

Biomedical Language Models are Robust to Sub-optimal Tokenization

Bernal Jiménez Gutiérrez Huan Sun Yu Su

The Ohio State University

{jimenezgutierrez.1, sun.397, su.809}@osu.edu

Abstract

As opposed to general English, many concepts in biomedical terminology have been designed in recent history by biomedical professionals with the goal of being precise and concise. This is often achieved by concatenating meaningful biomedical morphemes to create new semantic units. Nevertheless, most modern biomedical language models (LMs) are pre-trained using standard domain-specific tokenizers derived from large scale biomedical corpus statistics without explicitly leveraging the agglutinating nature of biomedical language. In this work, we first find that standard open-domain and biomedical tokenizers are largely unable to segment biomedical terms into meaningful components. Therefore, we hypothesize that using a tokenizer which segments biomedical terminology more accurately would enable biomedical LMs to improve their performance on downstream biomedical NLP tasks, especially ones which involve biomedical terms directly such as named entity recognition (NER) and entity linking. Surprisingly, we find that pre-training a biomedical LM using a more accurate biomedical tokenizer does not improve the entity representation quality of a language model as measured by several intrinsic and extrinsic measures such as masked language modeling prediction (MLM) accuracy as well as NER and entity linking performance. These quantitative findings, along with a case study which explores entity representation quality more directly, suggest that the biomedical pre-training process is quite robust to instances of sub-optimal tokenization.¹

1 Introduction

In order to communicate complex concepts precisely and efficiently, biomedical terminology has been designed by researchers and medical professionals by combining existing meaningful mor-

¹Our pre-trained model BioVocabBERT is publicly available at <https://huggingface.co/osunlp/BioVocabBERT>.

Ideal Tokenization	BERT Tokenization	PubMedBERT Tokenization
nephro-pathy	ne-ph-rop-athy	nephropathy
nephro-ectomy	ne-ph-re-ct-omy	nephrectomy
nephro-blastoma	ne-ph-ro-bla-sto-ma	nephroblastoma
nephro-calcin-osis	ne-ph-ro-cal-cino-sis	nephrocalcin-osis

Table 1: Sub-optimally tokenized biomedical terms containing the ‘nephro’ morpheme illustrate the limitations of current tokenization methods.

phemes to create new concepts. Many biomedical terms use general rules to combine meaningful morphemes taken from Greek and Latin (Banay, 1948). For example, these morphemes often have vowels that can be omitted such as the ‘-o-’ in ‘nephro’ from Table 1. This ‘-o-’ acts as a joint-stem to connect two consonantal roots (e.g. ‘nephro-’ + ‘-o-’ + ‘-pathy’ = ‘nephropathy’), but the ‘-o-’ is often dropped when connecting to a vowel-stem (e.g. ‘nephro-’ + ‘-ectomy’ = ‘nephrectomy’, instead of ‘nephro-ectomy’). Students in biomedical fields often learn the meaning of these elements as well as the word formation rules to be able to infer the meaning of unfamiliar words and recall complex terms more easily.²

Even though the agglutinating nature of biomedical terminology is well known, none of the existing pre-trained language models consider this information explicitly when building their tokenizers. As shown in Table 1, frequent words such as ‘nephropathy’ and ‘nephrectomy’ are tokenized by BERT (Devlin et al., 2019) into meaningless subwords (‘ne-ph-rop-athy’ and ‘ne-ph-re-ct-omy’) while remaining as whole words for PubMedBERT (Gu et al., 2021). For more infrequent but still important medical terms like ‘nephroblastomas’ and ‘nephrocalcinosis’, PubMedBERT encodes them as ‘nephroblastoma’ and ‘nephrocalcin-osis’. We

²The existence of popular books such as Collins (2007) emphasize the importance of understanding biomedical terminology design for medical professionals.

argue that there are more meaningful and efficient ways to tokenize both frequent and infrequent medical terms using meaningful morphemes like ‘nephro’ (of a kidney), ‘-pathy/-(o)sis’ (disease), ‘-ectomy’ (surgical procedure), ‘calcin’ (calcification) and ‘blastoma’ (type of cancer) which could help models transfer signal directly into infrequent and even out-of-vocabulary terms.

In this work, we first leverage large-scale morpheme segmentation datasets to more rigorously evaluate current tokenization methods both quantitative and qualitatively. Using the annotated morpheme segmentation dataset from the SIGMORPHON 2022 Shared Task (Batsuren et al., 2022), we are able to determine that current tokenizers such as BERT and the more biomedically relevant PubMedBERT are very poorly aligned with human judgments on morpheme segmentation, even when evaluating on biomedical terminology specifically.

Given that, although the PubMedBERT tokenizer exhibits low performance on biomedical morpheme segmentation, it shows some improvement over BERT’s tokenizer due to its use of biomedical corpus statistics, we hypothesize that using a tokenizer that aligns more strongly with standard biomedical terminology construction for pre-training could achieve improved performance in downstream tasks. In order to verify this hypothesis, we create a new tokenizer, BioVocabBERT, which uses a vocabulary derived from combining a fine-tuned morpheme segmentation model with biomedical domain-knowledge from the Unified Medical Language System (UMLS) (Bodenreider, 2004), a large scale biomedical knowledge base. Subsequently, we leverage BioVocabBERT, which greatly outperforms the PubMedBERT tokenizer on biomedical morpheme segmentation, to pre-train a biomedical language model by the same name and compare its performance with a replicated PubMedBERT model (to control for any potential differences in the pre-training process) on two downstream tasks: named entity recognition (NER) and entity linking.

Surprisingly, we find that the performance of our BioVocabBERT model is remarkably similar to the one obtained by our PubMedBERT replica throughout most datasets tested in fully supervised NER, low-resource NER and zero-shot entity linking. Small improvements arise in low-resource NER and zero-shot entity linking but results are inconsistent across datasets. Additionally, we exam-

ine the model’s robustness to segmentation failures in a small scale case-study which suggests that even if the model’s word embeddings are biased by tokenization errors, the model’s parameters are able to overcome such failures quite successfully. Finally, we measure our models’ language modeling accuracy by word frequencies and find a small word frequency trade-off whose exploration we leave for future work. Given these findings, we conclude that biomedical language model pre-training is quite robust to tokenization decisions which are not well aligned with human judgments, even when dealing with highly agglutinating biomedical terminology.

2 Related Work

2.1 Domain-Specific Pre-training

Recent work on domain-specific language models has demonstrated fairly conclusively that using domain-specific data for pre-training significantly improves language model performance on in-domain downstream tasks. Many different such strategies have been proposed with varying degrees of in-domain vs out-of-domain pre-training data in fields such as biomedicine (Peng et al., 2020; Lee et al., 2019; Gu et al., 2021; El Boukkouri et al., 2022), finance (Wu et al., 2023), law (Chalkidis et al., 2020), scientific research (Maheshwari et al., 2021), clinical practice (Alsentzer et al., 2019) and social media (DeLucia et al., 2022). For biomedical language models specifically, most work agrees that pre-training from scratch using an in-domain corpus, as done by Gu et al. (2021), leads to small but measurable performance improvements over other pre-training strategies.

Apart from introducing pre-training from scratch, Gu et al. (2021) demonstrated the limitations of general domain tokenization for domain-specific pre-training by showing downstream improvements obtained from using a domain-specific tokenizer, created by standard tokenizer building algorithms such as WordPiece (Schuster and Nakajima, 2012) and BPE (Gage, 1994) on an in-domain corpus. In-domain tokenizers have also been shown to improve performance in other domains such as law (Chalkidis et al., 2020) and more specific ones like cancer (Zhou et al., 2022). As a result, most recent biomedical LMs use tokenizers created from in-domain corpora statistics (Yasunaga et al., 2022; Luo et al., 2022).

2.2 Limits of Unsupervised Tokenization

Even though domain-specific tokenizers have become widely used for biomedical LMs, they are still constructed using mainly unsupervised algorithms which leverage information theoretic metrics from large-scale corpora to create subword vocabularies. However, as reported in the SIGMORPHON 2022 Shared Task for morpheme segmentation (Batsuren et al., 2022), these methods align little with morphological human judgments. Hofmann et al. (2021) explore how poor segmentation affects performance by injecting rule-based derivational morphology information into the tokenization process and showing improvements in word classification tasks, especially for low-frequency words. As far as we know, our work is one of the first to perform a similar morpheme segmentation analysis on biomedical tokenizers, even though biomedical terminology is highly agglutinating by design and should benefit from such analysis.

Furthermore, Hofmann et al. (2020) shows that introducing derivational morphology signal into BERT via fine-tuning improves its derivation generation capabilities, suggesting that performance of language models could be improved by adding morphologically relevant signal into their pre-training. Nevertheless, not much work apart from our current study has explored how introducing such signals could affect the pre-training process directly, especially not in biomedical language models.

3 Supervised Morpheme Segmentation

Recent work evaluating morphological segmentation at scale such as the SIGMORPHON 2022 Shared Task (Batsuren et al., 2022) demonstrates the impressive performance of supervised methods compared to unsupervised methods like BPE (Gage, 1994) or Morfessor (Creutz et al., 2005), even for languages with more limited annotated data than English. In the SIGMORPHON 2022 Shared Task, the organizers compile a large quantity of segmented morpheme data, over half a million English words obtained from Wiktionary and other sources using both hand-crafted and automated methods (Batsuren et al., 2021).

3.1 Evaluating Biomedical Segmentation

By comparing the SIGMORPHON dataset with words which appear frequently in the Unified Medical Language System (UMLS) (Bodenreider, 2004), a large scale biomedical knowledge base,

	Train	Dev	Test
English Set	458,692	57,371	57,755
Biomedical Subset	33,221	4,112	4,123

Table 2: Dataset statistics for the SIGMORPHON 2022 morpheme segmentation dataset and the biomedical dataset, as defined in §3.1.

we find that a small percentage (approximately 10%) of all annotated words are relevant biomedical terms. We therefore leverage this biomedical subset to evaluate the biomedical morpheme segmentation performance of several current tokenization methods. Due to the large difference in scale of the full dataset to the biomedical subset, we use the full SIGMORPHON dataset for training, including both general english and biomedical words. We use the same segmentation F1 score the SIGMORPHON Shared Task for evaluation. This score is calculated as the harmonic mean of *precision*, the ratio of correctly predicted morphemes over all predicted morphemes, and *recall*, the ratio of correctly predicted morphemes over all gold-label units. For more information about these evaluation metrics, we refer the interested reader to Section 2.3 of Batsuren et al. (2022).

	Segmentation F1
BERT Tokenizer	16.2
PubMedBERT Tokenizer	19.2
Fine-Tuned CANINE	74.1
BioVocabBERT Tokenizer	48.5

Table 3: Morpheme segmentation performance of baseline and novel tokenizers on the biomedical subset of the SIGMORPHON 2022 development set.

As seen in Table 3, both BERT and PubMedBERT achieve under 20% segmentation F1 performance on the SIGMORPHON biomedical development subset. In order to understand why current tokenizers obtain such dramatically low segmentation accuracy, we analyze 50 instances of sub-optimal tokenization. Apart from words which are not segmented because they exist in the PubMedBERT vocabulary, most errors are split into three main categories 1) missing units, 2) compound units and 3) ambiguous connecting vowels. Table 4 shows descriptions and examples of each type.

3.2 CANINE Fine-Tuning

As opposed to sub-word tokenization, morpheme segmentation does not require sub-word compo-

	Description	Example
Missing Units	Important biomedical morphemes missing from the vocabulary	onc-oneu-ral 'onco' (cancer-related) is not in the vocabulary
Compound Units	Splitting meaningful units while creating meaningless ones	neuroprot-ectant 'neuroprot' is meaningless and splits the meaningful morpheme 'protect'
Connecting Vowels	Vowels which connect two morphemes (more ambiguous)	bronch-olith optimal segmentation could split 'o' from 'lith'

Table 4: Sub-optimal segmentation types from the biomedical subset of SIGMORPHON 2022.

nents (morphemes) to map directly onto a word’s characters. For instance, in Table 5, SIGMORPHON annotations transform the root ‘neur’ into the word ‘neuron’, introducing further flexibility and complexity to the task. In order to adapt morpheme segmentation annotations to standard tokenization, we design rule-based heuristics that map each morpheme onto a subset of characters in the original word. Due to this new formulation and the success of transformer based models on this shared task, we choose a character based language model named CANINE (Clark et al., 2022) as the model to train for character tagging as morpheme segmentation. More formally, the segmentation task is re-framed as classifying each character into a B(egin) or I(nside) tag, where the B tag indicates the start of a new morpheme or token.

Original Word	Segmentation	BI Tags
onconeural	onco ##neur ##al	BIIIIBIIBI

Table 5: Example of a biomedical term segmented into morphemes and reformulated into BI tags for CANINE fine-tuning.

After fine-tuning CANINE on the full SIGMORPHON 2022 training set to create a supervised tokenization system, as seen in Figure 1 (left), we find that it achieves a 74% segmentation F1 score on the biomedical SIGMORPHON subset, a very strong result compared to current tokenizers. For reference, the best segmentation F1 score reported in the English word-level test set of the SIGMORPHON 2022 Shared Task is 93.7% by the DeepSpin team (Peters and Martins, 2022). Even though this score is not comparable to ours due to our use of a biomedical development subset for evaluation, we note that our fine-tuned CANINE model’s perfor-

mance is quite strong given that it is designed for pure tokenization as explained above.

3.3 BioVocabBERT: Domain Knowledge Injection

Despite its satisfactory segmentation performance, it is challenging to use our CANINE-based segmentation model as a language model tokenizer due to its vocabulary-less nature. Since this model can segment words arbitrarily using our character classification framework, unseen words can be split into subwords which have never been seen by the LM. Thus, pre-training a language model using this tokenizer would require allowing the model to increase its vocabulary size without limit while training and using an out-of-vocabulary token when unseen tokens are encountered during inference. This would lead to a language model with an exceedingly large vocabulary size (which would increase the cost of pre-training significantly) and potentially limited generalization ability to unseen tokens.

To tackle this problem, we introduce a tokenizer which uses the same left-to-right decoding algorithm used by BERT and PubMedBERT but replace its vocabulary with one designed for biomedical segmentation. In order to build a vocabulary which covers important biomedical tokens, we leverage the Unified Medical Language System, a medical knowledge base which contains approximately 15 million medical concept phrases. As shown in Figure 1 (center), we extract all single words from UMLS concept phrases and segment each of them using the CANINE-based tokenizer. This produces around 250,000 unique subwords which we further reduce to 55,580 by eliminating ones which only appear once in the CANINE segmented set of UMLS words. In order to avoid segmenting standard English words in unintuitive ways due to the higher proportion of biomedical subwords, we augment our 55,580 biomedical subwords with the original BERT vocabulary as seen in Figure 1 (right). After removing duplicate tokens, we are left with a vocabulary of 80,181 tokens. This carefully designed vocabulary enables our new tokenizer, BioVocabBERT, to obtain a segmentation score of 48.5% on the SIGMORPHON biomedical subset as seen in 3, outperforming the best scoring current wordpiece-based tokenizer PubMedBERT by almost 30 points.

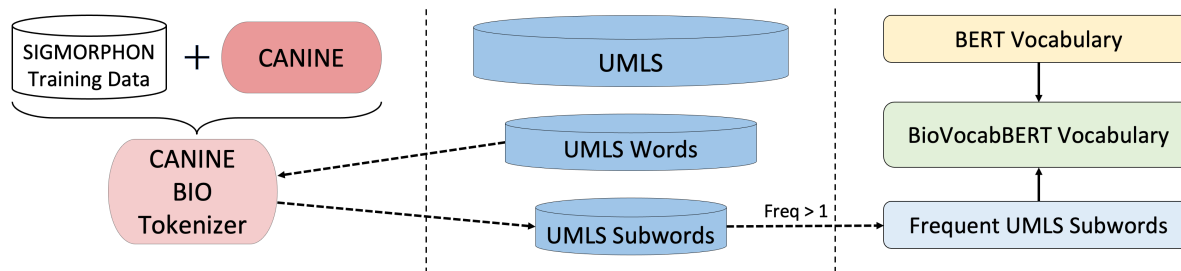


Figure 1: Overall process for creating our vocabulary for BioVocabBERT’s tokenizer. We first train a CANINE model on the SIGMORPHON training set (left). We then use this segmentation model to segment all unique UMLS words (center). Finally, we combine all UMLS subwords with frequency greater than 1 with the original BERT vocabulary to make our BioVocabBERT vocabulary (right).

4 Experimental Setup

4.1 Biomedical Pre-training

In order to discover whether morpheme segmentation performance has an effect on the biomedical language model pre-training process, we compare the downstream performance of two language models using tokenizers with very distinct segmentation performance but otherwise equivalent pre-training processes. The first model is pre-trained using the same tokenizer as PubMedBERT while the other one uses our designed BioVocabBERT tokenizer and is thus referred to by the same name.

As other BERT based architecture models, we pre-train these using the masked language modeling objective and choose standard token masking percentages used in previous work (Gu et al., 2021). We use the easily accessible and readily pre-processed corpus used for BlueBERT (Peng et al., 2020) pre-training which contains around 4 billion words³. For pre-training, we base our implementation on the work by Izsak et al. (2021) to obtain the most efficient pre-training possible. We describe the data, optimization steps, batch size, hardware and other pre-training details used for both models in Table 13 and Appendix B.

4.2 Tasks

In order to determine how tokenization improvements affect the quality of biomedical concept representation in language models, we narrow our task selection to those which are more closely related to entity understanding instead of overall sentence understanding as is the case with relation extraction, sentence similarity or natural language inference. We select named entity recognition

³<https://github.com/ncbi-nlp/BlueBERT/blob/master/README.md#pubmed>

(NER) and entity linking (EL), also referred to as concept normalization, as the two biomedical NLP tasks which most closely meet this criterion.

NER. We focus on evaluating our models in the more standard fully supervised fine-tuning NER setting as is done in previous work (Lee et al., 2019; Gu et al., 2021). We run hyperparameter tuning on the development set, the search space used can be found in Appendix A.

We also study our models’ low-resource NER performance using 500 and 1000 examples. We also carry out hyperparameter selection which can be found in Appendix A. We report results on the development set only for our low-resource NER experiments.

Entity Linking. For entity linking, we evaluate our models’ zero-shot performance as done by Liu et al. (2021) which allows us to measure entity representation quality as directly as possible.

4.3 Datasets

For NER, we use all datasets from BLURB (Gu et al., 2021), a comprehensive biomedical NLP benchmark. For entity linking, we follow previous work (Liu et al., 2021) and use four popular entity linking datasets, three of which are also included in BLURB as NER datasets. All dataset names and statistics can be found in Table 6. Below, we provide brief descriptions for each dataset we use, for more information about processing and training splits for these datasets, we refer the interested reader to the dataset descriptions in Gu et al. (2021) and Liu et al. (2021).

BC5CDR. The BioCreative V Chemical-Disease Relation corpus (Li et al., 2016) contains both disease and chemical annotations on PubMed ab-

	NER	EL	Train	Dev	Test
BC5CDR-disease	X	X	4,182	4,244	4,424
BC5CDR-chem	X	X	5,203	5,347	5,385
NCBI-Disease	X	X	5,134	787	960
JNLPBA	X		46,750	4,551	8,662
BC2GM	X		15,197	3,061	6,325
MedMentions		X	282,091	71,062	70,405

Table 6: Dataset statistics.

stracts. We evaluate disease and chemical entity extraction and linking separately following previous work (Gu et al., 2021).

NCBI-Disease. The Natural Center for Biotechnology Information Disease corpus (Doğan et al., 2014) contains disease name and concept annotations for 793 PubMed abstracts.

JNLPBA. The Joint Workshop on Natural Language Processing in Biomedicine and its Applications dataset (Collier and Kim, 2004) contains 2,000 abstracts from MEDLINE selected and annotated by hand for gene related entities.

BC2GM. The Biocreative II Gene Mention corpus (Smith et al., 2008) contains 17,500 sentences from PubMed abstracts labeled for gene entities.

MedMentions. MedMentions (Mohan and Li, 2019) is a large-scale entity linking dataset containing over 4,000 abstracts and around 350,000 mentions linked to the 2017AA version of UMLS.

5 Results & Discussion

5.1 Fully-Supervised NER

As seen in Table 7, our language models obtain competitive fully-supervised NER results compared to the results reported by Gu et al. (2021), validating our pre-training and fine-tuning process. We first find that the differences in performance between our PubMedBERT and BioVocabBERT models are very small and inconsistent across NER datasets. We note that the difference in performance between these models is often within the standard deviation reported within each dataset. Additionally, we see no pattern in performance differences based on entity types. For disease NER, BioVocab underperforms on NCBI-Disease but overperforms in BC5CDR-disease while in gene based NER, BioVocabBERT outperforms by a slightly larger margin on JNLPBA but performs only on-par on BC2GM. This seems to suggest that, at least when fine-tuning on a significant number of examples, PubMedBERT can very adequately compensate for instances of sub-optimal biomedical segmentation.

	PubMedBERT*	PubMedBERT	BioVocabBERT
NCBI-Disease	87.8	87.1 ± 0.8	86.7 ± 0.4
BC5CDR-disease	85.6	84.7 ± 0.2	85.2 ± 0.3
BC5CDR-chem	93.3	93.0 ± 0.3	93.4 ± 0.4
JNLPBA	79.1	78.2 ± 0.6	78.9 ± 0.1
BC2GM	84.5	83.4 ± 0.2	83.5 ± 0.3

Table 7: Comparison of fully supervised NER performance for the originally reported PubMedBERT, denoted by *, our PubMedBERT replica and our BioVocabBERT model. We report 3 runs for each of our models and provide the average entity-level F1 on the test set along with its standard deviation.

5.2 Low-Resource NER

To explore whether parity in fully-supervised NER comes from the effects of large scale fine-tuning or from the underlying models’ entity representation quality, we carry out a low-resource NER study using only 500 and 1,000 training examples. We present results only on the development set for this setting. Our results suggest that, even when fewer training examples are used for fine-tuning, the difference between models is small. As shown in Table 8, BioVocabBERT obtains small and inconsistent improvements in downstream performance over PubMedBERT tokenization across NER datasets in the low-resource setting. Nevertheless, we note that the largest gains for BioVocabBERT in these low data regimes come from chemical and genetic NER datasets (BC2GM and JNLPBA), suggesting that our tokenization strategy could be especially beneficial for irregular genetic entities.

	PubMedBERT		BioVocabBERT		Percent Δ	
	500	1000	500	1000	500	1000
NCBI-disease	77.2	81.2	77.7	80.6	0.5	-0.6
BC5CDR-disease	79.0	81.4	79.3	81.6	0.3	0.2
BC5CDR-chem	91.7	92.2	92.1	92.8	0.4	0.6
BC2GM	69.5	75.5	71.5	76.9	2.0	1.4
JNLPBA	75.6	77.2	76.3	77.6	0.7	0.4

Table 8: Comparing our models on low-resource NER with 500 and 1,000 examples. We report the entity-level F1 on the development set for this setting.

5.3 Zero-Shot Entity Linking

Following our low-resource NER results, we evaluate entity representation quality even more directly by measuring the zero-shot entity linking performance of both models. As shown in Table 10, performance of our models exceeds the original PubMedBERT results reported by Liu et al. (2021), validating the quality of our pre-training procedure.

Sub-optimal Tokenization	Word Embedding 5-NN	[CLS] Embedding 5-NN
epicarditis (epic-ardi-tis)	epicardiectomy (epic-ardi-ectomy) pancarditis (panc-ardi-tis) epicardin (epic-ardi-n) epicardium (epic-ardi-um) endopericarditis (endop-eric-ardi-tis)	pancarditis (panc-ardi-tis) perimyocarditis (peri-my-ocardi-tis) myopericarditis (myo-peri-car-di-tis) myoendocarditis (myo-end-ocardi-tis) pleuropericarditis (pleu-rop-eric-ardi-tis)
neuromodulation (neuromod-ulation)	neuromodulations (neuromod-ulations) neuromodulators (neuromod-ulators) neuromodulator (neuromod-ulator) immunomodulation (immunomod-ulation) immunoregulation (immunoreg-ulation)	neuromodulations (neuromod-ulations) neuromodulators (neuromod-ulators) neuromodulator (neuromod-ulator) neuroexcitation (neuro-exc-itation) neuroregulation (neuro-reg-ulation)

Table 9: In this table we show two PubMedBERT sub-optimal tokenization examples and their nearest neighbors with respect to word embeddings and ‘[CLS]’ token embeddings. Neighbors in bold are terms that were missed by the word embeddings but are retrieved correctly by the ‘[CLS]’ embeddings, repairing the sub-optimal tokenization bias.

We note that the main difference between our pre-training setup and the original PubMedBERT setting is the use of the masked language modeling (MLM) objective alone instead of both MLM and next-sentence prediction (NSP) objectives. This suggests that the use of the MLM objective only might be better aligned with obtaining high quality entity representations.

	PubMedBERT*		PubMedBERT		BioVocabBERT	
	R@1	R@5	R@1	R@5	R@1	R@5
NCBI-Disease	77.8	86.9	88.5	93.5	87.6	92.0
BC5CDR-disease	89.0	93.8	91.7	95.0	91.0	94.1
BC5CDR-chem	93.0	94.6	95.3	96.1	95.4	95.9
MedMentions	43.9	64.7	44.9	65.4	45.4	65.9

Table 10: Comparison of zero-shot entity linking performance for the originally reported PubMedBERT, denoted by *, our PubMedBERT replica and our BioVocabBERT model.

Additionally, when comparing our models, we find that BioVocabBERT slightly underperforms the PubMedBERT replica on all datasets except the more diverse MedMentions dataset. However, we note that the improvements obtained in the MedMentions dataset are also quite small at under 1%. Given the zero-shot nature of this experiment, it suggests that the entity representations obtained by these two models are of comparable quality and that PubMedBERT’s pre-training enables a high degree of robustness around sub-optimal tokenization.

5.4 Case Study: Tokenization Robustness

As shown in the NER and entity linking experiments above, the downstream performance of

biomedical language models appears to be mostly robust to biomedical concept segmentation which is not well-aligned with human judgments. To analyze this phenomenon, we take a closer look at how our pre-trained PubMedBERT model represents biomedical concepts which are segmented in apparently erroneous ways by the PubMedBERT tokenizer. Table 9 contains two words from the biomedical subset of SIGMORPHON which were sub-optimally segmented by PubMedBERT. For each of these words we include two sets of their 5 nearest neighbors according to different embedding types. The first set shows the nearest neighbors obtained using embeddings computed by averaging all subword embeddings that make up a specific word. The second set comes from using the ‘[CLS]’ token embedding of our PubMedBERT model, often used for downstream tasks in standard fine-tuning. The pool of words from which these neighbors are obtained consists of all the unique UMLS words in UMLS phrases used in the construction of BioVocabBERT.

Since it comes directly from subword embeddings, the first set of neighbors is meant to show whether the tokenizer’s sub-optimal segmentation introduces a bias which distracts the model from the true semantics of a biomedical term. The second neighborhood is meant to more faithfully show us how the model represents a biomedical concept. Comparing these two sets can let us determine if the bias introduced by subword embeddings is successfully regulated by the overall model.

We first observe that the bias we expected to find in the word embedding neighborhoods is evidently present. Most words in these first sets are

segmented in exactly the same ways as the original sub-optimally segmented word. As seen in Table 9, the word ‘neuromodulation’ is segmented by PubMedBERT as ‘neuromod-ulation’, splitting the meaningful ‘modulate’ morpheme down the middle, an example of the compound error in Table 4. Due to this, other words with the same subword but different semantics such as ‘immunomodulation’ (‘immunomod-ulation’) and ‘immunoregulation’ (‘immunoreg-ulation’) are added to the word embedding neighborhood. This is also seen in the second example, where the word embedding neighbors of ‘epicarditis’ (‘epic-ardi-tis’) all contain at least two of the three original subwords. If these word embeddings were the final model representations, this bias could lead to considerable errors in downstream tasks like entity linking by up-weighting terms based on sub-optimal subwords.

Fortunately, we observe that the language model is able to readily overcome the bias observed in the word embeddings when it comes to the final [CLS] representations. The second neighborhoods often contain semantically relevant words which were segmented differently than the original, such as ‘neuroexcitation’ (‘neuro-exc-itation’) and ‘neuroregulation’ (‘neuro-reg-ulation’) for ‘neuromodulation’ (‘neuromod-ulation’) or ‘perimyocarditis’ (‘peri-my-ocardi-tis’) for ‘epicarditis’ (‘epic-ardi-tis’), which both mean types of inflammation of the pericardium. This shows us that the language model successfully extracts the semantics of the morpheme ‘neuro’ from ‘neuromod’ as well as the cardiovascular related semantics from both ‘epic-ardi’ and ‘ocardi’, effectively mitigating the detrimental effects seen in the word embedding neighborhoods from sub-optimal tokenization. We thus conclude that this same robustness is responsible for the parity observed in downstream tasks between BioVocabBERT and the original PubMedBERT. More examples which show similar trends as the ones in Table 9 can be found in Table 14 in Appendix C.

5.5 Word Frequency Study

All the findings above suggest that biomedical language model pre-training yields entity representations which are fairly robust to tokenization failures. However, it is important to note that the distribution of rare vs. frequent entities in these small and medium scale datasets will be naturally skewed towards frequent entities if not intentionally manip-




Figure 2: MLM accuracy for our pre-trained models averaged across 10,000 word instances which fall under each word frequency bin.

ulated. Therefore, we design an experiment which explores whether the quality of representations in BioVocabBERT and our PubMedBERT replica varies with respect to word frequency. In this experiment, we obtain 10,000 instances of words from the pre-training corpus in each of the frequency bins listed in Table 2. We encode the sentence in which each instance is found and mask the word of interest. We report the percentage of words which are predicted correctly using the masked language modeling head’s prediction in each bin.

We note that both models perform very similarly across frequency bins until the two bins with the largest frequencies. Our BioVocabBERT model obtains a somewhat significant advantage in the second highest 500-5,000 frequency bin which then inverts to a similarly significant drop in the category with the most frequent words. This trade-off is likely due to many high frequency words having a single token in the PubMedBERT vocabulary, which leads the model to have a natural bias towards predicting these words. In medium frequency words, PubMedBERT’s bias towards high frequency words is likely detrimental and BioVocabBERT is able to easily outperform it. This trade-off appears to be small enough to have little effect on downstream performance for these models but we leave exploring its effect for future work.

6 Conclusion

In this work, we first note that current biomedical tokenization methods are not well aligned with human judgments and highlight that the agglutinating nature of biomedical terminology could be affected by this sub-optimal segmentation. To understand

whether this limited segmentation performance has an effect on downstream applications, we first build a biomedical tokenizer which is better aligned with human judgments using supervised morpheme segmentation and biomedical domain-knowledge. We then leverage this tokenizer to pre-train our BioVocabBERT model and compare it with a replicated PubMedBERT model on the NER and entity linking tasks. Surprisingly, we find that these models achieve almost exact parity in all datasets evaluated, suggesting that PubMedBERT’s domain specific tokenization and pre-training process was already quite robust to sub-optimal tokenization. We further verify this idea with a case-study which qualitatively confirms our observations. We hope that our work can give researchers and practitioners some insight into how instances of sub-optimal segmentation, which are often jarring to human experts, could have little effect on a model’s downstream applicability.

7 Limitations

Although our findings suggest that biomedical language model pre-training is quite robust to sub-optimal tokenization, we note that our work has a few potential limitations that should be explored further. The use of a biomedically relevant subset of the SIGMORPHON Shared Task dataset for evaluating biomedical term tokenization is a straight-forward and reasonable strategy, however, it is important to highlight that the resource was not created for this purpose and might not be perfectly aligned with ideal biomedical tokenization. Additionally, we would like to point out that even though our BioVocabBERT tokenizer outperforms other equivalent tokenizers like PubMedBERT’s, it severely underperforms the best possible segmentation accuracy (48.5 vs 74.1 for our fine-tuned CANINE model). It is therefore possible, although unexpected, that a tokenizer which performs biomedical tokenization at even higher levels could lead to sudden improvements in the pre-training process. Finally, we note that the effects of the BioVocabBERT’s much larger vocabulary size, almost three times larger than PubMedBERT’s, on the pre-training process were not explored in depth. Nevertheless, given that some previous work (Feng et al., 2022) argues that larger vocabularies lead to slight improvements in downstream tasks, our main conclusions are likely to hold.

Acknowledgements

The authors would like to thank the anonymous reviewers and colleagues from the OSU NLP group for their valuable feedback. This research was supported in part by NIH R01LM014199 and Ohio Supercomputer Center (Center, 1987).

References

- Emily Alsentzer, John Murphy, William Boag, Weihung Weng, Di Jindi, Tristan Naumann, and Matthew McDermott. 2019. [Publicly available clinical BERT embeddings](#). In *Proceedings of the 2nd Clinical Natural Language Processing Workshop*, pages 72–78, Minneapolis, Minnesota, USA. Association for Computational Linguistics.
- George L. Banay. 1948. An Introduction to Medical Terminology I. Greek and Latin Derivations. *Bulletin of the Medical Library Association*, 36,1:1–27.
- Khuyagbaatar Batsuren, Gábor Bella, Aryaman Arora, Viktor Martinovic, Kyle Gorman, Zdeněk Žabokrtský, Amarsanaa Ganbold, Šárka Dohnalová, Magda Ševčíková, Kateřina Pelegrinová, Fausto Giunchiglia, Ryan Cotterell, and Ekaterina Vylomova. 2022. [The SIGMORPHON 2022 shared task on morpheme segmentation](#). In *Proceedings of the 19th SIGMORPHON Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 103–116, Seattle, Washington. Association for Computational Linguistics.
- Khuyagbaatar Batsuren, Gábor Bella, and Fausto Giunchiglia. 2021. [MorphoNet: a large multilingual database of derivational and inflectional morphology](#). In *Proceedings of the 18th SIGMORPHON Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 39–48, Online. Association for Computational Linguistics.
- Olivier Bodenreider. 2004. The Unified Medical Language System (UMLS): Integrating Biomedical Terminology. *Nucleic acids research*, 32 Database issue:D267–70.
- Ohio Supercomputer Center. 1987. [Ohio supercomputer center](#).
- Ilias Chalkidis, Manos Fergadiotis, Prodromos Malakasiotis, Nikolaos Aletras, and Ion Androutsopoulos. 2020. [LEGAL-BERT: The muppets straight out of law school](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2898–2904, Online. Association for Computational Linguistics.
- Jonathan H. Clark, Dan Garrette, Iulia Turc, and John Wieting. 2022. [Canine: Pre-training an efficient tokenization-free encoder for language representation](#). *Transactions of the Association for Computational Linguistics*, 10:73–91.

- Nigel Collier and Jin-Dong Kim. 2004. [Introduction to the Bio-entity Recognition Task at JNLPBA](#). In *Proceedings of the International Joint Workshop on Natural Language Processing in Biomedicine and its Applications (NLPBA/BioNLP)*, pages 73–78, Geneva, Switzerland. COLING.
- C Edward Collins. 2007. *A Short Course in Medical Terminology*. Lippincott Williams and Wilkins, Philadelphia, PA.
- Mathias Creutz, K. Lagus, and Sami Virpioja. 2005. Unsupervised Morphology Induction Using Morfesator. In *Finite-State Methods in Natural Language Processing: 5th International Workshop (FSMNLP)*, pages 300–301.
- Alexandra DeLucia, Shijie Wu, Aaron Mueller, Carlos Aguirre, Philip Resnik, and Mark Dredze. 2022. [Bernice: A multilingual pre-trained encoder for Twitter](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 6191–6205, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Rezarta Islamaj Doğan, Robert Leaman, and Zhiyong Lu. 2014. [NCBI disease corpus: A resource for disease name recognition and concept normalization](#). *Journal of Biomedical Informatics*, 47:1–10.
- Hicham El Boukkouri, Olivier Ferret, Thomas Lavergne, and Pierre Zweigenbaum. 2022. [Re-train or train from scratch? comparing pre-training strategies of BERT in the medical domain](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 2626–2633, Marseille, France. European Language Resources Association.
- Zhangyin Feng, Duyu Tang, Cong Zhou, Junwei Liao, Shuangzhi Wu, Xiaocheng Feng, Bing Qin, Yunbo Cao, and Shuming Shi. 2022. Pretraining without wordpieces: Learning over a vocabulary of millions of words. *ArXiv*, abs/2202.12142.
- Philip Gage. 1994. A new algorithm for data compression. *The C Users Journal archive*, 12:23–38.
- 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 Trans. Comput. Healthcare*, 3(1).
- Valentin Hofmann, Janet Pierrehumbert, and Hinrich Schütze. 2020. [DagoBERT: Generating derivational morphology with a pretrained language model](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3848–3861, Online. Association for Computational Linguistics.
- Valentin Hofmann, Janet Pierrehumbert, and Hinrich Schütze. 2021. [Superbizarre is not superb: Derivational morphology improves BERT’s interpretation of complex words](#). 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 3594–3608, Online. Association for Computational Linguistics.
- Peter Izsak, Moshe Berchansky, and Omer Levy. 2021. [How to train BERT with an academic budget](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10644–10652, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2019. [BioBERT: a pre-trained biomedical language representation model for biomedical text mining](#). *Bioinformatics*, 36(4):1234–1240.
- Jiao Li, Yueping Sun, Robin J Johnson, Daniela Sciaky, Chih-Hsuan Wei, Robert Leaman, Allan Peter Davis, Carolyn J Mattingly, Thomas C Wieggers, and Zhiyong Lu. 2016. [BioCreative V CDR task corpus: a resource for chemical disease relation extraction](#). *Database (Oxford)*, 2016:baw068.
- Fangyu Liu, Ehsan Shareghi, Zaiqiao Meng, Marco Basaldella, and Nigel Collier. 2021. [Self-alignment pretraining for biomedical entity representations](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4228–4238, Online. Association for Computational Linguistics.
- Renqian Luo, Liai Sun, Yingce Xia, Tao Qin, Sheng Zhang, Hoifung Poon, and Tie-Yan Liu. 2022. [BioGPT: generative pre-trained transformer for biomedical text generation and mining](#). *Briefings in Bioinformatics*, 23(6). Bbac409.
- Himanshu Maheshwari, Bhavyajeet Singh, and Vasudeva Varma. 2021. [SciBERT sentence representation for citation context classification](#). In *Proceedings of the Second Workshop on Scholarly Document Processing*, pages 130–133, Online. Association for Computational Linguistics.
- Sunil Mohan and Donghui Li. 2019. [Medmentions: A large biomedical corpus annotated with UMLS concepts](#). In *Proceedings of the 2019 Conference on Automated Knowledge Base Construction (AKBC 2019)*, Amherst, Massachusetts, USA.

- Yifan Peng, Qingyu Chen, and Zhiyong Lu. 2020. [An empirical study of multi-task learning on BERT for biomedical text mining](#). In *Proceedings of the 19th SIGBioMed Workshop on Biomedical Language Processing*, pages 205–214, Online. Association for Computational Linguistics.
- Ben Peters and Andre F. T. Martins. 2022. [Beyond characters: Subword-level morpheme segmentation](#). In *Proceedings of the 19th SIGMORPHON Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 131–138, Seattle, Washington. Association for Computational Linguistics.
- Mike Schuster and Kaisuke Nakajima. 2012. Japanese and Korean voice search. *2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5149–5152.
- Larry Smith, Lorraine K Tanabe, Rie Johnson Nee Ando, Cheng-Ju Kuo, I-Fang Chung, Chun-Nan Hsu, Yu-Shi Lin, Roman Klinger, Christoph M Friedrich, Kuzman Ganchev, Manabu Torii, Hongfang Liu, Barry Haddow, Craig A Struble, Richard J Povinelli, Andreas Vlachos, William A Baumgartner, Jr, Lawrence Hunter, Bob Carpenter, Richard Tzong-Han Tsai, Hong-Jie Dai, Feng Liu, Yifei Chen, Chengjie Sun, Sophia Katrenko, Pieter Adriaans, Christian Blaschke, Rafael Torres, Mariana Neves, Preslav Nakov, Anna Divoli, Manuel Maña-López, Jacinto Mata, and W John Wilbur. 2008. [Overview of BioCreative II gene mention recognition](#). *Genome Biology*, 9(S2):S2.
- Shijie Wu, Ozan Irsoy, Steven Lu, Vadim Dabravolski, Mark Dredze, Sebastian Gehrmann, Prabhanjan Kam-badur, David Rosenberg, and Gideon Mann. 2023. [Bloomberggpt: A large language model for finance](#). *ArXiv*, abs/2303.17564.
- Michihiro Yasunaga, Jure Leskovec, and Percy Liang. 2022. [LinkBERT: Pretraining language models with document links](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8003–8016, Dublin, Ireland. Association for Computational Linguistics.
- Sicheng Zhou, Nan Wang, Liwei Wang, Hongfang Liu, and Rui Zhang. 2022. [CancerBERT: a cancer domain-specific language model for extracting breast cancer phenotypes from electronic health records](#). *Journal of the American Medical Informatics Association : JAMIA*, 29:1208 – 1216.

A NER Hyperparameter Tuning

We run hyperparameter tuning for each model in the fully-supervised and low-resource NER settings and they can both be found in Tables 11 and 12 below.

	Learning Rate	Batch Size	Warmup Ratio	Weight Decay	Total Epoch Number
Search	1e-5	16	0.06	0.1	5
Space	3e-5	32			10

Table 11: Hyperparameter search grid used for fully-supervised NER experiments.

	Learning Rate	Batch Size	Warmup Ratio	Weight Decay	Total Epoch Number
Search	1e-5	16	0.06	0.1	15
Space	3e-5	32			25

Table 12: Hyperparameter search grid used for low-resource NER experiments.

B Pre-training Details

Our models were pre-trained on 4 80GB A100s. The process took approximately 2 and 3 days respectively for PubMedBERT and BioVocabBERT given the larger computational requirements of using an 80,000 subword vocabulary.

	Objectives	Vocab. Size	Corpus Size	Gradient Steps	Batch Size	# of Examples
PubMedBERT (Original)	MLM & NSP	28,895	21GB	62,500	8,192	512M
PubMedBERT (Replica)	MLM	28,895	25GB	62,500	8,192	512M
BioVocabBERT	MLM	80,181	25GB	62,500	8,192	512M

Table 13: Pre-training details for the original PubMedBERT compared to our models.

C More Neighborhood Examples

The neighborhood examples shown in Table 14, help demonstrate that the general trends discussed in §5.4 hold more generally for many sub-optimal segmentation examples.

Sub-optimal Segmentation	Word Embedding 5-NN	[CLS] Embedding 5-NN
abdominopelvic (abdom-ino-pe-lv-ic)	abdominopelvis (abdom-ino-pe-lv-is) sacropelvic (sacro-pe-lv-ic) uteropelvic (utero-pe-lv-ic) abdomino (abdom-ino) midpelvic (mid-pe-lv-ic)	abdominocentesis (abdom-ino-cent-esis) thoracopelvic (thorac-ope-lv-ic) midpelvic (mid-pe-lv-ic) sacropelvic (sacro-pe-lv-ic) extrapelvic (extrap-elvic)
neuroradiography (neuroradi-ography)	roentgenography (roentgen-ography) ventriculography (ventricul-ography) neuroradiology (neuroradi-ology) electroretinography (electroretin-ography) herniography (herni-ography)	neuroradiology (neuroradi-ology) neuroradiologic (neuroradi-ologic) encephalography (encephal-ography) neurography (neuro-graphy) cerebroangiography (cerebro-angi-ography)
postinfectious (postin-fecti-on-al)	postinfection (postin-fecti-on) postin (postin) postinjection (postin-jecti-on) reinfection (rein-fecti-on) postinfusion (postin-fusi-on)	reinfection (rein-fecti-on) postinfection (postin-fecti-on) superinfection (super-infecti-on) reinfected (rein-fect-ed) superinfections (super-infecti-on-s)
neurogastrointestinal (neuro-ga-st-ro-intestinal)	extragastrointestinal (extra-ga-st-ro-intestinal) pangastrointestinal (pan-ga-st-ro-intestinal) myoneurogastrointestinal (myo-ne-uro-ga-st-ro-intestinal) gastrogastric (gastro-ga-st-ric) gastrogastrostomy (gastro-ga-st-rost-omy)	extragastrointestinal (extra-ga-st-ro-intestinal) pangastrointestinal (pan-ga-st-ro-intestinal) enteropancreatic (enter-opancre-atic) gastroenteropancreatic (gastroenter-opancre-atic) nasopancreatic (nas-opancre-atic)
adrenocorticosteroid (adrenocortic-oster-oid)	adrenocorticosteroids (adrenocortic-oster-oids) glucocorticosteroid (glucocortic-oster-oid) mineralocorticosteroid (mineralocortic-oster-oid) mineralocorticosteroids (mineralocortic-oster-oids) glucosteroid (gluc-oster-oid)	adrenocorticosteroids (adrenocortic-oster-oids) adrenocorticotropic (adrenocortic-otro-pic) corticoids (cortic-oids) corticoid (cortic-oid) glucosteroid (gluc-oster-oid)

Table 14: In this table we show more PubMedBERT sub-optimal segmentation examples and their nearest neighbors with respect to word embeddings and ‘[CLS]’ token embeddings. As in Table 9, bold neighbors were missed by the word embeddings but are retrieved correctly by the ‘[CLS]’ embeddings, repairing the sub-optimal segmentation bias.