consists of Reddit (Baumgartner et al., 2020), literature sources (fanfictions from ao3<sup>5</sup>, Gao et al., 2020’s books

<sup>5</sup>fanfictions until 2019 from Archive of Our Own <https://archiveofourown.org/>, downloaded from [https://archive.org/download/A03\\_story\\_dump\\_continuing](https://archive.org/download/A03_story_dump_continuing) in 2023, filtered for English language using AO3 tags. Dataset

Source	Train	Dev	Test
Wikipedia	1,469,999,792	15,000,029	15,000,087
Twitter	1,470,004,662	15,000,048	15,000,057
PubMed	1,469,999,499	15,000,106	15,000,501
Mixed	1,477,872,323	15,080,505	15,080,919

Table 7: **Fitting corpora with similar word counts.** We compare three fitting corpora for tokenizers. Word count is calculated using white-space splitting. The size of the fitting corpora are not exactly the same when it comes to word count. But variations in word count are below 1% and should not affect the vocabulary of the tokenizer fitted on them.

before 1919 from the Gutenberg project), news articles and comments (Zellers et al., 2019’s re- alnews, 2020 NYTimes articles and comments<sup>6</sup>, Kolhatkar et al., 2020’s sfu-soc), question answering (Gao et al., 2020’s StackExchange), reviews (Hou et al., 2024’s Amazon and Wan and McAuley, 2018’s GoodReads reviews), mails (Bevendorff et al., 2020’s Gmane), transcripts (YouTubeCom- mons<sup>7</sup> and Gao et al., 2020’s OpenSubtitles), blogs (Schler et al., 2006’s blogcorpus), raw text from webpages (Gao et al., 2020’s Common Crawl), sci- ence articles (Lo et al., 2020’s s2orc), code and mathematics (Gao et al., 2020’s GitHub and Deep- Mind Mathematics). See the share of different domains in Table 9.

### A.2 Pre-Tokenizer

See the regular expressions defining the different considered pre-tokenizers in Table 10. Differences affect mostly whitespace, contraction, punctuation and number handling.

was removed but should be re-creatable using tools like <https://github.com/nianeyna/ao3downloader>.

<sup>6</sup><https://www.kaggle.com/datasets/benjaminawd/new-york-times-articles-comments-2020>, minimum length filter of 250 was applied

<sup>7</sup><https://huggingface.co/datasets/PleIAs/YouTube-Commons>



Source	Text	word count	domain
Wiki	Mary Jane Christie Serrano (c. 1840 – 1923) was a writer, poet and considered ...	24	-
Twitter	Where are the top places in Broward or Palm Beach? [thinking-face][eyes]	10	-
PMed	... Myoelectrical activity of the gut has been studied in the postoperative period ...	132	-
Misc.	Israel, as usual, wants American forces to fight a bloody war against Iran. ...	66	nytimes
	&gt;In Israel, my grandfather fought for its life. The people down the street fought	581	reddit
	Q:\n\nHow can I determine the current focused process name and version in C#...	122	StackEx.
	I read the audio version of this story and loved it. ...	53	goodreads
	... always get 100 test cases (or whatever the default number of test is)?\n\nJanek ...	134	gmane
	there is Sydney waiting to enter. Their eyes meet. "Maggie! I was just ...	998	ao3
	... Abstract. How many n-orthants can be intersected in the n-dimensional ...	2048	s2orc
	from torch import optim as optim\n\n from geoopt.optim.mixin import ...	115	GitHub
	if you're looking for enhancement cores in the game one really useful way ...	102	YouTube
	I have a lot of ties. But my favorite one. My favorite tie is owned by ...	183	blogcorpus
	Witnesses interviewed: 3 (N°1141, 1142, 1143). Nikolay S.( N°1142): "After the	72	Pile-CC
	...		
	It came in perfect condition and it is very soft.	9	amazon
	I'm sure that the procurement people are doing the best job they can ...	94	sfu-socc
	... Sort -11, -1, 0, -3, 5 in decreasing order ...	1645	DM Maths
	... "But the killer isn't the Russian army." "It's the subzero temperatures." ...	1978	OpenSubtitles
	... THE\n\n LIFE\n\n OF\n\n GEORGE WASHINGTON,\n\n COMMANDER IN	2048	Gutenberg
	...		
	Rumson, NJ – December 2013 What started in 2003 as a group of mostly Christian	592	realnews
	...		

Table 8: **Dataset examples.** We show text examples for all used fitting corpora.

Genre	Domain	Train	Dev	Test
Forum	Reddit	245M	2.5M	2.5M
Literature	AO3	147M	1.5M	1.5M
Literature	Gutenberg before 1919	49M	0.5M	0.5M
News	Realnews	147M	1.5M	1.5M
News/Comments	NYTimes & Comments	24M	0.3M	0.3M
News/Comments	SFU-SOCC	3M	0.03M	0.02M
Q&A	StackExchange	196M	2.0M	2.0M
Reviews	Goodreads	49M	0.5M	0.5M
Reviews	Amazon	49M	0.5M	0.5M
Mails	Gmane	147M	1.5M	1.5M
Transcripts	YouTubeCommons	98M	0.9M	1.0M
Transcripts	OpenSubtitles	49M	0.5M	0.5M
Code	GitHub	49M	0.5M	0.5M
Science	S2ORC	98M	1.0M	0.9M
Blogs	BlogCorpus	10M	0.1M	0.1M
Raw Text Webpages	CommonCrawl	98M	1.0M	1.0M
Mathematics	DM Mathematics	20M	0.2M	0.2M
Total:		1,478M	15.1M	15.1M

Table 9: **Miscellaneous dataset statistics.**

name	RegEx	Example Text
NO	-	well..._\$3000_for_a_tokenizer_isn't_cheapz_#lol_ :) \n \n https://en.wikipedia.org/wiki/Sarcasm
WS	\s+	well... _ \$3000 _ for _ a _ tokenizer _ isn't _ cheapz _ #lol _ :)\n \n https://en.wikipedia.org/wiki/Sarcasm
_WS	\s+(?!\\S) \\s+	well... _ \$3000 _ for _ a _ tokenizer _ isn't _ cheapz _ #lol _ :) \n \n https://en.wikipedia.org/wiki/Sarcasm
GPT2	's 't 're 've 'm 'll 'd   ?\\p{L}+  ?\\p{N}+   ?[^\\s\\p{L}\\p{N}]+  \\s+(?!\\S) \\s+	well ... _\$ 3000 _for _a _tokenizer _isn ` t _cheapz _# lol _:) \n \n https :// en . wikipedia . org / wiki / Sarcasm
LLAMA3	(?i:'s 't 're 've 'm 'll 'd)   [^\\r\\n\\p{L}\\p{N}]?\\p{L}+   \\p{N}{1,3}   ?[^\\s\\p{L}\\p{N}]+[\\r\\n]*   \\s*[\\r\\n]+ \\s+(?!\\S) \\s+	well ... _\$ 300 0 _for _a _tokenizer _isn ` t _cheapz _# lol _:) \n \n https :// en . wikipedia .org /wiki /Sarcasm

**Table 10: Investigated Pre-tokenizers.** The pre-tokenizers we investigate can be described with regular expressions. We investigate using no pre-tokenizer (no), isolating whitespaces (ws), split on whitespaces including single leading whitespaces in non-whitespace tokens (\_ws), the pre-tokenizer used by GPT-2 (gpt2) and the pre-tokenizer used by Llama 3 (llama3). GPT-2 and Llama 3 mainly differ in contraction, URL, whitespace and number handling. We display how the investigated pre-tokenizer split an example text. We replace whitespaces with \_ to highlight pre-token boundaries with whitespace.

## B Evaluation Tasks

### B.1 Tasks Robust To Language Variation

**GLUE task selection.** We originally planned to pre-train BERT models and test them on the same GLUE tasks (Wang et al., 2018) as the ones used for the original BERT model (Devlin et al., 2019), i.e., CoLA, SST-2, MRPC, STS-B, QQP, MNLI, QNLI and RTE. We removed CoLA (Warstadt, 2019) as it is the task of classifying linguistic acceptability. Models have to classify morphological, syntactic and semantic “violations” and the task can thus be expected to be sensitive to language variation. Further, we removed STS-B (Cer et al., 2017). STS-B is a regression task for the semantic similarity between two sentences. However, we aimed to focus on classification tasks that could be modeled with logistic regression. After pre-training and fine-tuning on the remaining GLUE tasks, we also removed MRPC (Dolan and Brockett, 2005) and RTE from the evaluated tasks. We removed RTE because the standard deviation of .04 made the differences in performance for models unclear, the variation could be related to the small training dataset of only 2.5k instances. We further removed MRPC as it showed almost no variation between different tokenizers – not even for the 500 vocabulary size. All remaining tasks had a train set size of at least 67k. Note that SST-2 was released in a parsed format resulting in a lowercased and pre-tokenized text which might affect results.

**GLUE+dialect.** We transform GLUE using Multi-VALUE (Ziems et al., 2023). Multi-VALUE has strong requirements on text formatting. When Multi-VALUE perturbations fail we leave the GLUE text in that row as is. Depending on the task this concerns between 3% and 18% of instances.

### B.2 Tasks Requiring Sensitivity To Language Variation

See an example for each of the tasks requiring sensitivity to language variation in Table 11.

**Authorship Verification.** Similar to the Miscellaneous corpus A.1, the authorship verification corpus consists of datasets taken from Amazon (Hou et al., 2024), AO3, GMane (Bevendorff et al., 2020), 2020 NYTimes articles and comments, realnews (Zellers et al., 2019), Reddit (Baumgartner et al., 2020), StackExchange (Gao et al., 2020) and Wikipedia articles. Additionally, the corpus includes texts from Bookcorpus (Zhu et al., 2015) and PubMed (Gao et al., 2020). It totals about

40.8k train pairs, 2.5k dev pairs and 4.8 test pairs. For the contrastive training task, we use the Supervised Contrastive Loss<sup>8</sup> (Khosla et al., 2020) with a siamese setup, a batch size of 128, and a learning rate of 0.00001. We find the threshold best separating same author and distinct author pairs on the dev set and report accuracy on the test set. We release the dataset on Hugging Face.<sup>9</sup>

**PAN.** The dataset was extracted from Reddit and preprocessed by removing citations, markdown, emojis, hyperlinks, multiple line breaks and extra whitespace (Bevendorff et al., 2024). Compared to the Authorship Verification task (where the classifiers may learn to rely on content cues, apostrophe encodings, whitespace encoding, etc., cf. Wegmann et al., 2022), this may be more difficult. The dataset was downloaded from <https://pan.webis.de/clef24/pan24-web/style-change-detection.html>.

**NUCLE.** The dataset was downloaded from <https://www.comp.nus.edu.sg/~nlp/corpora.html>.

**CORE.** We use 8 main register labels for multi-class prediction. We display an example for all considered CORE (Laippala et al., 2023) labels in Table 12. The original CORE consists of main as well as sub-labels that make up a total of 56 labels in Laippala et al. (2023). However using all 56 in a multi-label setup proved too difficult for our BERT as well as the logistic regression models without further hyperparameter tuning. We decided for a multi-classification setup, limiting ourselves to 8 out of 9 main register labels, specifically we excluded the “OTHER” category. This reduced the train dataset from about 34k instances to 30k. We split up texts of length > 250 to chunks of a maximum of 250 using whitespace splitting. Then, we perform stratified sampling with replacement for each class to additionally upweigh small classes. Note that this results in duplicates for the small classes up to a maximum of 10 occurrences of the same text.

<sup>8</sup>SupConLoss as implemented in `pytorch_metric_learning`, see <https://kevinmusgrave.github.io/pytorch-metric-learning/losses/>

<sup>9</sup>Future HuggingFace Link.

Name	Task	Source	Text 1	Text 2	label
CORE	Register Classification	(Laippala et al., 2023)	[...] What do you do when you just cant seem to mix something? Hi, You have mixed it...you just didn't know when to stop and move on! [...]	-	Interactive Discussion
AV	Authorship Verification	our contribution	Hi,\n\nI am currently evaluating OTRS for use as a Helpdesk System.\n\nHowever I am a little confused about how best to set it up. [...]	Hi,\n\nI am setting up Mailscanner not to deliver any spam to my users, but to\nsend a spam report once a day [...]	Same Author
PAN	Author Change	(Bevendorff et al., 2024)	I'm not gonna watch the video. I gotta keep my sanity. With that being said, what could we call for that they haven't done? The cops have been fired and charged with murder. [...]	Maybe police agencies should be not federalized but under one agency. No more sherriff's no more local police. Just state police. [...]	Author Change
NUCLE	Error Classifications	(Dahlmeier et al., 2013)	Chernobyl accident, happened in 1986, was a nuclear reactor accident.	-	[ArtOrDet, Trans]
Dialect	Dialect Classification	transformations with (Ziems et al., 2023)	What the best things to do in Hong Kong one?	-	CollSgE

Table 11: **Tasks Sensitive to Language Variation.** For each task, we show an example.

Text	label
... Sometimes, people just don't feel well. But if you don't feel well more than sometimes, it may be helpful to talk to someone about it. ...	opinion
... A transportation advocacy group is circulating a list of 100 questions aimed at broadening the British Columbia government's consultation on coastal ferry services. ...	narrative
... I'm sure many people have hit this brick wall. What do you do when you just cant seem to mix something? Hi, You have mixed it...you just didn't know when to stop and move on! ...	interactive discussion
'Always think of home': an introduction to the Buenos Ayres Notebook The Buenos Ayres Notebook takes its name from the city Buenos Aires ('Good Air' or 'Fair Winds'), ...	informational description/explanation
... It would be a pleasure just to know just a little bit moreoh oh I could grow quite fond of your acquaintance ...	lyrical
An Acadian-style cabin constructed completely of rough sawed Southern Yellow Pine, surrounded by split rail fence. ...	informational persuasion
... it can be easy or even enjoyable. Here is a guide on how to give an oral presentation in front of your class. Decide on a topic ...	how-to/instructional
... Kareem Ettouney, the art director at Media Molecule always said, Mash up, not mish mash! ——— What kinds of challenges did you face with replicating LBP's iconic 2D puppet aesthetic into a 3D space? ...	spoken

Table 12: **Examples for main CORE labels.** We focus on the 8 main CORE (Laippala et al., 2023) labels.



GPU h	P	T	loss	steps	batch
0h	4.6M	10M	7.7	600	32
1h	42M	100M	6.5	6k	32
4h	110M	250M	6.2	15k	32
9.0	110M	750M	3.9	11k	128
11.0	110M	750M	<b>2.7</b>	45k	32
12.3	4.6M	330M	4.1	80k	256
13.2	110M	750M	3.2	14k	128
13.6	4.6M	3300M	4.1	75k	256
14.8	4.6M	3300M	4.1	80k	256
22.0	11.6M	3300M	3.2	75k	256

Table 13: **Hyperparameters for BERT Pre-Training.** We compare the evaluation loss, and GPU hours while varying the number of parameters, tokens, steps and batch size. For similar GPU hours (between 11h-15h), using more pre-training tokens does not seem to improve performance as much as increasing model size. Balancing the number of model parameters and tokens, as well as the number of steps seems crucial.

## C Modeling

**Compute Optimal BERT.** We originally evaluated tokenizers on tiny BERT models, using 80k steps on a training dataset with 3.3B tokens during pre-training (second to last row in Table 13). This corresponds to more than three epochs on the relatively large training dataset for a tiny BERT model with only 4.6M parameters. Using this setup we found that tokenizers with the largest vocabulary sizes repeatedly outperformed all other settings. For tiny BERT, models with larger vocabulary sizes also need to use orders of magnitude more parameters because of the larger embedding matrix. For tiny BERT, the model size rises to 17M for a vocabulary size of 128k. Were we using a non-optimal ratio between number of tokens and parameters for our invested 15 GPU hours? We experimented with different BERT model sizes P (tiny – 4.6M, small – 11.6M, base – 110M), number of tokens T, batch sizes and number of steps. See the results in Table 13. We evaluate performance with the eval loss on the held out set of the pretraining corpus consisting equally of Bookcorpus3 (Zhu et al., 2015) and OpenWebText2 (Gao et al., 2020). We use one tokenizer fit on the Miscellaneous corpus, using the gpt-2 pre-tokenizer and a vocabulary size of 32k tokens. For further hyperparameter choices, see the Hyperparameters paragraph. We make the following observations: For similar GPU hours (between 11h-15h), using exponentially more training corpus tokens T than parameters P for tiny BERT does not improve performance as much as increasing model size P. This and other model results might hint

Source	Word Count	Excerpt
Book-Corpus	1,687,724,544	... visit and they all swore on a second blood oath that it wasn't them.\n\n“‘What about Phantom?’” Hell Girl asked. ...
OpenWeb-Text2	1,650,000,384	... The future of SA-affiliated club sports, a cappella, and Greek groups is uncertain after the All-Campus Judicial Council ruled Friday ...
Total	3,337,724,928	

Table 14: **Pretraining Corpora.** Word count is calculated using white-space splitting.

at Chinchilla’s scaling law, that is, optimal token count scaling with the parameter size of a model for a fixed compute, specifically  $T_{OPT} \approx P^{27/23}$  (Hoffmann et al., 2024), also holding for smaller encoder models. Further, for our low resource setting, using more steps, and thus weight updates, seems to be more important than using a large batch size. These conclusions seems promising but very tentative. Exhaustive pretraining experiments were out of scope for this work. Nevertheless, we think that finding compute-optimal settings to train small transformer models is crucial to efficiently evaluate tokenizers going forward.

**WebBook Corpus.** Our training consists of equal parts Bookcorpus3<sup>10</sup> (Zhu et al., 2015) and OpenWebText2 (Gao et al., 2020). See the statistics in Table 14. We randomly sample sequences of 512 words, totaling 3.3 billion words as in the original BERT paper (Devlin et al., 2019). We ensure English excerpts by removing books and web text that are not predicted as English language by using langdetect.<sup>11</sup> We split of 2% of the sampled data for held out dev and test sets to evaluate BERT pretraining.

**Hyperparameters.** We keep the following pre-training settings the same as BERT (Devlin et al., 2019): Adam with learning rate of 1e-4,  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ , L2 weight decay of 0.01, learning rate warmup over the first 1% of steps, linear decay of the learning rate, dropout probability of 0.1 on all layers. However, we do not use the next sentence prediction objective and only train on MLM with mlm probability at 15% as NSP proved to be in-

<sup>10</sup><https://twitter.com/theshawwn/status/1320282149329784833>

<sup>11</sup><https://github.com/Mimino666/langdetect>

ferior to MLM in later models (Liu et al., 2019). Originally, pre-training was performed on a set of 3.3 billion words over 40 epochs, we experiment with different number of tokens and steps in §5. We use the architecture of the tiny, small, medium and base BERT model (Turc et al., 2019) which consists of 4.6, 11.6, 42, and 110 million parameters respectively.<sup>12</sup>

## D Intrinsic Evaluation

**Logistic regression.** We tokenize the task texts. Then, we use the resulting tokens as features for the logistic regression. We use a bag-of-tokens approach and do not consider word order. We do not consider frequency but use binary information on token presence. We use Cartesian products of tokens for NLI tasks. Specifically, given two sentences in an NLI setting, we generate all possible combinations of individual tokens from both sentences and use these as features. For PAN and Authorship Verification, we limit the Cartesian product to tokens that appear in both sentences and tokens that only appear in one sentence.<sup>13</sup> For multi-classification and multi-label task we train separate one-vs-rest logistic regression models. We use l1-regularization with  $C=0.4$ . Note that one could probably increase logistic regression performance by tuning parameters and features more specifically to the investigated tasks.

**Rényi Efficiency.** We record some observations: Using NO pre-tokenizer consistently got the highest Rényi Efficiency and the ws pre-tokenizer consistently got the lowest Rényi Efficiency values, even though both were usually not among the best performing pre-tokenizers in the downstream tasks.

**Varying Vocabulary Size.** Corpus Token Count is sensitive to the vocabulary size of the tokenizer. A tokenizer with a vocabulary size of 128k will almost always have a lower Corpus Token Count than a tokenizer with the vocabulary size of 32k. Independent of how many tokens might be a better fit for a given corpus or task. Similarly, Rényi efficiency (Zouhar et al., 2023) assumes the same vocabulary size to make the efficiency values comparable across tokenizers. Note that even when tokenizers have the same vocabulary size, the vocabulary coverage on the downstream corpus (i.e., the actual number of tokens that appear in the down-

stream corpus) might be smaller. For example, a tokenizer that was fitted on the Twitter corpus might include vocabulary that never appears in the original GLUE tasks. As a result Corpus Token Count and Rényi efficiency might be skewed for tokenizers that have a very low overlap in vocabulary with the downstream task corpus.

<sup>12</sup>This might affect performances for models with especially big vocabulary sizes.

<sup>13</sup>PAN and AV had considerably longer texts than the GLUE NLI tasks leading to otherwise too many features.

## E Modeling Results on Diverging Pre-training and Fitting Corpus

We experimented with diverging pre-training and fitting corpora. Specifically, for the pre-training corpus, we sampled 750 words from **WebBook** (§C). We expect WebBook corpus to show less spelling and syntactic variation than Miscellaneous used in the main paper. For the fitting corpus, we compared using PubMed, Wikipedia, Twitter and Miscellaneous. Further, we experimented with fitting on the Miscellaneous corpus but varying the pre-tokenizer and vocabulary size. See results in Table 15, Table 16 and Figure 3. In contrast to the results in the main paper, the performances are not averaged over three seeds and should be considered preliminary.

For mismatched pre-training and fitting corpus, the vocabulary size needs to be higher for tasks requiring robustness to language variation. This is intuitive as more tokens during tokenizer fitting increase the likelihood of overlap with the pre-training corpus.

Further, the pre-tokenizers NO, WS and \_WS seem to have more trouble leveraging their tokens having been fitted on a corpus different from the pre-training corpus. This is intuitive as their tokens can be expected to be especially dependent on the fitting corpus. This is probably also the explanation for LLAMA3 performing worse than GPT2 for tasks requiring sensitivity to language variation. GPT2 separates Unicode character categories the most and might thus have the largest overlap with the pre-training corpus.

For mismatched pre-training and fitting corpus, the choice of fitting corpus seems more influential for tasks requiring robustness to language variation. Potentially the lack of overlap between the fitting and the pre-training corpus change what tokens can be leveraged effectively. Wikipedia performs surprisingly well for tasks that require sensitivity to language variation. Wikipedia is a corpus using a very standardized version of English. We theorize that this might help with recognizing deviations from that norm.

	Model	org	+typo	+dialect	AVG
Corpus	PMed	81.3	69.4	79.2	76.6
	Wiki	82.1	68.9	79.6	76.8
	Twitter	<b>82.4</b>	<b>70.4</b>	<b>80.5</b>	<b>77.8</b>
	Misc.	81.6	69.2	79.9	76.9
Pre-Tok	NO	72.3	61.6	70.9	68.3
	WS	79.5	66.9	78.1	74.8
	_WS	80.7	68.6	79.0	76.1
	LLAMA3	<b>82.2</b>	<b>69.4</b>	79.7	<b>77.1</b>
	GPT2	81.6	69.2	<b>79.9</b>	76.9
Size	500	79.6	<b>72.8</b>	78.1	76.8
	4k	80.9	70.8	78.9	76.8
	32k	81.6	69.2	79.9	76.9
	64k	<b>82.5</b>	69.3	<b>80.7</b>	<b>77.5</b>
	128k	81.9	68.5	79.9	76.7

Table 15: **Mismatched Pre-training and Fitting Corpus on Tasks Robust to Language Variation.** We use the WebBook pre-training corpus and fit on the Miscellaneous, PubMed, Wikipedia and Twitter corpora.




Figure 3: **Pairwise significance testing for tasks robust and sensitive to language variation with Mismatched Pre-training and Fitting Corpus.** We use the WebBook pre-training corpus and fit on the Miscellaneous, PubMed, Wikipedia and Twitter corpora. We use (McNemar, 1947)’s test to test how different the BERT models trained with different tokenizers classify the tasks robust to language variation (first row) and tasks sensitive to language variation (second row). Tokenizers are sorted by mean performance. Blue colors show statistical significance, while red colors are above the 0.05 threshold.

	Model	AV (acc $\uparrow$ )	PAN (acc $\uparrow$ )	CORE (acc $\uparrow$ )	NUCLE (F1 $\uparrow$ )	DIAL (F1 $\uparrow$ )	Mean
Fitting Corpus	PubMed	80.6	65.5	55.6	23.8	88.3	62.8
	Wikipedia	80.9	67.2	56.8	<b>24.8</b>	88.7	63.7
	Twitter	80.8	64.7	57.2	23.5	<b>89.1</b>	63.0
	Misc.	<b>82.0</b>	<b>68.3</b>	<b>57.9</b>	24.1	89.0	<b>64.3</b>
Pre-Tokenizer	NO	79.8	52.4	51.5	16.9	76.6	55.4
	WS	74.9	65.2	55.7	16.6	85.3	59.5
	_WS	80.4	65.2	56.8	22.1	87.9	62.5
	LLAMA3	81.3	66.9	56.8	23.4	<b>89.0</b>	63.5
	GPT2	<b>82.0</b>	<b>68.3</b>	<b>57.9</b>	<b>24.1</b>	<b>89.0</b>	<b>64.3</b>
Vocabulary Size	500	77.6	62.7	53.2	15.0	86.6	59.0
	4k	80.3	65.1	56.0	21.3	88.1	62.2
	32k	82.0	<b>68.3</b>	57.9	24.1	89.0	<b>64.3</b>
	64k	82.1	67.2	<b>58.0</b>	<b>25.4</b>	88.9	<b>64.3</b>
	128k	<b>82.6</b>	66.6	54.3	24.1	<b>89.5</b>	63.4

Table 16: **Mismatched Pre-training and Fitting Corpus on Tasks Sensitive to Language Variation.** We use the WebBook pre-training corpus and fit on the Miscellaneous, PubMed, Wikipedia and Twitter corpora.



## F Intended Use and Licenses for used Datasets

We discuss intended use and licenses for the datasets we re-used and created for this work.

### F.1 Datasets curated from different sources

**The Pile.** We include several datasets extracted from Gao et al. (2020)’s The Pile. The Pile consists of newly collected datasets as well as datasets from other sources. There is no license information included in the paper or original data release, but the Pile is described as “open source language modelling data”<sup>14</sup> and, even though not explicitly stated, the intended use should be for open source language modelling and research. We use the Pile to access datasets from originally other sources: books from the 1919 Gutenberg project (Rae et al., 2019)<sup>15</sup>, StackExchange<sup>16</sup>, OpenSubtitles (Tiedemann, 2016)<sup>17</sup>, Common Crawl<sup>18</sup> and DeepMind<sup>19</sup> (Saxton et al., 2019). We further use the following datasets collected by the authors of the Pile: GitHub, OpenWebText2 and PubMed<sup>20</sup>.

**Reddit.** We use dataset originally downloaded from Pushshift (Baumgartner et al., 2020). While the original release was public, and used in many research publications, Reddit updated their terms and the original Pushshift releases are not publicly accessible anymore. Reddit mentions at least partial support of academic research after agreeing to their terms of service.<sup>21</sup>

**Ao3.** We downloaded a public release of Archive of Our Own from [https://archive.org/download/AO3\\_story\\_dump\\_continuing](https://archive.org/download/AO3_story_dump_continuing) in 2023. It did not include license nor intended use descriptions. The dataset was removed by now but should be re-creatable using tools like <https://github.com/nianeyna/ao3downloader>. Note that the license situation remains unclear. AO3 has a complex license where authors retain their rights and the website

is granted a ‘a world-wide, royalty-free, nonexclusive license to make your Content available’. AO3 terms of service forbids use of fanfictions for commercial generative AI.<sup>22</sup>

**Amazon Reviews.** The Amazon Reviews dataset (Hou et al., 2024) was downloaded from [https://cseweb.ucsd.edu/~jmcauley/datasets/amazon\\_v2/](https://cseweb.ucsd.edu/~jmcauley/datasets/amazon_v2/). While the dataset is publicly available, the license for the data remains with Amazon but the customers who wrote the reviews retain the copyright. There is no general site that provides guidance on the license and constraints for this data when used in the academic or research space. The guidelines for Amazon Services are noted here <https://www.amazon.com/gp/help/customer/display.html?nodeId=508088>

**GoodReads.** GoodReads was publicly released with Wan and McAuley (2018); Wan et al. (2019) for academic use. We downloaded it through <https://mengtingwan.github.io/data/goodreads>. The Goodreads license is available at <https://www.goodreads.com/about/terms>. The license includes the text ‘This license does not include any resale or commercial use of any part of the Service, or its contents; any collection and use of any book listings, descriptions, reviews or other material included in the Service; any derivative use of any part of the Service or its contents; any downloading, copying, or other use of account information for the benefit of any third party; or any use of data mining, robots, or similar data gathering and extraction tools.’

**GMane.** Public mailing list emails collected from the gmane.io server, available at <https://webis.de/data/webis-gmane-19.html>. Released with Bevendorff et al. (2020). Accessed through <https://zenodo.org/records/3766985> after submitting a request. No license stated but not publicly available without request. Terms of use are documented on the dataset website <https://zenodo.org/records/3766985>.

**Blogcorpus.** Released with Schler et al. (2006). The paper does not discuss license or intended use. Accessed through <https://www.kaggle.com/datasets/ratatman/blog-authorship-corpus>. Schler et al. (2006) downloaded from <https://www.blogger.com/>.

**Bookcorpus3.** Originally released with Zhu et al. (2015). Bandy and Vincent (2021) released a retrospective datasheet. Zhu et al. (2015) did not

<sup>14</sup><https://pile.eleuther.ai/>

<sup>15</sup>Project Gutenberg consists mostly of public US ebooks, see <https://www.gutenberg.org/policy/permission.html>

<sup>16</sup>anonymized data shared with a cc-by-sa 4.0 license, see <https://archive.org/details/stackexchange>

<sup>17</sup>No licensing or intended use information included. Originally extracted from <https://www.opensubtitles.org/>.

<sup>18</sup>see Terms of Use: <https://commoncrawl.org/terms-of-use>

<sup>19</sup>No licensing information included in the paper.

<sup>20</sup>See terms of the original dataset here: <https://ftp.ncbi.nlm.nih.gov/pubmed/baseline/README.txt>

<sup>21</sup>Developer Platform & Accessing Reddit Data

<sup>22</sup>[https://archiveofourown.org/tos\\_faq](https://archiveofourown.org/tos_faq)

discuss intended use or licenses. However, the license for the data can be expected to remain with the original book copyright holders, except in cases where the copyright has expired.

**NYTimes.** The dataset is publicly available at <https://www.kaggle.com/datasets/benjaminawd/new-york-times-articles-comments-2020> shared with a CC BY-NC-SA 4.0 license. However, in all likelihood the license still belongs to the NYTimes while the copyright remains with the commenter. Some details are available here <https://help.nytimes.com/hc/en-us/articles/360039332111-The-New-York-Times-Content-Agreement>.

**Realnews.** Published with Zellers et al. (2019). Downloaded from <https://github.com/rowanz/grover/tree/master/realnews>. License can be found at this [Google Docs Form](#). It is intended only for research and education use and can not be distributed.

**SFU-Socc.** Released with Kolhatkar et al. (2020). Downloaded from <https://github.com/sfu-discourse-lab/SOCC>. Shared with Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License.

**s2orc.** Released with Lo et al. (2020) mentioning research and development as intended use. Downloaded through <https://github.com/allenai/s2orc/?tab=readme-ov-file>. License is given as ODC-By 1.0.

**YouTubeCommons.** Downloaded through <https://huggingface.co/datasets/PleIAs/YouTube-Commons>. Released with CC-By license. All transcripts are part of a video shared under a CC-By license.

**GLUE.** Released a collection of tasks with Wang et al. (2018). Downloaded from <https://huggingface.co/datasets/nyu-mll/glue>. GLUE is a common public dataset used to evaluate language models. We use MNLI (Williams et al., 2018), SST-2 (Socher et al., 2013), QQP and QNLI.

**PAN.** This PAN 2024 dataset was extracted from Reddit (Bevendorff et al., 2024). We downloaded the dataset from <https://zenodo.org/records/10677876>. Reddit’s terms of use might apply.

**NUCLE.** We use the NUCLE 3.3 corpus (Dahlmeier et al., 2013), downloaded from <https://www.comp.nus.edu.sg/~nlp/corpora.html> after submitting a request. It is available for research purposes. License information can be found at <https://sterling8.d2.comp.nus.edu.sg/>

[nucle\\_download/nucle.php](#) and does not allow for distribution of the corpus.

**CORE.** Released with (Laippala et al., 2023), downloaded from <https://github.com/TurkuNLP/CORE-corpus>. It is released with a CC BY-SA 4.0 license.

## F.2 Datasets collected by us

**Twitter.** Sampled in 2023 with Twitter research API access using the Decahose sampling stream. Distribution of tweet texts was not granted under research access. Intended use was academic research.

## G Personally Identifying Info Or Offensive Content in Datasets

Some of the used datasets can be expected to include personally identifying information or offensive content. We did not take steps to remove identifiable cues or offensive content. This was out of scope for the extensive amount of datasets used. We hope that the effect is negligible as for all datasets, except for Twitter, datasets were already publicly accessible. We acknowledge that re-distributing it might, however, make it more widely accessible. We do not release the Twitter dataset publicly.

## H Model Size and Budget

We used single A100s to run modeling. We pre-trained 24 distinct BERT models for our main experiments (taking less than 360 GPU hours), and fine-tuned each model for all evaluation tasks ( $\approx 24 * (6h * 3 \text{ [GLUE tasks]} + 3h \text{ [tasks requiring sensitivity language variation]}) = 24 * 21h = 504h$ ).

## I Use of AI Assistants

We used ChatGPT and GitHub Copilot for coding, to look up commands and sporadically to generate individual functions. Generated functions were tested w.r.t. expected behavior. We used AI assistants for rephrasing and grammatical error correction in writing.