

# MEDVOC: Vocabulary Adaptation for Fine-tuning Pre-trained Language Models on Medical Text Summarization

Gunjan Balde<sup>1</sup>, Soumyadeep Roy<sup>1,2</sup>, Mainack Mondal<sup>1</sup> and Niloy Ganguly<sup>1</sup>

<sup>1</sup>Indian Institute of Technology Kharagpur

<sup>2</sup> L3S Research Center, Germany

balde.gunjan0812@kgpian.iitkgp.ac.in, soumyadeep.roy9@iitkgp.ac.in, {mainack, niloy}@cse.iitkgp.ac.in

## Abstract

This work presents a dynamic vocabulary adaptation strategy, MEDVOC, for fine-tuning pre-trained language models (PLMs) like BertSumAbs, BART, and PEGASUS for improved medical text summarization. In contrast to existing domain adaptation approaches in summarization, MEDVOC treats vocabulary as an *optimizable parameter* and optimizes the PLM vocabulary based on *fragment score* conditioned only on the downstream task’s reference summaries. Unlike previous works on vocabulary adaptation (limited only to classification tasks), optimizing vocabulary based on summarization tasks requires an extremely costly intermediate fine-tuning step on large summarization datasets. To that end, our novel *fragment score*-based hyperparameter search very significantly reduces this fine-tuning time — from 450 days to less than 2 days on average. Furthermore, while previous works on vocabulary adaptation are often primarily tied to single PLMs, MEDVOC is designed to be deployable across multiple PLMs (with varying model vocabulary sizes, pre-training objectives, and model sizes) — bridging the limited vocabulary overlap between the biomedical literature domain and PLMs. MEDVOC outperforms baselines by 15.74% in terms of Rouge-L in zero-shot setting and shows gains of 17.29% in high Out-Of-Vocabulary (OOV) concentrations. Our human evaluation shows MEDVOC generates *more faithful* medical summaries (88% compared to 59% in baselines).

## 1 Introduction

Medical text summarization is useful for many real-life use-cases such as summary generation of clinical records [Kanwal and Rizzo, 2022], health-related queries [He *et al.*, 2021], and radiology reports [Dai *et al.*, 2021]. Most medical summarization approaches are based on pre-trained language models (PLMs) that are trained on text from open-domain

We make the codebase and supplementary materials publicly available at <https://github.com/gb-kgp/MEDVOC>.

<p style="text-align: center;"><i>Source Document</i></p> <p>what is the role of prokinetic agents for <b>constipation</b>? the efficacy of <b>erythromycin</b> was assessed in the treatment of 14 children aged 4 to 13 with refractory chronic <b>constipation</b>, and presenting megarectum and fecal impaction ... the patients were randomized to receive placebo for 4 weeks followed by <b>erythromycin</b> estolate, 20 mg kg-1 day-1, divided into four oral doses for another 4 weeks, or vice versa. .... mean laxative requirement was lower when patients ingested erythromycin (p &lt; 0.05). no <b>erythromycin</b>-related side effects occurred. <b>erythromycin</b> was useful in this group of severely constipated children...</p>
<p style="text-align: center;"><i>Reference Summary</i></p> <p>a small <b>rct</b> of 14 children between 4 and 13 years of age showed that <b>erythromycin</b> improved symptoms of <b>constipation</b> and decreased <b>laxative</b> use ( number needed to treat [ nnt ] = 10 )</p>
<p style="text-align: center;"><i>IFT-CNN Generated Summary</i></p> <p>a <b>double-blind</b> , <b>placebo-controlled</b> crossover <b>study of 14 children aged 4 to 13 years</b> compared <b>erythromycin</b> estolate with placebo for 4 weeks followed by <b>erythromycin</b> estolate , 20 mg/1 day-1 , divided into 4 oral doses for another 4 weeks . patients were randomized to <b>receive placebo</b>, followed by a 20 mg/d of <b>erythromycin</b> estolate 20mg/d , and then received <b>placebo</b>.</p>
<p style="text-align: center;"><i>MEDVOC Generated Summary</i></p> <p>a <b>double-blind</b> , <b>placebo-controlled</b> crossover <b>study of 14 children aged 4 to 13 years</b> showed that <b>erythromycin</b> was more effective than <b>placebo</b> ( p &lt; .05 ) .</p>
<p style="text-align: center;"><i>Medical OOV words (before &gt; after)</i></p> <p>con ##sti ##pati ##on &gt; constipation er ##yt ##hr ##omy ##cin &gt; erythromycin consti ##pate ##d &gt; constipated gas ##tro ##enter ##ology &gt; gastro ##enter ##ology</p>

Figure 1: Illustrative example of BertSumAbs model from EBM dataset. Purple text color: indicates semantic or surface form overlaps with RS, Bold text: indicates medical (UMLS) concept-bearing words, Yellow highlight: OOV words that are ultimately added to the updated vocabulary, and Orange highlight: medical concept-bearing word(s) that overlap with reference summary.

sources. Thus, their performance is sub-optimal because they did not incorporate medical knowledge into their models [He *et al.*, 2021; Zhang *et al.*, 2021]. Domain adaptation approaches for summarization in general [Fabbri *et al.*, 2021a; Laskar *et al.*, 2022; Xu *et al.*, 2023] and medical summarization [He *et al.*, 2021; Lamproudis. *et al.*, 2022; Xie *et al.*, 2023b], in particular, have garnered reasonable research interest. However, we identified two prominent research gaps in existing works.

First, despite poor domain similarity of 33% in Figure 3(a) between CNN/DailyMail [See *et al.*, 2017] (open domain) and PubMed Abstracts Collection (medical domain), none of

the domain adaptation approaches for medical summarization update the PLM’s vocabulary during fine-tuning. Figure 1 demonstrates the challenges that arise when generating medical summaries without vocabulary adaptation. Figure 3(b) demonstrates that medical concept-bearing words mostly lose their meaning due to poor representation because they are tokenized into four or more subwords. However, the undesirable tokenization actually happens at the decoder level during the generation of such medical concept-bearing words in a summary. Vocabulary adaptation is successful in the classification settings where it updates the PLM vocabulary adding a target domain-specific vocabulary [Hong *et al.*, 2021; Xu *et al.*, 2021; Lamproudis *et al.*, 2022]. Unfortunately, the vocabulary set construction algorithms of the classification setting are quite restrictive as they rely on fixed empirical thresholds [Tai *et al.*, 2020; Hong *et al.*, 2021] and fail to adapt themselves to a new, significantly different PLM architecture. In this paper, we are **first to explore vocabulary adaptation techniques for summarization** and design the MEDVOC fine-tuning strategy. However, adapting vocabulary adaptation strategies (earlier classification) to summarization (a generative setting) is non-trivial for MEDVOC: (i) the decoder needs to be additionally trained, (ii) rare sub-words get included in the target domain-specific vocabulary, and (iii) absence of large-scale medical summarization datasets for intermediate fine-tuning purposes. We address these challenges in this work.

The second research gap is that most of the vocabulary adaptation till now is evaluated on a single PLM, as shown in Table 1. To this end, as part of MEDVOC, we develop an **efficient, dynamic vocabulary construction step** that adapts to any encoder-decoder-based (PLM) summarization model and target (downstream task) datasets. We treat vocabulary construction as a hyperparameter tuning step and show that optimizing the fragment score of reference summaries in a given setting closely resembles the same optimization as the downstream task performance. This helps to avoid the extremely time-consuming step of intermediate fine-tuning. Our fragment score-based hyperparameter search very significantly reduces this fine-tuning time, from 450 days to  $\sim 45$  hours, averaged across three PLMs over four downstream medical summarization tasks.

To address the two key research gaps, we need to first train the PLMs on the downstream medical summarization tasks. However, the size of the training datasets of the downstream medical summarization task is quite small, in the range of 700 to 1525 data points. Directly fine-tuning pre-trained models on such small *target datasets* could lead to sub-optimal performance [Phang *et al.*, 2019]. The standard approach relies on introducing an intermediate fine-tuning stage using large datasets [Chang and Lu, 2021; Suresh *et al.*, 2023]. In this work, we show that the **task of biomedical paper title generation serves as a good intermediate fine-tuning task** [Fabri *et al.*, 2021a] for medical text summarization.

MEDVOC outperforms baselines by 15.74%, 4.80%, and 5.99% in zero-shot, few-shot, and full dataset settings respectively on average across four medical summarization datasets and three PLMs such as BertSumAbs [Liu and Lapata, 2019], BART [Lewis *et al.*, 2020] and PEGASUS [Zhang *et al.*,

Related Works	PLMs	Task	VA	IFT
Tai <i>et al.</i> , [2020]	BERT	Classification	✓	✗
Diao <i>et al.</i> , [2021]	RoBERTa	Classification	✗	✗
Hong <i>et al.</i> , [2021]	BERT	Classification	✓	✗
Lamproudis <i>et al.</i> , [2022]	BERT	Classification	✓	✗
Xu <i>et al.</i> , [2023]	BERT	Classification	✗	✗
Xu <i>et al.</i> , [2021]	Transformer-Big	Machine Translation	✓	✗
Liu <i>et al.</i> , [2023]	GPT-2, BART	Question Answering	✓	✗
Fabri <i>et al.</i> , [2021a]	BART-L	Summarization	✗	✓
Xie <i>et al.</i> , [2022]	BERT	Summarization	✗	✗
MEDVOC (Ours)	BART-L, Pegasus-L, BertSumAbs	Summarization	✓	✓

Table 1: Comparison of experimental setup with related works. VA stands for *Vocabulary Adaptation* and IFT stands for *Intermediate fine-Tuning*. ‘-L’ refers to the *Large* model variant.

2020a] in terms of Rouge-L. MEDVOC produces more informative (5.99% Rouge-L improvement on average) and more faithful medical summaries (5.96% Concept Score improvement on average). We observe gains of 10.81% and 17.29% across challenging scenarios of long-form medical summary generation and reference summaries with high out-of-vocabulary (OOV) concentration respectively. We make the codebase and supplementary materials publicly available at <https://github.com/gb-kgp/MEDVOC>.

## 2 Related Works

To the best of our knowledge, this is the first work to explore vocabulary adaptation strategies for the summarization task. However, recent works have explored vocabulary adaptation strategies for classification tasks and domain adaptation techniques for summarization, which we present in Table 1.

**Vocabulary Adaptation Strategies For Classification.** To handle the vocabulary mismatch issue, BioBERT [Lee *et al.*, 2019], retrained the model from scratch using a domain-specific corpus and showed performance improvement over the base pre-trained models. While other works like VOLT [Xu *et al.*, 2021] and AVocaDo [Hong *et al.*, 2021] aim to optimize the model’s vocabulary by adding a set of subwords to the existing vocabulary using some utility scoring function. AVocaDo uses a fragment score-based [Rust *et al.*, 2021] threshold, whereas T-DNA [Diao *et al.*, 2021] selects n-grams with high Pointwise Mutual Information and iteratively merges them. exBERT [Tai *et al.*, 2020] adopts an ad-hoc approach to determine the size of  $V_{TGT}$  and fix it as 17K (56% of pre-trained vocabulary size), and does not perform any vocabulary optimization. However, all these vocabulary adaptation works have two major drawbacks: (i) they are limited to classification tasks, and (ii) they show results on a single model type and their algorithm is not flexible enough to handle different model types. To the best of our knowledge, this is the first work that explores vocabulary adaptation strategies for summarization over multiple PLMs. We focus on encoder-decoder-based PLMs because it is computationally infeasible to re-train medical LLMs [Chen *et al.*, 2023; Wu *et al.*, 2023].

### 3 Proposed Methodology

Here, we describe the MEDVOC fine-tuning strategy for adapting PLMs to medical text summarization tasks in Figure 2. We present the dynamic vocabulary construction step of MEDVOC in Section 3.1, and then explain the intermediate fine-tuning details in Section 3.2.

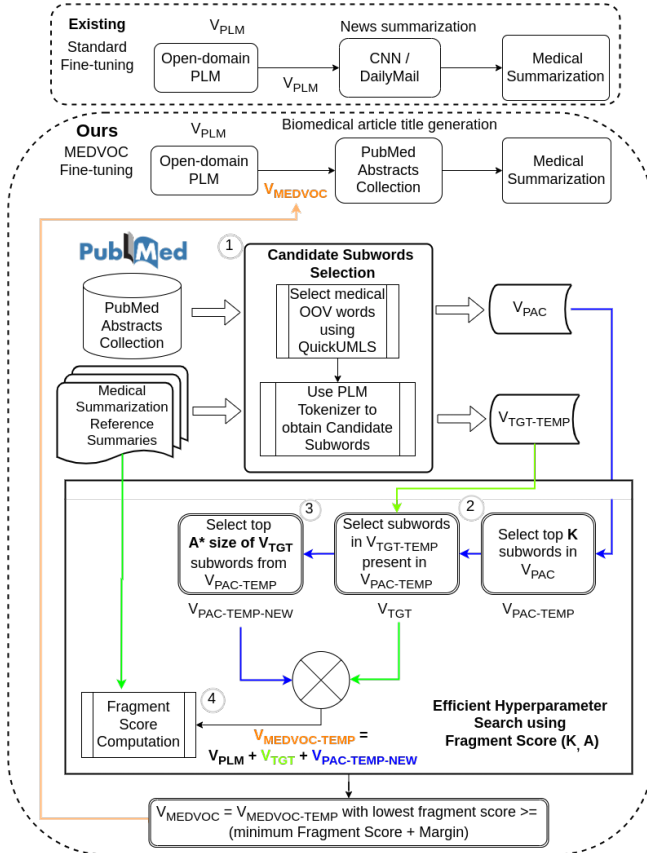


Figure 2: Methodological overview of MEDVOC and existing fine-tuning strategy.

#### 3.1 Dynamic Task-Aware Vocabulary Adaptation

The key challenge is that the base PLM’s vocabulary ( $V_{PLM}$ ) remains unchanged during intermediate fine-tuning. Since this model vocabulary is obtained from training on open-domain data, we observe that  $V_{PLM}$  misses (important) medical terms (Figure 3(b)) present in the target dataset, which occurs quite frequently across datasets considered in this study. This causes the PLM model tokenizer to split relevant medical terms into too many meaningless subwords, which ultimately is not able to capture the term semantics. This phenomenon is well-observed in prior studies w.r.t classification tasks [Xu *et al.*, 2021; Hong *et al.*, 2021] and results in poor downstream task performance. Therefore, we update  $V_{PLM}$  by adding a set of target domain-specific subwords to the vocabulary, we refer to this updated model vocabulary as  $V_{MEDVOC}$ .

**Candidate Subwords Generation.** We first identify a set of words present in reference summaries (RS) of target tasks

that are poorly tokenized, i.e., split into more than 3 subwords. However, this resulted in poor coverage as it constituted a small fraction of 21.99% out of the total set of out-of-vocabulary (OOV) words on average, i.e., words split into more than one subword. Therefore, we also include medically relevant OOV words. We use the *matcher.match* function of QuickUMLS [Soldaini and Goharian, 2016] tool, with the parameters - (i) similarity measure as ‘*cosine*’, and (ii) similarity threshold at 95%, to obtain the set of such medically relevant words. We run the PLM’s tokenization scheme on these selected words for the target task to obtain  $V_{TGT-TEMP}$ . We apply the same procedure on the source documents of the Pubmed Abstracts Collection (PAC) dataset to obtain  $V_{PAC}$  (see Section 4.4 for more details).

**Vocabulary Construction Using Target Datasets ( $V_{TGT}$ ).** We observe that  $V_{TGT-TEMP}$  contain subwords that are infrequent in PAC and it is well known that adding infrequent downstream task-specific subwords may lead to the rare word problem [Schick and Schütze, 2020; Hong *et al.*, 2021]. Since PAC is used for intermediate fine-tuning, these infrequent words would appear in fewer contexts during training, and thus lead to a sub-optimal (noisier) representation. Therefore, we only consider subwords in  $V_{TGT-TEMP}$  that overlap with the top  $K$  subwords of  $V_{PAC}$ , thus mitigating the rare word issue. This value of  $K$  is the first hyperparameter for MEDVOC. However, we empirically observe that the size of  $V_{TGT}$  is quite small as compared to the PLM’s vocabulary size (3.66%, 1.58%, and 0.13% for BertSumAbs, BART, and PEGASUS respectively). This leads to marginal performance improvement as the added sub-words are overshadowed by the PLM vocabulary during summary generation.

**Optimal Subset Selection From  $V_{PAC}$ .** Since this step is upper-bounded by the size of  $V_{PAC}$ , which is quite large (3 times the model vocabulary), a large vocabulary causes parameter explosion and token sparsity problems, which hurts model learning [Allison *et al.*, 2006; Xu *et al.*, 2021]. Therefore, we put an upper limit equal to the  $|V_{PLM}|$ . This also makes sure that the added vocabulary size does not exceed the PLM vocabulary size. Thus, determining the optimal subset size of  $V_{PAC}$  lies in a tradeoff between model vocabulary size (large enough w.r.t PLM vocabulary) and model performance (small enough to not degrade downstream task performance). Therefore, we also include top  $P$  subwords from  $V_{PAC}$ , where  $P = A \times |V_{TGT}|$ . The value of  $A$  thus forms the second hyperparameter for MEDVOC.

**Efficient Hyperparameter Search Using Fragment Score.** The standard hyperparameter search to find optimal values of  $A$  and  $K$  based on downstream task performance is extremely time-consuming. This is because it involves an additional intermediate fine-tuning step using PAC (312K data points) that takes 45 hours to run across 3 Tesla V100 32 GB GPUs (averaged across the three PLMs over four datasets). An exhaustive grid search over 240 settings takes 450 days. Our efficient fragment score-based hyperparameter search performs the same in 2 days on average, thus leading to a 240x speedup. Upon extensive evaluation, we observe that by optimizing the fragment score for a model type and dataset, the downstream task performance also gets optimized.

---

**Algorithm 1: Vocabulary Construction of MEDVOC**


---

```

1 Input: Pre-trained Vocabulary ( $V_{PLM}$ ), Model tokenizer type
    $T$ , Source documents of Pubmed Abstracts
   Collection  $d_{PAC}$ , Reference Summaries of Target
   dataset  $d_{TGT}$ 
Output: MEDVOC vocabulary  $-V_{MEDVOC}$ 
Initialization:  $K$ : Selecting Top- $K$  subwords in  $V_{PAC}$ ,  $A$ :
   Factor over size of  $V_{TGT}$ , Margin: 0.04
2 Function FragmentScore( $d_{train}, V$ ):
3    $f_C(V) \leftarrow \frac{\text{total count of subwords tokenized by } V}{\text{word count in } d_{train}}$ 
4   return  $f_C(V)$ 
5  $V_{TGT-TEMP} \leftarrow \text{CandidateSubwordsSelection}(d_{TGT}, T)$ 
6  $V_{PAC} \leftarrow \text{CandidateSubwordsSelection}(d_{PAC}, T)$ 
7 for  $A \leftarrow 0.25$  to 10 by 0.25 do
8   for  $K \leftarrow 5000$  to  $\min(V_{PLM}, V_{PAC})$  by 5000 do
9      $V_{PAC-TEMP} \leftarrow V_{PAC}[0 : K]$  // Select top- $K$ 
      subwords
10     $V_{TGT} \leftarrow V_{TGT-TEMP} \cap V_{PAC-TEMP}$  // Mitigating
      the rare word issue
11     $P = \min(V_{PLM}, A \cdot |V_{TGT}|)$  // Size to
      sample from  $V_{PAC-TEMP}$ 
12     $V_{PAC-TEMP-NEW} \leftarrow V_{PAC}[0 : P]$  // Select
      top- $P$  subwords from  $V_{PAC}$ 
13     $V_{MEDVOC-TEMP} \leftarrow V_{PLM} \cup V_{TGT} \cup V_{PAC-TEMP-NEW}$ 
14     $frag\_score(A, K) \leftarrow$ 
      FragmentScore( $d_{TGT}, V_{MEDVOC-TEMP}$ )
15  $\text{minFragScore} \leftarrow$  Minimum fragment score across all values
      of  $A, K$ 
16  $V_{MEDVOC} \leftarrow V_{MEDVOC-TEMP}$  with smallest vocabulary size
      with fragment score within a Margin of  $\text{minFragScore}$ 
17 return  $V_{MEDVOC}$ 

```

---

Fragment score [Hong *et al.*, 2021; Rust *et al.*, 2021] is defined as the average number of subwords that a word gets tokenized into by a base model’s tokenizer. We consider only those words which are split into more than one subword for fragment score computation on the target dataset. Since computing the fragment score requires the downstream task’s reference summaries and is independent of the intermediate fine-tuning step, the time taken for hyperparameter search drastically reduces to only a few hours. We find that the lowest time taken is 1 hour and 30 minutes on a single core of Intel i5 12-core CPU for BertSumAbs on the CHQSum dataset, whereas the highest time taken is 5 hours and 45 minutes for PEGASUS on the EBM dataset. Therefore, the speedup observed by MEDVOC is proportional to the grid search space used for hyperparameter optimization. Our final vocabulary  $V_{MEDVOC}$  comprises subwords that lead to fragment scores within a certain range of best-achievable (minimum) fragment scores on the target dataset’s reference summaries. Algorithm 1 further explains the dynamic vocabulary construction step.

**Time Complexity Of MEDVOC.** The candidate subword selection step as well as the fragment score computation is proportional to the size of the input corpus ( $d_{PAC}$  and  $d_{TGT}$ ). The hyperparameter search space is dependent on a constant set of values of  $A$ , while  $K$  is conditioned on  $|V_{PAC}|$ , with the upper limit being  $|V_{PLM}|$ , which results in time complexity

of  $O(|V_{PLM}| * |d_{TGT}|)$ . The time complexity of MEDVOC is  $O(|d_{PAC}|) + O(|d_{TGT}|) + O(|V_{PLM}| * |d_{TGT}|)$ . Since, our target task datasets have limited size ( $|d_{TGT}|$  ranges between 700 to 1525 documents) and  $|d_{TGT}| \ll |d_{PAC}|$ . The final time complexity is  $O(|V_{PLM}| * |d_{TGT}|) + O(|d_{PAC}|)$ .

### 3.2 Intermediate Fine-Tuning With Biomedical Article Title Generation

Intermediate fine-tuning (IFT) is known to help PLMs when the downstream task has limited training (fine-tuning) data [Chang and Lu, 2021; Fabbri *et al.*, 2021a]. In our case, the training dataset sizes range between 700 and 1525 data points (see Table 2), which is too small for training purposes and would easily overfit the PLMs, leading to poor performance. However, large-scale summarization datasets (similar to CNN/DailyMail in the open domain) are required for meaningful intermediate fine-tuning, which is absent in the medical domain. We show that *biomedical article title generation* satisfies the properties of a good intermediate fine-tuning task before fine-tuning on the downstream task of medical abstractive text summarization. Given that a good intermediate task aims to capture the knowledge or characteristics of the target task [Chang and Lu, 2021; Suresh *et al.*, 2023], we use PAC as an intermediate fine-tuning task because it closely reflects key characteristics of the downstream summarization datasets (see Section 5.2).

**Fine-tuning Details.** Intermediate fine-tuning is performed using PAC (once for every PLM), whereas standard fine-tuning is done using downstream summarization tasks (once for every target dataset and PLM combination). Except for difference in the training dataset, we follow the standard fine-tuning procedure for the summarization task. We observe a marginal increase in the model’s parameter count on adopting the MEDVOC strategy as only the embedding matrix corresponding to the added vocabulary needs to be additionally trained. The parameter count increments by 0.15%, 1.15% and 1.59% in case of PEGASUS, BART, and BertSumAbs.

## 4 Experimental Setup

We describe datasets and evaluation metrics, followed by baselines and training details (more details in Appendix A).

### 4.1 Target Task Datasets

We evaluate MEDVOC on two medical document summarization and two medical question summarization tasks. We provide detailed description of datasets in Appendix A.2.

**Medical document summarization.** In *BioASQ* [Tsatsaronis *et al.*, 2015] and *EBM* [Mollá and Santiago-Martínez, 2011], each data point contains a query and PubMed (a biomedical database) abstract as the source document (SD) and an answer to the query as a reference summary (RS).

**Medical question summarization.** *MeQSum* [Ben Abacha and Demner-Fushman, 2019] and *CHQSum* [Yadav *et al.*, 2022] contain consumer health questions posed by (non-medical) users as the SD and a short question (manually curated by medical experts) as the RS. *MeQSum* and *CHQSum* comprise questions provided by the U.S National Library of Medicine and the Yahoo! Answers L6 corpus.

Dataset	Document count			Word count		OOV %		
	Train	Val	Test	SD	RS	BSA	BART	PEGASUS
CNN/Dailymail	287,227	13,368	11,490	700	57	7.5	11.0	17.4
PAC-Summ	391,618	21,754	21,756	276	15	25.0	44.4	26.7
EBM	1423	209	424	298	58	14.3	11.5	18.2
BioASQ	1525	491	496	505	40	20.0	9.4	26.0
MeQSum	700	150	150	70	12	12.5	5.7	16.7
CHQSum	1000	107	400	184	12	8.3	6.3	12.5

Table 2: Dataset statistics of intermediate fine-tuning datasets (CNN/DailyMail, PAC) and downstream medical summarization datasets. *OOV%* refers to the median fraction of unigrams in RS that are absent from the PLM vocabulary.

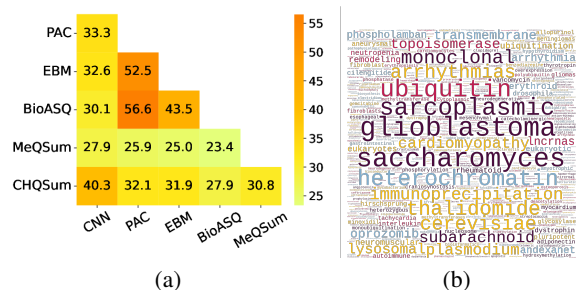


Figure 3: (a) Heatmap to show vocabulary overlap among different training datasets, computed based on the overlap between the top 10K most frequent words in each dataset. CNN corresponds to CNN/DailyMail dataset. (b) Words in the BioASQ dataset across three PLMs are split into four or more subwords; we observe that most of them are medical terms.

## 4.2 Evaluation Metrics

We report Rouge [Lin, 2004] and BertScore [Zhang *et al.*, 2020b] (BSr) to assess general summarization quality. We use Rouge-L as the main evaluation metric in line with prior works [Yuan *et al.*, 2022; Zhang *et al.*, 2023a]. *Concept F1-score* (CSr) is used to measure the faithfulness of medical summaries [Zhang *et al.*, 2023b] and is computed as the overlap of UMLS medical concepts between the generated and reference summaries. We extract the medical concepts using the QuickUMLS [Soldaini and Goharian, 2016] tool. We also discuss an additional metric (MedRouge) and parameters of Rouge in Appendix A.3.

## 4.3 Baseline Models

We provide further implementation details for the baseline models in Appendix A.1 and A.4.

**Vocabulary Adaptation Baseline Models.** We obtain the *BSA-PubMedBERT* baseline by replacing the encoder of BertSumAbs (BSA) with PubmedBERT [Gu *et al.*, 2021] respectively. In the same manner, we adapt a recent vocabulary adaptation model, AVocaDo [Hong *et al.*, 2021] from classification to the summarization setting. AVocaDo selects subwords from a vocabulary and iteratively builds on a downstream dataset until the fragment score stays above a particular threshold (taken as 3).

**Intermediate Fine-tuning Baseline Models.** BioBART [Yuan *et al.*, 2022] is obtained when BART is continuously pre-trained using PubMed abstracts corpora using only text-infilling as the pre-training objective. It

achieves state-of-the-art performance on the MeQSum and CHQSum datasets. *IFT-CNN* describes existing PLMs that only perform intermediate fine-tuning with the CNN/DailyMail dataset. Unlike *IFT-CNN*, *IFT-PAC* uses the PubMed Abstracts Collection (PAC) dataset for intermediate fine-tuning. *IFT-PAC* is equivalent to MEDVOC without vocabulary adaptation. *BSA-BioBERT* is obtained similar to *BSA-PubMedBERT* where the encoder of BertSumAbs (BSA) is replaced with BioBERT [Lee *et al.*, 2019].

## 4.4 Training Details

We obtain the biomedical paper abstracts from the official dump<sup>1</sup> of PubMed dated 2020. It comprises 450K data points (PubMed abstract and title). We refer to it as the *PubMed Abstracts Collection (PAC)* in the paper. We remove data points that overlap with downstream datasets (as EBM and BioASQ contain PubMed abstracts as the source document), as a decontamination step to prevent memorization issues [Radford *et al.*, 2019]. We randomly select 312K data points to form the final dataset. We keep the dataset size of PAC similar to that of CNN/Dailymail for a fair comparison. Appendix A.5 provides further implementation details and the optimal hyperparameters and vocabulary size details.

## 5 Experimental Results

We show the performance comparison results of MEDVOC in Table 4. We observe an average Rouge-L improvement of 15.74% across datasets over baselines in a zero-shot setting. We further observe gains of 10.81% and 17.29% across challenging scenarios of long-form medical summary generation and reference summaries with high OOV concentration respectively. The consistent improvement also holds for Concept Score that captures the faithfulness aspect [Zhang *et al.*, 2023b]. MEDVOC performs quite well for short-form summaries where it achieves SOTA performance on the MeQSum (even outperforming BioMedGPT [Zhang *et al.*, 2023a]) and CHQSum data.

### 5.1 Performance Evaluation of MEDVOC

We investigate the MEDVOC performance using five research questions (RQs).

**RQ1: MEDVOC outperforms vocabulary adaptation baselines.** We observe that MEDVOC outperforms the vocabulary adaptation baselines of BSA-PubMedBERT and BSA-AVocaDo by a good margin of 33.10% and 3.94% respectively in terms of Rouge-L. We thus observe the effectiveness of designing the vocabulary adaptation as a hyperparameter tuning search (as explained in Section 3.1). We show consistent improvement due to the vocabulary adaptation step alone (MEDVOC versus IFT-PAC) across three different PLMs in the case of EBM and BioASQ, where MEDVOC improves over IFT-PAC by 3.55% and 5.38% respectively. We further observe that MEDVOC outperforms proportionately to the percentage of OOV words (higher the OOV%, MEDVOC outperforms more). We thus observe limited improvement in the case of MeQSum and CHQSum

<sup>1</sup><https://ftp.ncbi.nlm.nih.gov/pubmed/updatefiles>

Model	MeQSum	CHQSum	EBM	BioASQ
BSA-PubMedBERT	39.79	30.59	17.76	26.65
BSA-AVocaDo	49.30	34.49	18.43	45.86
MEDVOC (BSA)	51.49	35.11	19.51	47.54

Table 3: Rouge-L comparison with vocabulary adaptation baselines. MEDVOC outperforms BSA-AVocaDo by 3.94% on average.




Figure 4: (a) Zero-shot and few-shot performance in terms of Rouge-L scores averaged across the three PLMs. MEDVOC shows statistically significant improvement over IFT-CNN in most settings, except BioASQ. (b) Performance improvement of MEDVOC over IFT-CNN in high OOV concentration and long-form reference summaries (top-ten percentile); PEGASUS shows the highest jump.

where the OOV percentage is 11.63% and 9.03%, which is much lower as compared to EBM and BioASQ where the OOV percentage is 14.67% and 18.46%. We extensively analyze the impact of the added vocabulary of MEDVOC and AVocaDo in terms of fragment score in the Appendix B.1.

**RQ2: MEDVOC outperforms baselines even in zero and few-shot summarization tasks.** We observe from Figure 4(a) that MEDVOC consistently outperforms IFT-CNN across the full dataset, leading to improved zero-shot and few-shot (10 and 100-shot) abstractive summarization for CHQSum, MeQSum and EBM, with average performance gains of 28.94% and 8.13% in terms of Rouge-L, respectively. Remarkably, the advantage of MEDVOC is more pronounced in zero and few-shot settings compared to training on the entire dataset (15.74% versus 5.99%). In contrast, for the BioASQ dataset, IFT-CNN exhibits higher zero-shot performance than MEDVOC (Rouge-L score of 35.36), increasing to 42.31 in the full data setting (MEDVOC outperforms in full data setting by 8.67%) as reference summaries of BioASQ is extractive (characterized by unigram and bigram overlaps between SD and RS of 96.72% and 84.30%, respectively), similar to the well-known extractive nature of CNN/DailyMail dataset.

**RQ3: MEDVOC outperforms baselines even when reference summaries have high OOV concentration.** We select the top ten percentile of data points that have the highest OOV concentration in reference summaries. These points represent the most difficult data points in terms of vocabulary mismatch. MEDVOC shows a high improvement of 17.29% on average over IFT-CNN in the BioASQ and EBM

datasets that have an average OOV concentration of 41.95% in the selected data points. Figure 4(b) shows that higher the OOV concentration, higher the performance improvement as the PEGASUS model type with BioASQ dataset shows the highest performance jump of 40.53% over IFT-CNN.

**RQ4: MEDVOC outperforms baselines in case of longer reference summaries.** Generating long-form medical summaries is under-explored [Liu *et al.*, 2023]. Here, we select data points where RS length is greater than 30 tokens (which is the 95 percentile of length of summaries for MeQSum and CHQSum) and limit our evaluation to EBM and BioASQ. We observe that across these two datasets and three model types, MEDVOC outperforms IFT-CNN by 10.81% on average, which is quite higher than 8.23% improvement on the entire dataset. Therefore, MEDVOC improves the generation of long-form medical summaries (Figure 4(b)).

**RQ5: MEDVOC trains the model decoder to incorporate more relevant medical words during generation and produce more faithful medical summaries.** We observe that BertSumAbs shows the highest performance improvement (based on Rouge-L) on average of 6.26% due to the vocabulary addition step, among the three PLMs. This is because the number of subwords required to generate a medical concept-bearing word (fragment score) significantly reduced from 2.09 to 1.58, the highest drop of 25% as compared to BART and PEGASUS of 3.47% and 3.14% respectively. We further observe that the *faithfulness* aspect measured using *Concept Score* (CSr) in a medical context [Zhang *et al.*, 2023b] improved significantly on average across EBM and BioASQ, by 8.72%, 6.76% respectively. This indicates that MEDVOC generates more faithful summaries. Furthermore, we observe that a large percentage of top-3 candidate beams in the BertSumAbs model type contains 85.92%, 85.37%, and 86.16% of newly added MEDVOC vocabulary in the case of first, second, and third candidate beam respectively. This further highlights the positive impact of vocabulary adaptation for medical abstractive summarization.

## 5.2 Ablation Analysis

We observe an average Rouge-L improvement of 29% for IFT-PAC when compared with the IFT-CNN model in the case of EBM and BioASQ, in the high OOV concentration in the RS setting (as done in RQ3), where PEGASUS model on BioASQ dataset shows the highest performance improvement of 64.51%, as also observed in RQ3.

**Biomedical paper title generations serve as a good intermediate fine-tuning task for medical summarization.** We observe the following characteristics of PAC that are similar to the downstream datasets. First, the source document length of EBM and BioASQ, of 276 and 505 words, is more similar to the PAC dataset (276 words) as compared to CNN/DailyMail (700 words). Second, in the case of MeQSum and CHQSum, the length of RS is 12 words on average, and it is almost the same as PAC (15 words) as compared to CNN/DailyMail (57 words). Third, the abstractive nature of summaries of PAC aligns more with the downstream datasets. PAC, EBM, and MeQSum have a bigram overlap between SD and RS of 33.33%, 15.39% and 10.10%, respectively.

Model	BertSumAbs					BART				PEGASUS					Overall		
	R-1	R-2	R-L	BSr	CSr	R-1	R-2	R-L	BSr	CSr	R-1	R-2	R-L	BSr	CSr	R-L	CSr
<b>EBM</b>																	
SOTA	25.40	7.06	18.88	85.19	18.44	29.41	9.15	20.62	85.97	24.14	25.46	6.82	18.09	85.75	20.31	19.19	20.96
IFT-CNN	26.37	6.37	18.79	84.76	18.14	27.06	7.66	19.08	85.76	21.72	25.46	6.82	18.09	85.75	20.31	18.65	20.06
IFT-PAC	27.22	7.53	19.26	84.79	19.58	28.30	7.98	19.78	85.84	21.82	26.69	7.93	19.01	85.47	22.78	19.35	21.39
MEDVOC	27.67	8.01	19.51	85.05	20.36	29.22	8.62	20.65	86.17	22.66	29.12	8.41	19.95	85.71	22.41	20.03	21.81
<b>BioASQ</b>																	
SOTA	49.05	37.15	45.84	90.55	50.57	51.78	39.91	47.36	90.77	52.05	44.63	29.77	39.48	89.44	47.82	44.23	50.14
IFT-CNN	45.65	33.48	42.17	89.27	44.61	48.84	37.41	45.29	90.31	49.48	44.63	29.77	39.48	89.44	47.82	42.31	47.30
IFT-PAC	49.58	37.82	44.76	89.81	50.96	50.32	38.26	45.00	90.45	49.53	45.37	34.80	41.06	89.70	46.85	43.60	49.02
MEDVOC	52.03	40.44	47.54	90.48	52.20	52.48	39.16	48.02	91.16	52.87	47.44	33.49	42.39	89.94	46.42	45.98	50.50
<b>MeQSum</b>																	
SOTA	52.64	37.66	49.99	93.55	53.56	55.53	40.31	52.67	93.99	58.10	53.87	38.65	51.03	93.88	55.84	51.23	55.83
IFT-CNN	46.92	30.53	44.33	91.72	47.48	59.49	43.24	56.16	94.83	60.65	53.87	38.65	51.03	93.88	55.84	50.51	54.63
IFT-PAC	49.44	33.31	46.28	92.89	49.76	59.09	42.76	55.73	93.90	61.30	58.24	43.45	55.39	94.31	61.07	52.47	57.38
MEDVOC	54.65	38.70	51.49	93.62	53.44	58.44	44.40	55.88	94.20	60.52	56.30	40.86	53.52	93.18	59.25	53.63	57.73
<b>CHQSum</b>																	
SOTA	35.99	16.96	33.72	91.01	33.45	40.44	21.04	38.51	91.98	38.73	43.07	24.11	40.44	92.04	42.33	37.56	38.17
IFT-CNN	37.81	19.14	35.64	91.02	34.82	41.07	22.18	39.02	92.09	42.19	43.07	24.11	40.44	92.04	42.33	38.36	39.78
IFT-PAC	38.58	18.56	36.24	91.36	34.28	40.53	21.16	38.75	91.92	41.02	42.73	23.55	40.35	91.92	41.44	38.45	38.91
MEDVOC	37.87	18.25	35.11	91.10	33.56	42.58	24.02	40.59	92.05	45.63	43.10	24.09	40.57	92.02	43.43	38.75	40.84

Table 4: Performance comparison of MEDVOC with Rouge-L (R-L) as the primary metric; we highlight the **best** and **second-best** settings. IFT-PAC is equivalent to MEDVOC without vocabulary adaptation. The improvements wherever observed in MEDVOC over IFT-CNN for **R-L** are **statistically significant** across all settings (using paired t-test;  $p < 0.01$ ). MEDVOC generates **more informative** (improves overall R-L by 5.99% on average) and **more faithful** medical summaries (improves overall Concept Score by 5.96% on average). For SOTA, we use BioBERT, BioBART, and IFT-CNN for BertSumAbs, BART, and PEGASUS respectively.

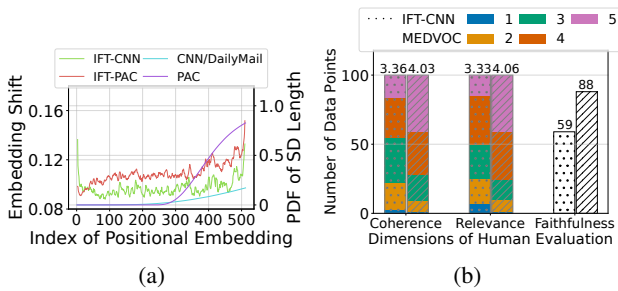


Figure 5: (a) Shift observed in the positional embedding for BertSumAbs in terms of Euclidean distance. (b) Human evaluation scores comparison over 100 randomly selected test data points. MEDVOC produces more relevant, coherent, and faithful summaries during human evaluation with medical experts.

**Biomedical paper title generation successfully aligns the positional embeddings to the target (medical) domain.** The most informative part of a PubMed abstract is predominantly located in the *Conclusion* and *Results* section [Mollá and Santiago-Martínez, 2011; Jin *et al.*, 2019], which naturally makes these sections more probable and desirable to be a part of a reference summary. Therefore, we observe in Figure 5(a) that the positional embedding shift due to IFT-CNN over occurs the most at the start, whereas a consistently higher shift is seen for IFT-PAC after the initial around fifty tokens and the highest shift occurs towards the end of a SD of a PubMed abstract (token index  $\geq 400$ ). Thus, the domain mismatch is well-captured by positional embedding shift.

### 5.3 Human Evaluation

We randomly select 100 test data points uniformly across the four datasets and use the Prolific platform to recruit medi-

cal experts (self-reported to have a Masters or Doctoral degree in Medicine and Biomedical Sciences and are older than 24 years) for annotating summary pairs of MEDVOC and IFT-CNN across the standard aspects [Fabbri *et al.*, 2021b; Zhang *et al.*, 2023b] of *relevance*, *coherence* (on a Likert scale of 1 to 5), and *faithfulness* (binary). Each annotator was given 30 minutes to evaluate 10 summaries and was compensated at a rate of 9 UK pounds per hour (see Appendix B.2 for more details). Figure 5(b) shows the human evaluation results where MEDVOC generates more faithful summaries (88% versus 59% of summaries are faithful), and more relevant summaries, where 76% of data points get a positive score of 4 or 5 in Likert scale, as compared to 50% by IFT-CNN.

## 6 Conclusion

We present a dynamic vocabulary adaptation strategy, called MEDVOC, for fine-tuning PLMs for improved medical text summarization. To the best of our knowledge, this is the first work that uses vocabulary adaptation techniques for summarization and opens up an interesting potential research direction. Through extensive experimentation, we observe that MEDVOC consistently outperforms vocabulary adaptation baselines and significantly outperforms standard fine-tuning strategy (IFT-CNN) in full data setting. MEDVOC outperforms even in zero and few-shot settings, as well as when reference summaries have high OOV concentration or are very long ( $> 30$  words). MEDVOC outperforms IFT-CNN by a high margin in terms of relevance and faithfulness in human evaluation with medical experts. As an immediate future work, we will extend to the multi-document summarization setting. Given that MEDVOC leads to more faithful summaries, we will incorporate vocabulary adaptation to improve state-of-the-art models [Alambo *et al.*, 2022; Zhang *et al.*, 2023b; Xie *et al.*, 2023a] that improve the factual consistency of summaries.

## Acknowledgements

We thank the Ministry of Education, Govt of India, for supporting Gunjan Balde with Prime Minister Research Fellowship during his Ph.D. tenure. Soumyadeep Roy is supported by the Institute Ph.D. Fellowship at the Indian Institute of Technology Kharagpur. This research was partially funded by the Federal Ministry of Education and Research (BMBF), Germany under the project LeibnizKILabor with grant No. 01DD20003 and by a Google India Faculty Research Award.

## Contribution Statement

Gunjan Balde and Soumyadeep Roy make equal contributions and share co-first authorship. Gunjan Balde is the corresponding author of this paper.

## References

- [Alambo *et al.*, 2022] Amanuel Alambo, Tanvi Banerjee, et al. Improving the factual accuracy of abstractive clinical text summarization using multi-objective optimization. In *2022 44th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, pages 1615–1618. IEEE, 2022.
- [Allison *et al.*, 2006] Ben Allison, David Guthrie, et al. Another look at the data sparsity problem. In *Text, Speech and Dialogue: 9th International Conference, TSD 2006, Brno, Czech Republic, September 11-15, 2006. Proceedings 9*, pages 327–334. Springer, 2006.
- [Ben Abacha and Demner-Fushman, 2019] Asma Ben Abacha and Dina Demner-Fushman. On the summarization of consumer health questions. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2228–2234, July 2019.
- [Chang and Lu, 2021] Ting-Yun Chang and Chi-Jen Lu. Re-thinking why intermediate-task fine-tuning works. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 706–713, November 2021.
- [Chen *et al.*, 2023] Zeming Chen, Alejandro Hernández Cano, et al. Meditron-70b: Scaling medical pretraining for large language models. *arXiv preprint arXiv:2311.16079*, 2023.
- [Dai *et al.*, 2021] Songtai Dai, Quan Wang, et al. BDKG at MEDIQA 2021: System report for the radiology report summarization task. In *Proceedings of the 20th Workshop on Biomedical Language Processing*, pages 103–111, June 2021.
- [Diao *et al.*, 2021] Shizhe Diao, Ruijia Xu, et al. Taming pre-trained language models with n-gram representations for low-resource domain adaptation. 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 3336–3349, August 2021.
- [Fabbri *et al.*, 2021a] Alexander Fabbri, Simeng Han, et al. Improving zero and few-shot abstractive summarization with intermediate fine-tuning and data augmentation. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 704–717, June 2021.
- [Fabbri *et al.*, 2021b] Alexander R. Fabbri, Wojciech Kryściński, et al. SummEval: Re-evaluating summarization evaluation. *Transactions of the Association for Computational Linguistics*, 9:391–409, 2021.
- [Gu *et al.*, 2021] Yu Gu, Robert Tinn, Cheng, et al. Domain-specific language model pretraining for biomedical natural language processing. *ACM Trans. Comput. Healthcare*, 3(1), oct 2021.
- [He *et al.*, 2021] Yifan He, Mosha Chen, et al. damo\_nlp at MEDIQA 2021: Knowledge-based preprocessing and coverage-oriented reranking for medical question summarization. In *Proceedings of the 20th Workshop on Biomedical Language Processing*, pages 112–118, June 2021.
- [Hong *et al.*, 2021] Jimin Hong, TaeHee Kim, et al. AVocaDo: Strategy for adapting vocabulary to downstream domain. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4692–4700, November 2021.
- [Jin *et al.*, 2019] Qiao Jin, Bhuwan Dhingra, et al. PubMedQA: A dataset for biomedical research question answering. In *Proceedings of the EMNLP-IJCNLP 2019*, pages 2567–2577, November 2019.
- [Kanwal and Rizzo, 2022] Neel Kanwal and Giuseppe Rizzo. Attention-based clinical note summarization. In *Proceedings of the 37th ACM/SIGAPP Symposium on Applied Computing, SAC '22*, page 813–820, 2022.
- [Lamproudis. *et al.*, 2022] Anastasios Lamproudis., Aron Henriksson., et al. Vocabulary modifications for domain-adaptive pretraining of clinical language models. In *Proceedings of the 15th International Joint Conference on Biomedical Engineering Systems and Technologies (BIOSTEC 2022) - HEALTHINF*, pages 180–188. INSTICC, 2022.
- [Laskar *et al.*, 2022] Md Tahmid Rahman Laskar, Enamul Hoque, et al. Domain Adaptation with Pre-trained Transformers for Query-Focused Abstractive Text Summarization. *Computational Linguistics*, 48(2):279–320, 06 2022.
- [Lee *et al.*, 2019] Jinhyuk Lee, Wonjin Yoon, et al. BioBERT: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, 36(4):1234–1240, 09 2019.
- [Lewis *et al.*, 2020] Mike Lewis, Yinhan Liu, et al. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, 2020.
- [Lin, 2004] Chin-Yew Lin. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81, 2004.



- [Liu and Lapata, 2019] Yang Liu and Mirella Lapata. Text summarization with pretrained encoders. In *Proceedings of EMNLP-IJCNLP*, pages 3730–3740, 2019.
- [Liu *et al.*, 2023] Siyang Liu, Naihao Deng, et al. Task-adaptive tokenization: Enhancing long-form text generation efficacy in mental health and beyond. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 15264–15281, 2023.
- [Mollá and Santiago-Martínez, 2011] Diego Mollá and María Elena Santiago-Martínez. Development of a corpus for evidence based medicine summarisation. In *Proceedings of the Australasian Language Technology Association Workshop*, pages 86–94, 2011.
- [Phang *et al.*, 2019] Jason Phang, Thibault FÉvry, et al. Sentence encoders on stilts: Supplementary training on intermediate labeled-data tasks, 2019.
- [Radford *et al.*, 2019] Alec Radford, Jeffrey Wu, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- [Rust *et al.*, 2021] Phillip Rust, Jonas Pfeiffer, et al. 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, August 2021.
- [Schick and Schütze, 2020] Timo Schick and Hinrich Schütze. Rare words: A major problem for contextualized embeddings and how to fix it by attentive mimicking. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):8766–8774, Apr. 2020.
- [See *et al.*, 2017] Abigail See, Peter J. Liu, et al. Get to the point: Summarization with pointer-generator networks. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, Volume 1: Long Papers*, pages 1073–1083, 2017.
- [Soldaini and Goharian, 2016] Luca Soldaini and Nazli Goharian. Quickumls: a fast, unsupervised approach for medical concept extraction. In *MedIR workshop, SIGIR*, pages 1–4, 2016.
- [Suresh *et al.*, 2023] Shilpa Suresh, Nazgol Tavabi, et al. Intermediate domain finetuning for weakly supervised domain-adaptive clinical NER. In *The 22nd Workshop on Biomedical Natural Language Processing and BioNLP Shared Tasks*, pages 320–325, July 2023.
- [Tai *et al.*, 2020] Wen Tai, H. T. Kung, et al. exBERT: Extending pre-trained models with domain-specific vocabulary under constrained training resources. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1433–1439, November 2020.
- [Tsatsaronis *et al.*, 2015] George Tsatsaronis, Georgios Balikas, et al. An overview of the bioasq large-scale biomedical semantic indexing and question answering competition. *BMC Bioinformatics*, 16:138, 2015.
- [Wu *et al.*, 2023] Chaoyi Wu, Weixiong Lin, et al. Pmc-llama: Towards building open-source language models for medicine, 2023.
- [Xie *et al.*, 2022] Qianqian Xie, Jennifer Amy Bishop, et al. Pre-trained language models with domain knowledge for biomedical extractive summarization. *Knowledge-Based Systems*, 252:109460, 2022.
- [Xie *et al.*, 2023a] Qianqian Xie, Jinpeng Hu, et al. Factreranker: Fact-guided reranker for faithful radiology report summarization. *arXiv preprint arXiv:2303.08335*, 2023.
- [Xie *et al.*, 2023b] Qianqian Xie, Zheheng Luo, et al. A survey for biomedical text summarization: From pre-trained to large language models, 2023.
- [Xu *et al.*, 2021] Jingjing Xu, Hao Zhou, et al. Vocabulary learning via optimal transport for neural machine translation. 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 7361–7373, August 2021.
- [Xu *et al.*, 2023] Benfeng Xu, Chunxu Zhao, et al. Retrieval-augmented domain adaptation of language models. In *Proceedings of the 8th Workshop on Representation Learning for NLP (RepLANLP 2023)*, pages 54–64, July 2023.
- [Yadav *et al.*, 2022] Shweta Yadav, Deepak Gupta, et al. Chq-summ: A dataset for consumer healthcare question summarization. *arXiv preprint arXiv:2206.06581*, 2022.
- [Yuan *et al.*, 2022] Hongyi Yuan, Zheng Yuan, et al. Bio-BART: Pretraining and evaluation of a biomedical generative language model. In *Proceedings of the 21st Workshop on Biomedical Language Processing*, pages 97–109, May 2022.
- [Zhang *et al.*, 2020a] Jingqing Zhang, Yao Zhao, et al. PE-GASUS: Pre-training with extracted gap-sentences for abstractive summarization. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 11328–11339, 13–18 Jul 2020.
- [Zhang *et al.*, 2020b] Tianyi Zhang, Varsha Kishore, et al. Bertscore: Evaluating text generation with BERT. In *8th International Conference on Learning Representations, ICLR*, 2020.
- [Zhang *et al.*, 2021] Longxiang Zhang, Renato Negrinho, et al. Leveraging pretrained models for automatic summarization of doctor-patient conversations. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 3693–3712, November 2021.
- [Zhang *et al.*, 2023a] Kai Zhang, Jun Yu, et al. Biomedgpt: A unified and generalist biomedical generative pre-trained transformer for vision, language, and multimodal tasks. *arXiv preprint arXiv:2305.17100*, 2023.
- [Zhang *et al.*, 2023b] Nan Zhang, Yusen Zhang, et al. Fame-sum: Investigating and improving faithfulness of medical summarization. *arXiv preprint arXiv:2311.02271*, 2023.