

Pre-trained Language Models in Biomedical Domain: A Systematic Survey

BENYOU WANG, SRIBD & SDS, The Chinese University of Hong Kong, Shenzhen, China

QIANQIAN XIE*, Department of Computer Science, University of Manchester, United Kingdom

JIAHUA PEI, University of Amsterdam, Netherlands

ZHIHONG CHEN, SRIBD & SSE, The Chinese University of Hong Kong, Shenzhen, China

PRAYAG TIWARI, School of Information Technology, Halmstad University, Sweden

ZHAO LI, The University of Texas Health Science Center at Houston, USA

JIE FU, Mila, University of Montreal, Canada

Pre-trained language models (PLMs) have been the de facto paradigm for most natural language processing (NLP) tasks. This also benefits the biomedical domain: researchers from informatics, medicine, and computer science (CS) communities propose various PLMs trained on biomedical datasets, e.g., biomedical text, electronic health records, protein, and DNA sequences for various biomedical tasks. However, the cross-discipline characteristics of biomedical PLMs hinder their spreading among communities; some existing works are isolated from each other without comprehensive comparison and discussions. It is nontrivial to make a survey that not only systematically reviews recent advances in biomedical PLMs and their applications but also standardizes terminology and benchmarks. This paper summarizes the recent progress of pre-trained language models in the biomedical domain and their applications in downstream biomedical tasks. Particularly, we discuss the motivations of PLMs in the biomedical domain and introduce the key concepts of pre-trained language models. We then propose a taxonomy of existing biomedical PLMs, which categorizes them from various perspectives systematically. Plus, their applications in biomedical downstream tasks are exhaustively discussed, respectively. At last, we illustrate various limitations and future trends, which aims to provide inspiration for the future research of the research community.

CCS Concepts: • **Computing methodologies** → **Natural language processing**; **Natural language generation**; **Neural networks**; **Bio-inspired approaches**.

Additional Key Words and Phrases: Biomedical domain, pre-trained language models, natural language processing

ACM Reference Format:

Benyou Wang, Qianqian Xie, Jiahua Pei, Zhihong Chen, Prayag Tiwari, Zhao Li, and Jie Fu. 2021. Pre-trained Language Models in Biomedical Domain: A Systematic Survey. 1, 1 (July 2021), 57 pages. <https://doi.org/10.1145/nnnnnnnn.nnnnnnnn>

*Qianqian Xie is the corresponding author: xqq.sincere@gmail.com.

Authors' addresses: Benyou Wang, wangbenyou@cuhk.edu.cn, SRIBD & SDS, The Chinese University of Hong Kong, Shenzhen, China; Qianqian Xie, qianqian.xie@manchester.ac.uk, Department of Computer Science, University of Manchester, United Kingdom; Jiahua Pei, j.pei@uva.nl, University of Amsterdam, Netherlands; Zhihong Chen, zhihongchen@link.cuhk.edu.cn, SRIBD & SSE, The Chinese University of Hong Kong, Shenzhen, China; Prayag Tiwari, prayag.tiwari@ieee.org, School of Information Technology, Halmstad University, Sweden; Zhao Li, lizhao.informatics@gmail.com, The University of Texas Health Science Center at Houston, USA; Jie Fu, jie.fu@polymtl.ca, Mila, University of Montreal, Canada.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2021 Association for Computing Machinery.

Manuscript submitted to ACM

CONTENTS

Abstract	1
Contents	2
1 Introduction	3
2 Background: Pre-trained Language Models	7
2.1 Backbone Networks in Language Models	7
2.2 Pre-training for texts	9
2.3 Pre-training for images	12
2.4 Fine-tuning Paradigm in PLMs	12
3 PLMs in Biomedical Domain	13
3.1 Motivation	13
3.2 Biomedical Data for Pre-training	14
3.3 How to tailor PLMs to the Biomedical Domain	17
3.4 Biomedical Pre-trained Language Models	18
3.5 Beyond Text: Biomedical Vision-and-Language Models	21
3.6 Beyond Text: Language Models for Proteins/DNA	23
4 Fine-tuning PLMs for Biomedical Downstream Tasks	25
4.1 Information Extraction	26
4.2 Text Classification	30
4.3 Sentence Similarity	31
4.4 Question Answering	32
4.5 Dialogue Systems	33
4.6 Text Summarization	35
4.7 Natural Language Inference	36
4.8 Proteins/DNAs Prediction	37
4.9 Competitions and Venues	38
5 Discussion	40
5.1 Limitations and Concerns	40
5.2 Future trends	41
6 Conclusion	44
References	44

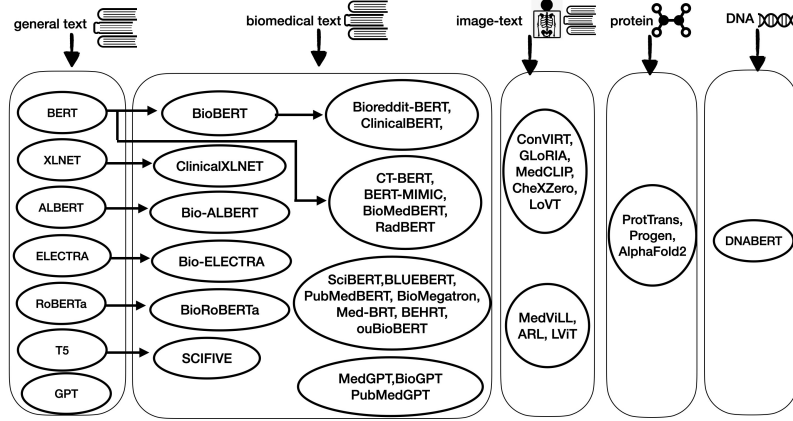


Fig. 1. Overview of selected released Biomedical pre-trained language models. One can see a more detailed list in Sec. 3. Note that there is a BERT-like language model embedded in the overall architecture of AlphaFold 2.

1 INTRODUCTION

As the principal method of communication, humans usually record information and knowledge in a format of *token* sequences, e.g., natural languages, time series, constructed knowledge base, etc. For biomedical information and knowledge, tokens in sequences could be of various types, including words, disease codes, amino acids, and DNAs. Tremendous biomedical information and knowledge in nature and human history are implicitly encapsulated in these natural token sequences in nature (*a.k.a.*, data).

There exist many data that involve biomedical information with different abstraction degrees of biomedical knowledge. However, there is a trade-off between the high abstraction degree and its scale. For data that explicitly conveys biomedical knowledge (i.e. at a high abstraction degree), it is usually small-scaled, see biomedical knowledge bases and EHR data (maybe in multi-modality). One example of data that may not directly convey biomedical knowledge could be protein and DNA sequences, since one can hardly know what a short protein or DNA sequence really means for humans and it needs more effort for abstraction. Fortunately, these data are usually tremendous. In the current stage, existing work pays more attention to data at a high abstraction level (biomedical knowledge-intensive data, e.g., EHR, biomedical knowledge bases, and biomedical encyclopedia); however, it is usually relatively small-scale. We argue that biomedical knowledge on various abstraction degrees should be paid attention to. To capture and mine the biomedical information and knowledge from various abstraction degrees, there is recently growing attention in the biomedical natural language processing (NLP) community to adopt pre-trained language models (PLMs); since PLMs could leverage these massive sequences without biomedical knowledge abstraction and human annotations, including but not limited to plain biomedical text, biomedical images, general text, protein sequences, and DNA sequences.

The biomedical NLP is a cross-discipline research direction from various communities such as bioinformatics, medicine, and computer science (especially a major frontier of artificial intelligence, *i.e.*, natural language processing *a.k.a.* NLP). The computational biology community [142] and biomedical informatics community [57] have made a substantial effort to make use of NLP tools for information mining and extraction of widespread-adopted electronic health records, medical scientific publications, medical WIKI pages, etc. For many decades, NLP has been investigating various biomedical tasks [56, 58] such as classification, information extraction, question answering, drug discovery et al.

Meanwhile, the approaches in the NLP community are changing rapidly, as one can witness exponentially increasing submitted papers in top conferences like ACL, EMNLP, and NAACL. Tailoring these NLP approaches that have been evidenced effectively in the NLP community to a specific biomedical domain is beneficial.

Unfortunately, there is usually a delay for newly proposed NLP approaches being applied to the biomedical domain. Especially, since the adoption of various pre-trained language models (e.g., ELMo [230], GPT [240], BERT [66], XLNET [114], RoBERTa [185], T5 [241] and ELECTRA [54]) [237] have nearly shifted the paradigm in NLP, their biomedical variants trained using biomedical data comes sooner or later. With this hot trend of the biomedical pre-trained language model, this survey aims to bridge the gap between pre-trained language models and their applications in the biomedical domain.

Motivation of pre-trained language models in biomedical domain. The current NLP paradigm is gradually shifting to a two-stage (pre-training and fine-tuning) paradigm, thanks to recently proposed pre-trained language models. Compared to the previous paradigm with purely supervised learning that relies on feature engineering or neural network architecture engineering [182], the current two-stage paradigm is more friendly to the scenario when supervised data is limited while large-scaled unsupervised data is tremendous. Fortunately, the biomedical domain is a typical case of such a scenario.

The motivation to use pre-trained language models in the biomedical domain is pretty straightforward. First, annotated data in the biomedical domain is usually not large-scale. Therefore, a well-trained pre-trained language model is more crucial to provide a richer feature extractor, which may slightly reduce the dependence on annotated data. Second, the biomedical domain is more knowledge-intensive than the general domain. At the same time, pre-trained language models could serve as an easily-used soft knowledge base [231] that captures implicit knowledge from large-scale plain documents without human annotations. More recently, GPT3 has been shown to have the potential to ‘remember’ many complicated common knowledge [38]. Lastly, large-scaled biomedical corpora and biomedical sequences (including proteins and DNAs), which are previously thought as difficult to handle, can be effectively handled by pre-trained language models (especially transformers networks).

As shown in Fig. 2, in recent three years, we have witnessed a rapid development of pre-trained language models (e.g., ELMo [230], GPT [240], BERT [66], XLNet [114], RoBERTa [185], T5 [241] and ELECTRA [54]) in the general NLP domain. Following these progresses, there are efforts to tailor these pre-trained language models to their corresponding biomedical variants, via in-domain data. For example, BERT, the most typical pre-trained language, has many variants in the biomedical domain, e.g., Med-BERT [248], BioBERT [156], publicly available Clinical BERT Embeddings [13], SciBERT [23], ClinicalBERT [113], and COVID-twitter-BERT [210] et al. We draw an overview for these models in Fig. 1. It shows that the extensions of general domain pre-trained language models to the biomedical domain attract great attention from researchers in both NLP and bioinformatics communities. Interestingly, we can observe that once the general NLP community develops a new variant of PLM, it usually leads to a biomedical counterpart after some months. This parallel development between general PLMs and biomedical PLMs shows a strong demand and even a necessity to summarize the existing works, which could help beginners to start their contributions in this field easily.

Difference with existing surveys. There are a few reviews to summarize the NLP applications in the biomedical, clinical, bioinformatic domain, such as an early one [276] and recent ones [228, 326, 360]. They cover many general methods and applications of biomedical/clinical NLP. Specifically, [276] mainly discuss either based on statistics-based NLP pipeline (including lexicon, co-occurrence patterns, syntactic/semantic parsing), or word embeddings based neural network approaches (it was mentioned that 60.8% of them are based on recurrent neural networks) [326] for NLP




Fig. 2. Parallel development of general and biomedical pre-trained language models. The time is determined by the released date of the paper, for example, in arXiv. General pre-trained language models are shown below in the timeline, and biomedical pre-trained language models are shown above the timeline (refer to Tab. 4 for detailed dates).

applications (e.g., information extraction, text classification, named entity recognition, and relation extraction et al). Especially, two reviews [130, 134] discuss the word embeddings used in biomedical NLP.

All the above reviews made thorough summarization of existing work before the pre-trained language model era of NLP. The NLP techniques in these reviews are mainly about feature engineering, or architecture engineering [182]. However, the NLP recently has been shifted to a pre-training and then fine-tuning paradigm with large-scale pre-trained language models (see existing surveys [32, 96, 182, 183, 237] for pre-trained language model in the general domain). [32] called these pre-trained models as ‘foundation models’ to underscore their critically central. We believe the biomedical NLP applications have benefited and will continually benefit from the development of pre-trained language models.

More recently, [129] reviews biomedical textual pre-training, especially using BERT. The difference between [129] and this review is that Our paper provides a more inclusive taxonomy of biomedical PLMs than [129], which are three fold.. First, biomedical PLMs summarized in our review are not limited to that trained on texts like [129], but also other data resources including protein, DNA, and even biomedical text-image pairs. In general, any data that involves biomedical information could be used in biomedical PLMs. Second, in contrast to [129] which only discusses Transformer-based pre-trained language models, this review also discusses RNN-based language models (like ELMO [121], which is typically considered as the first pre-trained language model in NLP). We also summarize decoder involved *generative* pre-trained language models (like GPT [146] and T5 [232]), while [129] mainly discusses encoder-based PLMs (BERT or BERT variants). Third, to the best of our knowledge, this is the first survey paper to discuss pre-trained **vision-language** models in the biomedical domain. Last, our paper provides a more comprehensive overview of the applications of PLMs in the biomedical domain compared with [129]. Except for biomedical NLP tasks such as natural language inference, text summarization [334], relation extraction et al that are summarized in [129], our paper further reviews recent PLMs-based methods for event detection, dialogue systems, as well as protein and DNA sequence. Moreover, compared with [129] that only reviews recent methods of biomedical NLP tasks coarsely, we make a thorough categorization and discussion of PLMs-based methods for biomedical NLP tasks and their benchmark datasets. Our paper also introduces competitions and venues such as shared tasks. Therefore, we believe there is a requirement for a more thorough survey paper to review the recent progress of pre-trained language models in the biomedical domain from a multi-scale perspective.

Contribution. The contributions of the paper can be summarized as follows:

- We give a comprehensive review to summarize existing PLMs-based methods for the biomedical domain, which thoroughly categorizes and discusses biomedical data sources, biomedical PLMs, model variants, downstream tasks, shared competitions, etc.



Fig. 3. Architecture of this survey.

- We propose a taxonomy of biomedical PLMs, which classifies existing PLMs in the biomedical domain from various perspectives: training data sources, model architecture, etc.
- We enumerate existing resources for PLMs and their detailed configuration, facilitating their spreading for beginners.
- We discuss the limitations of existing methods and prospect future trends.
- To the best of our knowledge, this is the first survey paper to summarize generative pre-trained language models, protein/DNA language models, pre-trained **vision-language** models in the biomedical domain.

How do we collect the papers? In this survey, we collected over a hundred related papers. We used Google Scholar as the main search engine, and also adopted MedPub, Web of Science, as an essential tool to discover related papers. In addition, we screened most of the related conferences and journals such as ACL, EMNLP, NAACL, AAAI, Bioinformatics, JAMIA, AMIA, etc. The major keywords we used included medical pre-trained language model, clinical pre-trained language model, biological language model, etc. Plus, we take Med-BERT [248], BioBERT [156], SciBERT [23], ClinicalBert [113], COVID-twitter-BERT [210] as the seed papers to check papers that cited them.

Organization. The overall architecture of this paper is shown in Figure 3. The paper is organized as below: Sec.2 introduces the general pre-trained language models including backbone networks, pre-training objective, pre-training

corpora, fine-tuning, and categorization of PLMs. Sec.3 introduces the pre-trained language models for the biomedical domain and proposes a taxonomy, including motivations for using PLMs, biomedical data sources, domain-specific pre-training, biomedical PLMs, and their categorization. Sec.4 summarizes the applications of biomedical PLMs for various downstream tasks and categorizes existing methods for these tasks respectively. More discussions about limitations and future directions are in Sec. 5. We conclude in Sec.6.

2 BACKGROUND: Pre-trained Language Models

Pre-trained language models (PLMs) have been widely used in natural language processing, etc., due to their effectiveness to learn useful representations from unannotated data such as natural languages. In this paper, we mainly discuss pre-trained language in sequential tokens¹. We will introduce the textual pre-training in Sec. 2.2, one can read the review paper of PLMs in [237] for more details. Thanks to the popularity of CLIP, pre-trained language models are also usually jointly trained with a visual pre-trained model in the image-text pre-training scenario. We will also discuss visual pre-training in Sec. 2.3. Note that models in the visual pre-training usually treats image patches as visual tokens, this makes it language model-like pre-training; we, therefore, include visual pre-training models in this survey.

In this section, we will introduce the basic ingredients of pre-training models: the training objective with self-supervised tasks and corpora in Sec. 2.2 and Sec. 2.3 for text and images respectively, basic neural network models in Sec. 2.1, and training paradigm in Sec. 2.4.

2.1 Backbone Networks in Language Models

The success of pre-trained language models is also attributed to the development of their base backbone network, from LSTM [108] to Transformer [301]. Before Transformer was invented, LSTM was widely used as the base architecture of pre-trained language models such as ELMO. However, because of its recurrence structure, it is computationally expensive to scale up LSTM to be deeper in layers. To this end, Transformer is proposed and becomes the backbone of modern NLP. Transformers are better architecture can be attributed to: 1) efficiency: a recurrent-free architecture that could compute the individual token in parallel, 2) effectiveness: attention allows spatial interaction across tokens that dynamically depends on the input itself. In this section, we briefly introduce the two typical architectures in pre-trained language models, namely, LSTM and Transformers.

2.1.1 Previous backbone networks in texts.

LSTM. Long short-term memory (LSTM) is a recurrent neural network (RNN) architecture for sequential modeling. Unlike standard feed-forward neural networks processing single data points (such as images), LSTM can deal with entire sequences of data (such as text, speech, or video). A common LSTM unit is composed of a cell, an input gate, an output gate, and a forget gate. The cell learns hidden states over arbitrary time intervals and the three gates regulate the flow of information into and out of the cell. LSTM networks are well-suited for time series data and were developed to deal with the vanishing gradient problem that can be encountered when training traditional RNNs. Peters et al [230] tried to adopt a Long and Short term memory network (LSTM) in pre-trained language, which naturally processes tokens sequentially.

2.1.2 Previous backbone networks in images.

¹Tokens usually refers to words or subwords in NLP, and also protein sequences in the biomedical domain.

CNNs. Convolutional neural networks [155] (CNNs) are a type of neural networks that are particularly suited for vision tasks. Typically, CNNs are made up of four main types of layers: convolution, pooling, activation, and fully connected layers. The convolution layers are trainable filters that can learn to recognize patterns in images, such as edges, textures, and objects; The pooling layers are used to reduce the dimensionality of the data; The activation layers are used to introduce non-linearity to the network; The fully connected layers are used to make predictions based on the extracted features. Note that CNNs are also a good choice for language understanding [139].

2.1.3 The current backbone networks in texts and images.

Transformer. The backbone of most pre-trained language models (e.g., BERT, its variants, GPT, T5 et al) is a neural network called ‘Transformer’ building upon self-attention networks (SANs) and feed-forward networks (FFNs). SAN is used to facilitate interaction between tokens, while FFN is used to refine the token presentation using non-linear transformation. Since Transformer has been the de facto backbone to replace recurrent and convolutional units, almost all language models adopt the Transformer as the backbone network. The transformer is superior in terms of capacity and scalability thanks to, 1) discarding recurrent units and process tokens more efficiently in parallel with the position embeddings[308, 309], 2) relieving saturation issue of expressive power with large-scale data and very deep layers due to the well-designed architecture including residual connections, layer normalization, and etc.

A Transformer layer consists of a self-attention (SAN) module and a feed-forward network (FFN) module. An input X ² for SAN will be linearly transformed into query, key, value, and output space $\{Q, K, V\}$ as below³:

$$\begin{bmatrix} Q \\ K \\ V \end{bmatrix} = X \times \begin{bmatrix} W^Q \\ W^K \\ W^V \end{bmatrix} \quad (1)$$

The self-attention mechanism (a.k.a Scaled Dot-Product Attention) is calculated as

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK}{\sqrt{d_k}}\right)V \quad (2)$$

For a multi-head version of the self-attention mechanism, it linearly projects Q, K, V with h times using individual linear projections to smaller dimensions (e.g. $d_k = \frac{d_{\text{model}}}{h}$), instead of performing a single attention function with d_{model} -dimensional keys, values and queries. Finally, the output of SAN is

$$\begin{aligned} \text{SAN}(X) &= [\text{head}_1; \dots; \text{head}_h]W^O \\ \text{head}_i &= \text{Attention}(Q_i, K_i, V_i), \end{aligned} \quad (3)$$

where $Q = [Q_1; \dots; Q_h]$, $K = [K_1; \dots; K_h]$, and $V = [V_1; \dots; V_h]$. The individual attention heads are independently calculated. Since the output of SAN is a linear transformation (using W^O) of V , which is a weighted sum of V . A stack of many purely SAN layers is not expressive [71], since it is equivalent to a single linear transformation. To this end, a feed-forward network with non-linear activation is alternately used with each SAN layer,

$$\text{FFN}(X) = \delta(XW^{\text{in}})W^{\text{out}}. \quad (4)$$

Since some neurons after the activation function (e.g., δ is ReLU or GELU [104]) become inactivated (zero), d_{in} is usually bigger than d_{model} to avoid the low-rank bottleneck, typically, $d_{\text{in}} = 4 \times d_{\text{model}} = d_{\text{out}}$. Other tricks, such

² X is the word embedding of each individual input token which are tokenized using subword tokenization. Moreover, the input is usually concatenated with position embeddings [317] to perceive word order

³For all linear transformation in this paper, the bias term is in default omitted

as layer normalization, residual connection, dropout, and weight decay are also adopted to relieve the optimization and overfitting problems when it goes deeper, resulting in better stability when training large neural networks. It is generally believed that Transformer is better than LSTM in terms of generalization since its performance usually does not get to saturation as early as LSTM. When models become large, the performance of the Transformer is consistently increasing when feeding more data while LSTM gets saturation if a certain amount of data is fed.

Interestingly, the computer vision [171] and computational biology communities also borrow some insights to design their models, see ViT [171] for vision and AlphaFold2 [127] for protein. In Table ??, we introduce some typical pre-trained language models in general NLP domains, based on these two backbone neural networks.

2.2 Pre-training for texts

Previously, there were many typical methods to build token representation (e.g., word vectors) from plain corpora. For example, [200, 227] build a one-to-one mapping between words and their vectors, which is called ‘static word embedding’ since it is static and not related to word context. However, it is well known that words often express different meanings in different contexts. To achieve this, most recently many pre-trained language models [230] are proposed to learn ‘contextualized word embedding’ that models the bi-directional contexts of words. For ‘contextualized word embedding’, the vector for a word depends on its specific usage in a context. For example, the meanings of ‘bank’ in ‘river bank’ and in ‘money bank’ are supposed to have some difference. Compared with ‘static word embedding’, the ‘contextualized word embedding’ largely improves the quality of word representation in various tasks [66].

A *language model* aims to assign a probability to a given piece of text (e.g., a sentence or an n-gram.) [128], see below:

$$\Theta : \mathbb{V}^N \rightarrow \mathbb{R}^+ \quad (5)$$

While, in the scenario of natural language processing, a generally-called *language model* is usually a *conditional language model* that assigns a probability to a next word w_n given some conditioning context (denoted as $[w_1, \dots, w_{n-1}]$). A conditional language model is a generalization of *language model* in a sense the former could be obtained by dividing the probability of the concatenated sentence (i.e., $[w_1, \dots, w_{n-1}, w_n]$) by that of the context, namely

$$P(w_n | w_1, \dots, w_{n-1}, w_n) = \frac{\Theta(w_1, \dots, w_{n-1})}{\Theta(w_1, \dots, w_{n-1})} \quad (6)$$

In the earliest, neural language models [25, 201] and their variants such as Skip-Gram [200], CBow [200] and Glove [227], were the backbones of modern NLP to provide pre-trained word features. The pre-training task of classical neural language models [25] is the unidirectional language modeling (ULM), that predicts the next word conditionally on history words. To learn better word embeddings, several classical models further improved the pre-training task. For example, the training objective of Skip-Gram [200] is predicting context words given the input word. CBow [200] aims to predict the next word based on its bidirectional context words. The training task of Glove [227] is to predict the log co-occurrence of words. These models typically use shallow neural network architecture to conduct calculations between word vectors, for efficient training.

Language models could be considered as an instance of self-supervision. Compared to data-hungry supervised learning, which usually needs annotations from humans, language models could make use of massive amounts and cheap plain corpora from the internet, books, etc. In language models, a next word is a natural label for a context sentence as a next word prediction task, or one can artificially mask a known word and then predict it. The paradigm that uses the unstructured data itself to generate labels (for example, the next word or the masked word in language models) and train language models to predict labels thereof is called ‘self-supervision learning’. Language model pre-training is

Table 1. Typical ways for word vectors and language models. $X = \{a, b, c, d, e\}$ is an example text sequence. ELMO, BERT, and GPT usually work on much longer sequences than neural language models (NLMs), Skip-gram and CBOW.

Model	Type	Architecture	Task	Loss function
NLM [25]	static	1-layer MLP	$(a, b) \rightarrow c$ predicting the next word	$-\sum_{i=1}^T \log p(x_i \{x_1, \dots, x_{i-1}\})$
Skip-Gram [200]	static	1-layer MLP	$b \rightarrow c, b \rightarrow a$ predicting neighboring words	$-\sum_{i=1}^T \log p(\{x_{i-o}, \dots, x_{i-1}, x_{i+1}, \dots, x_{i+o}\} x_i), (o \text{ is the window size})$
CBow [200]	static	1-layer MLP	$(a, c) \rightarrow b$ predicting central words	$-\sum_{i=1}^T \log p(x_i \{x_{i-o}, \dots, x_{i-1}, x_{i+1}, \dots, x_{i+o}\}), (o \text{ is the window size})$
Glove [227]	static	1-layer MLP	$\vec{w}_i^T \vec{w}_j \propto \log p(\#(w_i w_j))$ predicting the log co-occurrence count	$-\sum_{i=1, j=1}^T f(x_{ij})(\vec{w}_i^T \vec{w}_j + b_i + c_j - \log x_{ij}), (x_{ij} = p(\#(w_i w_j)))$
ELMO [230]	contextualized LSTM		$(a, b, c, d) \rightarrow e, (e, d, c, b) \rightarrow a$ bi-directional language model	$-\sum_{i=1}^T \log p(x_i \{x_1, \dots, x_{i-1}\}) + \log p(x_i \{x_{i+1}, \dots, x_T\})$
BERT [66], Roberta [185] ALBERT [154], XLNET [350]	contextualized	Transformers or Transformer-XL	$(a, [\text{mask}], c) \rightarrow (b, _)$ predicting masked words	$-\sum_{x \in \text{mask}(x)} \log p(x \hat{X}), \hat{X} \text{ is the corrupted sentence with masks}$
Electra [54]	contextualized	Transformer	$(a, b, c, d) \rightarrow (0, 1, 0, 1)$ replaced token prediction	$-\sum_{i=1}^T \log p(b_i \hat{X}), b_i \text{ indicates whether } x_i \text{ is replaced.}$
T5 [241] BART [158]	contextualized	Transformers	$(a, b, c, _) \rightarrow (d, e)$ predicting the sequence	$-\sum_{i=1}^T \log p(y_i X, y_0, \dots, y_{i-1}), X \text{ and } Y = \{y_1, \dots, y_T\} \text{ are the input/output}$
GPT [240]	contextualized	Transformers	$(a, b, c, d) \rightarrow e$ autoregressively predicting the next word	$-\sum_{i=1}^T \log p(x_i \{x_1, \dots, x_{i-1}\}), \{x_1, \dots, x_T\} \text{ is the sequence}$

therefore referred to as an ‘auxiliary task’, in which the learned representations in language models can be used as an initial model for various downstream supervised tasks. The pre-training objective/task is critical for learning efficient representations that are generalizable and universal for downstream tasks.

Recently, efforts have been proposed to learn contextualized word representations based on deep neural networks, such as the pioneer method ELMO [230], GPT [240], and the breakthrough work: BERT [66]. Similar to traditional neural language models, GPT uses the unidirectional language model task as the pre-training objective. ELMO proposed the pre-training task for bidirectional language modeling based on both the forward language model and backward language model task. The forward language model task aims to model the probability of the word given its previous words, while the backward language model task predicts the word based on its future words. To better model bi-directional contexts during pre-training, BERT proposed the masked language model (MLM) pre-training objective with the inspiration of the Cloze task. It randomly masks tokens of input sequences and aims to predict masked tokens with the masked text sequences. Different from ELMO which concatenates the forward and backward language model, MLM can train the deep bidirectional contextual representations with only one language model. Based on MLM, Encoder-Decoder language models such as T5 [241], proposed the pre-training objective of generating the given sequences in an auto-regressive way taking the masked sequences as input. The language models based on the auto-regressive pre-training objective are more suitable for the text generation tasks such as abstractive summarization and question answering. The overview of pre-training tasks is shown in Table 1. Recently, Open AI have released many API services on their trained model, including GPT 3, InstuctGPT, Codex, and ChatGPT. Especially, ChatGPT could interact in a conversational and makes it possible to answer follow-up questions, admit mistakes, challenge incorrect premises, and reject inappropriate requests.

These pre-training tasks in language modeling are sometimes called ‘pretext tasks’. In conclusion, by pre-training multi-layer transforms in plain text using pretext tasks, it learns general text representation that can easily be adapted to downstream tasks.

Pre-training corpora. Except for the superior pre-training objective, it usually requires a large scale of raw texts to pre-training language models effectively. On the internet, unlabelled raw texts are abundant ranging from news texts,

and web pages, to online encyclopedias. The training corpora for pre-trained language models mainly include: 1) online encyclopedia like Wikipedia⁴, which was widely used for training BERT and its variants. 2) existing books and stories that have been digitized like BooksCorpus [378] and , 3) web texts extracted from online websites/URL, such as crawled online corpora⁵. PLMs trained by these corpora are usually able to capture the common sense knowledge inherited in the raw training texts. For specific domains such as the biomedical domain, it, therefore, needs other efforts such as domain-specific pre-training with domain-specific texts, to capture the domain knowledge (will further be introduced in the next section). Moreover, the vocabulary with limited words is unable to cover all words in the large-scale training texts. To address the out-of-vocabulary (OOV) problem, they proposed to split words into sub-words to formulate the vocabulary via the Byte-Pair Encoding (BPE) [262] or WordPiece [152] methods.

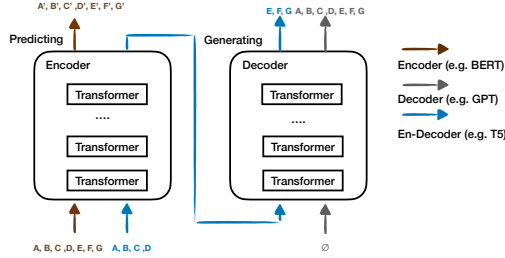


Fig. 4. The difference between Encoder, Decoder and En-Decoder pre-trained language models.

Category	Data	Task
Pre-training	general domain	pre-training task
Domain adaption	target domain	pre-training task
Task adaption	general domain	downstream task
Fine-tuning	target domain	downstream task

Table 2. Categories to tailor pre-trained language models

Representative PLMs. Pre-trained language models can generally be categorized into three principal types, based on whether the input or output constitutes a text sequence or label: Encoder-only, Decoder-only, and Encoder-Decoder models. Models such as BERT [66], RoBERTa [185], and ALBERT [154] fall under the Encoder-only category and are primarily utilized for text classification and sequence labeling tasks. RoBERTa [185] is a BERT variation that has undergone a more extended training phase and employs additional data. ALBERT [154] serves as a lightweight BERT variant but features shared weights and a factorized word embedding.

Pre-trained models equipped with the decoder such as GTP series, T5, BART, could deal with generation-related tasks like translation, summarization, and language models⁶. See Fig. 4 for the difference: an Encoder model predicts labels for each input tokens (in brownish yellow); a Decoder model generates a sequence of tokens w.r.t. a probability distribution (in blue); an En-Decoder model predicts a new sequence conditioned on a given sequence (in grey), *a.k.a.* Seq2Seq.

Knowledge in PLMs. As a pioneer, LAMA [231] has explored the ability about how much PLMs could capture factual and commonsense knowledge (in the format of triplets in knowledge bases). It concludes that large PLMs (e.g., BERT-Large) can recall knowledge slightly better than small competitors and remarkably better than with non-neural and supervised alternatives [231]. However, [39] revise the ability that PLMs can potentially be a reliable knowledge source. Cao et al [39] claims that the way PLMs capture knowledge is vulnerable; it might overfit dataset artifacts and make use of answer leakage. In the biomedical domain, it needs more domain knowledge and it is therefore more knowledge-intensive than the general domain. Some existing work (e.g. [118]) has explored injecting biomedical domain knowledge in PLMs.

⁴<https://dumps.wikimedia.org/>

⁵<https://commoncrawl.org/>

⁶XLNet [350] provides a generalization of autoregressive pre-training by leveraging bidirectional contexts to conduct masked word prediction akin to BERT. It could also deal with text generation.

2.3 Pre-training for images

Deep neural networks have achieved excellent performance in the imaging domain on various vision tasks, e.g., image classification, object detection, and instance segmentation. One of the major reasons behind this is pre-training. However, different from language models in the NLP field, ‘pre-training’ in the earliest means training vision models on large annotated image datasets, e.g., ImageNet [64]. Subsequently, different self-supervised learning approaches are proposed to overcome the shortcoming of supervised learning, e.g., generalization error and spurious correlations. Next, we detail different types of pre-training for images.

Supervised pre-training. In supervised pre-training, the most commonly-used dataset is ImageNet which contains over one million labeled images. Supervised pre-training [100, 150] involves training a deep learning model on the entire ImageNet dataset to learn generic features that can be useful for various downstream tasks. Once the model has been pre-trained on the large dataset, it can be fine-tuned on a smaller, task-specific dataset relevant to the specific task. This can help the model learn valuable features that can be generalized to different tasks at hand.

Contrastive self-supervised Learning. Different from supervised pre-training, contrastive self-supervised learning [48, 89, 99] is a method for representation learning without needing labeled data. It involves training a model to distinguish between different variations of a given input image. For example, the model might be trained to identify whether two images are a rotated version of the same image or whether they are two completely different images. By learning to predict these labels, the model can learn useful features that can be applied to various tasks, such as object detection and semantic segmentation.

Masked self-supervised Learning. Motivated by BERT in NLP, masked self-supervised learning has attracted attention in the computer vision field [21, 98, 336]. It is a type of generative pre-training approach. Models are trained to reconstruct images from incomplete data, in which part of the input image is removed or masked before it is fed into the model. This allows the model to learn the underlying structure of the image.

Contrastive language-image pre-training. Contrastive language-image pre-training [238] (CLIP) aims to train a vision model on a wide variety of image-text datasets. The model is trained to pair images and texts in a mini-batch through contrastive learning. CLIP showed excellent zero-shot transfer ability, where the pre-trained model can achieve comparable results with the original ResNet [100] on ImageNet in a zero-shot manner. One of the primary reasons is that texts provide rich, detailed information about the visual content of an image. For example, a text description of an image can include information about the objects and scenes depicted in the image, as well as their spatial relationships and attributes. This information can help a machine learning model to identify and understand an image’s visual content. Additionally, texts can be easily generated and collected in large quantities, making them a convenient and scalable source of supervision for visual representation learning.

2.4 Fine-tuning Paradigm in PLMs

One challenge to use PLMs in downstream tasks is that there are two gaps between PLMs and downstream tasks, the *task gap* and *domain gap*. The *task gap* means the meta-task in PLMs (usually masked language model in BERT or causal language model in GPT) usually can not directly be tailored to most downstream tasks (e.g. sentimental classification). The *domain gap* refers to the difference between the trained corpora in PLMs and the needed domain in a specific downstream task. The adaptation of both *task gap* and *domain gap* is crucial.

Adaption. To use the pre-trained language model in a downstream task, it is suggested to adopt both the domain and task adaption [90, 94, 253, 366], see Table. 2 for the difference. The domain adaption suggests continuing training pre-trained models trained from a general domain, in the target domain, *e.g.*, biomedical domain. Task adaption refers to fine-tuning on similar downstream tasks. In this paper, without specifying, we mainly discuss the domain-adapted pre-trained models in various downstream tasks. Task adaption is not the main concern in this review. Take BERT as an example, BERT is first trained using next-sentence predictions (NSP) and masked language models in the pre-training phase. Such pre-trained BERT will be used as the initial feature extractor. BERT with an additional classifier layer is then fine-tuned to optimize the objective of down-stream tasks (like MNLI [324], NER [294], and SQuAD [242]).

3 PLMS IN BIOMEDICAL DOMAIN

Recently, the pre-trained language models have been widely applied to various NLP tasks and achieved significant improvement in performance, because: 1) Pre-training on the huge text corpus can learn universal language representations and help with the downstream tasks. 2) Pre-training provides a better model initialization, which usually leads to a better generalization performance and speeds up convergence on the target task. 3) Pre-training can be regarded as a kind of regularization to avoid overfitting on small data [237]. Self-supervised learning, which pre-trained language models rely on, usually adopts plain unstructured corpora in a format of a sequence of tokens. At first, most pre-trained language models focus on pre-training in general plain corpora from the Internet, like Wikipedia or crawled webpages. Except for the general domain, efforts have been proposed to extend PLMs in specific domains such as: [80] trains CodeBERT in the programming language and [23] trains SciBERT on scientific publications and biological sequence. This paper aims to discuss pre-trained language models in the biomedical domain. It is believed that the pre-trained language model can always benefit from more training corpora [90]. To achieve better performance in the domain-specific downstream tasks, it is also intuitive that the in-domain data pre-training is necessary.

We will first introduce the motivation of using pre-trained language models in the biomedical domain in the Sec. 3.1. Then, we will illustrate the main components on tailoring PLMs to the biomedical domain including the in-domain data in the Sec. 3.2, and the pre-training and fine-tuning strategy in the Sec. 3.3. Next, in the Sec. 3.4, we will introduce existing pre-trained models in the biomedical domain, which are pre-trained from the in-domain data as introduced in the Sec. 3.2. We will give an overview of these models, categorization of them, and discussion differences between them. We expect to help one from both the bioinformatics and computer science communities to get knowledge of the biomedical domain-specific pre-trained language model quickly.

3.1 Motivation

In the biomedical domain, the motivation for using pre-trained language models is manifold.

- Firstly, the biomedical domain involves biomedical data in the format of sequential tokens (like biomedical texts and the history of electronic health records) that usually lack annotations. However, these sequential data were previously thought of as difficult to model. Thanks to pre-trained language models, it has been empirically demonstrated to train these sequential data in a self-supervised manner effectively. This would open a new door for processing biomedical data with pre-trained language models.
- Second, annotated data in the biomedical domain is usually limited at scale. Some extreme cases in machine learning are called ‘zero-shot’ or ‘few-shot’. More recently, language models such as GPT3 show that language models have the potential for few-shot learning and even zero-shot learning [38]. Therefore, a well-trained

dataset	types	size	characteristics
MIMIC III	EHR	58,976 hospital admissions for 38,597 patients	from Beth Israel Deaconess Medical Center in 2001-2012
CPRD	EHR	11.3M patients	anonymized medical records from 674 UK GP practices
BREATHES	Scientific Publications	6M articles and about 4 billion words	sources are diverse.
PubMed	Scientific Publications	35M citations and abstracts of biomedical literature	It provide only links to journal articles
COMETA in Reddit	Social Media	800K Reddit posts	68 health-themed subreddits with entity annotation
Tweets	Social Media	up-to-date	one could crawl real-time Tweets using its official API
UMLS	Knowledge Bases	2M names for 900K concepts	well-organized medical knowledge source
IU-Xray	image-Text Pairs	3,955 reports and 7,470 images	XML reports with findings, indications, comparisons, etc.
MIMIC-CXR	image-Text Pairs	77,110 images	images corresponding to 227,835 radiographic studies
ROCO	image-Text Pairs	81,000 radiology images and corresponding captions	figures and their corresponding captions in PubMed articles
MedICaT	image-Text Pairs	17,000 images includes captions	open-access biomedical papers and their captions

Table 3. Summary of Biomedical Data for pre-training.

pre-trained language model in the biomedical domain is more crucial to provide a richer feature extractor, which may slightly reduce the dependence on annotated data.

- Plus, the biomedical domain is more knowledge-intensive than the general domain, since most tasks may need domain expert knowledge, while pre-trained language models could serve as an easily-used soft knowledge base [231] that captures implicit knowledge from large-scale plain biomedical corpora without human annotations. More recently, GPT3 has been shown to have the potential to ‘remember’ many complicated common knowledge [38].
- Lastly, beyond text, there exist various types of biological sequential data in the biomedical domain, like protein and DNA sequences. Using these data to train language models has shown great success in biological tasks like protein structure predictions. Therefore, it is expected that pre-trained language models could solve more challenging problems in biology.

3.2 Biomedical Data for Pre-training

Unstructured plain data for pre-trained language models mainly include electronic health records, scientific publications, social media text, biomedical image-text pairs, and other biological sequences like protein, see Tab. 3. An overview of EHR mining can be seen in [76, 340], and [87] discussed both health records and social media text. One can also check [130] for some systematic overview of biomedical textual corpora.

3.2.1 Electronic Health Record. Electronic health record (EHR) is a collection of patient and population electronically-stored health information in a digital format that may include demographics, medical history, medication and allergies, immunization status, laboratory test results, radiology images, vital signs, personal statistics like age and weight, and billing information. One can check [274, 322] for details about EHR with deep learning. Assessing such records may be restricted to limited organizations, which hinders its widespread to the public. The reason may involve some privacy issues.

MIMIC III. Medical Information Mart for Intensive Care III dataset [126]⁷ is one of the most popular EHR datasets, which consists of 58,976 unique hospital admissions from 38,597 patients in the intensive care unit of the Beth Israel Deaconess Medical Center between 2001 and 2012. In addition, there are 2,083,180 de-identified notes associated with the admissions.

⁷<https://mimic.mit.edu/>

CPRD. Clinical Practice Research Datalink (CPRD) [107] is the primary care database of anonymized medical records from 674 general physicians (GP) practices in the UK, which involves over 11.3 million patients. It consists of data on demographics, symptoms, tests, diagnoses, therapies, and health-related behaviors. It is also linked to secondary care (*i.e.*, hospital episode statistics, or HES) and other health and administrative databases (*e.g.*, office for national statistics' death registration). With 4.4 million active (alive, currently registered) patients meeting quality criteria, approximately 6.9% of the UK population are included, this shows that patients are broadly representative of the UK general population in terms of age, sex, and ethnicity. As a result, CPRD has been widely used across countries and spawned a lot of scientific research output.

3.2.2 Scientific Publications. Scientific publications are another source for biomedical pre-trained language models since we expect that biomedical knowledge may be encapsulated in scientific publications.

BREATHE. Biomedical Research Extensive Archive To Help Everyone (BREATHE)⁸, is a large and diverse dataset collection of biomedical research articles from leading medical archives. It contains titles, abstracts, and full-body texts. The dataset collection process was done with public APIs that were used when available. The primary advantage of the BREATHE dataset is its source diversity. BREATHE is from nine sources including BMJ, arXiv, medRxiv, bioRxiv, CORD-19, Springer Nature, NCBI, JAMA, and BioASQ [42]. BREATHE v1.0 contains more than 6M articles and about 4 billion words. BREATHE v2.0 is the most recent version.

PubMed. PubMed⁹ is a free search engine accessing the MEDLINE database of references and abstracts on life sciences and biomedical topics primarily. PubMed comprises more than 32 million citations for biomedical literature from MEDLINE, life science journals, and online books. Citations may include links to full-text content from PubMed Central and publisher websites. PubMed abstracts (PubMed) have 4.5B words, and PubMed Central full-text articles (PMC) have 13.5B words.

3.2.3 Social Media. Users post information on social media, which may contain biomedical information. We mainly introduce Reddit and Tweets as examples.

Reddit. Reddit is an American social news aggregation, web content rating, and discussion website. Registered members submit content to the site, such as links, text posts, images, and videos, then voted up or down by other members. Posts are organized by subject into user-created boards called "communities" or "subreddits", which cover a variety of topics such as news, politics, religion, science, movies, video games, music, books, sports, fitness, cooking, pets, and image-sharing. Submissions with more up-votes appear towards the top of their subreddit and, if they receive enough up-votes, ultimately on the site's front page. Despite strict rules prohibiting harassment, Reddit's administrators have to moderate the communities and, on occasion, close them. COMETA corpus [22] crawled health-themed forums on Reddit using Pushshift (Baumgartner et al., 2020) and Reddit's own APIs.

Tweets. Twitter is an American micro-blogging and social networking service on which users post and interact with messages known as "tweets". Registered users can post, like, and retweet tweets. Tweets were originally restricted to 140 characters, but the limit was doubled to 280 for non-CJK languages in November 2017. Audio and video tweets remain limited to 140 seconds for most accounts. The COVID-twitter-BERT [210] is trained on a corpus of 160M tweets about the coronavirus collected through the Crowdbreaks platform [211] during the period from January 12 to April 16, 2020.

⁸<https://cloud.google.com/blog/products/ai-machine-learning/google-ai-community-used-cloud-to-help-biomedical-researchers>

⁹<https://pubmed.ncbi.nlm.nih.gov/>

3.2.4 Online Medical Knowledge Sources. Other than unstructured text, there is some online medical knowledge source that is well-organized. For example, UMLS provides biomedical concepts that may benefit biomedical pre-trained language models.

UMLS. Unified Medical Language System (UMLS) [30] (<http://umlsks.nlm.nih.gov>) is a repository of biomedical vocabularies developed by the US National Library of Medicine. The UMLS has over 2 million names for 900,000 concepts from more than 60 families of biomedical vocabularies, as well as 12 million relations among these concepts. These vocabularies include the NCBI taxonomy, the Medical Subject Headings (MeSH), Gene Ontology, OMIM, and the Digital Anatomist Symbolic Knowledge Base. The UMLS knowledge sources are updated every quarter. In addition, all vocabularies are freely available for research purposes within an institution if a license agreement is signed.

3.2.5 Biomedical Image-Text Pairs. Besides texts, there are many medical texts paired with their corresponding images. This type of data is a good resource for learning the cross or joint representations of medical images and texts.

IU-Xray. IU-Xray [62] has a collection of chest X-Ray images from the Indiana University hospital network. The data includes two files: one for the images and the other for the XML reports of the radiography. Each report may have multiple images, typically having two views: frontal and lateral. The XML reports contain information such as findings, indications, comparisons, and impressions. In total, there are 3,955 reports and 7,470 images.

MIMIC-CXR. Medical Information Mart for Intensive Care Chest X-Ray [125] is a large publicly available dataset of chest radiographs with free-text radiology reports. It contains 377,110 images corresponding to 227,835 radiographic studies performed at the Beth Israel Deaconess Medical Center in Boston, MA.

ROCO. Radiology Objects in COntext [225] is a large-scale medical and multimodal imaging dataset from the articles of PubMed Central, an open-access biomedical literature database. They are figures and their corresponding captions in articles. It has over 81,000 radiology images (from various imaging modalities) and their corresponding captions.

MedICaT. MedICaT [282] is also a dataset of medical figure-caption pairs also extracted from PubMed Central. Different from ROCO, 74% of its figures are compound figures, including several sub-figures. It contains more than 217,000 images from 131,000 open-access biomedical papers and includes captions, inline references, and manually annotated sub-figures and sub-captions.

3.2.6 Biological Sequences. Other than text, there are various types of biomedical token sequences, e.g., amino acids for proteins. The structure of each protein is fully determined by a sequence of amino acids [15]. These amino acids are from a limited-size amino acid vocabulary, of which 20 are commonly observed. This is similar to text that is composed of words in a lexicon vocabulary. In this subsection, we introduce a protein dataset called ‘Pfam’ and a DNA sequence dataset from Human Genome Project.

Pfam Protein Dataset. The Pfam database ¹⁰ is a large collection of protein families, in which each protein is represented by multiple sequence alignments using hidden Markov models. The newest version is Pfam 34.0, which was released in March 2021 and contains 19,179 families (or called ‘entries’) and 645 clans ¹¹. The original purpose of the Pfam database is for the classification of protein families and domains. It creates the database using a semi-automated

¹⁰<http://pfam.xfam.org/>

¹¹Clans are the generated higher-level groupings of related entries in Pfam. A clan is a collection of entries that are related by sequence similarity, structure, or profile-HMM.

method of curating information on known protein families. Pfam 34.0 contains 47 million sequences, which could be used to train protein language models.

DNA Dataset. The DNA sequence is composed of a genomic sequence. The Human Genome Project was the international research effort to determine the DNA sequence of the entire human genome. Human Genome Project Results. In 2003, an accurate and complete human genome sequence was finished two years ahead of schedule and at a cost less than the original estimated budget. [119] uses the reference human genome GRCh38.p13 primary assembly from GENCODE Release ¹². The total sequence length is about 3 Billion.

3.3 How to tailor PLMs to the Biomedical Domain

The pre-trained language model [66] is a new two-stage paradigm for NLP. In the first phase, it trains a language model (e.g., masked language model and casual language model) with a self-supervised meta-task in task-agnostic corpora. In the second phase, it fine-tunes the pre-trained language model to a (usually small-scaled) specific downstream task. To tailor pre-trained language models on the biomedical domain, methods [90, 113, 156] have explored conducting the domain-specific adaptation on both the pre-training and fine-tuning stage. In the pre-training stage, the domain-specific adaption of existing efforts involves in the continual pre-training or training from scratch with a large scale of raw biomedical data. This yield many efficient foundation models in the biomedical domain such as BioBERT [156] and PubMedBERT [90] et al, that can be directly used for downstream domain-specific tasks in the fine-tuning stage.

3.3.1 Biomedical Language Model Pre-training. One challenge in the biomedical domain is that medical jargon and abbreviations consist of many terms that are composed of Latin or Greek parts. Moreover, clinical notes have different syntax and grammar from books or encyclopedias. These lead to the semantic and domain-knowledge gap between the general pre-trained language models and the biomedical domain. Therefore, many existing approaches have investigated the biomedical language models pre-training on the basis of pre-trained language models in the general domain, to tailor pre-trained language models to the biomedical domain.

Continual pre-training. The general way used by many methods [113, 156, 226] is to conduct the continual pre-training based on the general pre-trained language models such as BERT. They directly initialize the model with existing general PLMs and further pre-training it with the self-supervised task and domain-specific corpora such as PubMed texts and MIMIC-III et al. The representative works include the BioBERT [156] that conducts continual pre-training based on the BERT with the PubMed abstracts and PubMed Central full-text articles, BlueBERT [226] that uses PubMed texts and MIMIC-III, Clinical BERT [113] that further pre-trains BERT with clinical notes. In this case, they use the same vocabulary as the general PLMs, which cover words in a corpus of the general domain such as Wikipedia and BookCorpus. However, as mentioned before, biomedical texts consist of many domain-specific terms. Using the same vocabulary as the general PLMs can be ineffective for modeling biomedical texts [90].

Pre-training from scratch. To conduct better pre-training for biomedical language models, some efforts [23, 90] have explored the way of pre-training from scratch. Different from the continual pre-training, they propose to build the new vocabulary from the raw biomedical training corpora. SciBERT [23] is the representative work, that constructs the new vocabulary with the size of 30K and trains the model with the mix-domain corpora, where 18% training texts from the computer science domain, and 82% from the biomedical domain. However, one recent work [90] has argued that

¹²https://www.ncbi.nlm.nih.gov/assembly/GCF_000001405.39/

the mixed domain pre-training doesn't make sense for the biomedical domain, since the target data of downstream applications in the biomedical domain is highly domain-specific. Instead, they proposed the superior domain-specific pre-training from scratch that uses the training corpora from only the biomedical domain.

Summary. Our observation is that the core factors that affect the decision between training from scratch or continuously training are twofold: the scale of pre-training biomedical corpora and the domain specificity for biomedicine, where we need to make a trade-off. Pre-training is in general data-hungry, one could fully leverage a large amount of biomedical corpora without inheriting parameters from a well-trained general PLM if there already exist enough biomedical corpora. Early work (e.g., [354]) tends to continuously train biomedical PLMs from an initial BERT. Nowadays, it becomes more popular to directly train biomedical PLMs from scratch thanks to the large scale of collected data and adequate computing resources [187]. Interestingly, [270] reused and tailored a giant general PLM (PaLM) to a clinical one, since giant models are economically expensive. We might expect some approach to decompose existing models and reuse part of them; afterward one can inject biomedical modules into it.

3.3.2 Fine-tuning. Based on well-trained biomedical language models, one has to adapt them to downstream tasks. This is typically implemented to replace the mask language model prediction head and next sentence prediction head with a downstream prediction head, e.g., classification head, or sequence labeling heads.

Since the downstream tasks usually have much less training data than those used in pre-training, fine-tuning is an unstable process. Sun et al [284] investigate different fine-tuning methods of BERT on the natural text classification tasks. Mosbach et al [208] argues that the fine-tuning instability is due to vanishing gradients. Merchant et al [197] observe that fine-tuning mainly modifies the top layers of BERT. Unfortunately, the solutions (e.g. hyper-parameters of which layer to fine-tune) proposed in those papers cannot be easily translated to other settings. To automate this process, automatic hyper-parameter tuning (e.g. Bayesian optimization [37, 298]) can come into help. Tinn et al [292] systematically study fine-tuning stability in biomedical NLP. Particularly, it finds that freezing lower layers is beneficial for small models, while layerwise decay is beneficial for larger models. In most cases, it facilitates robust fine-tuning by using domain specific vocabulary and pre-training.

3.4 Biomedical Pre-trained Language Models

Based on the types of training corpora in the biomedical domain as introduced in the above section 3.2, we mainly introduce two groups of biomedical pre-trained language models: biomedical textual language models and protein language models. Based on the types of training corpora in the biomedical domain as introduced in the section 3.2, we mainly introduce biomedical pre-trained language models in three scenarios: pure language models, vision-and-language modeling, and protein/DNA language models.

3.4.1 Overview of Existing Biomedical Textual Language Models. Since BERT was released, various biomedical pre-trained language models have been proposed via continued training with in-domain corpora based on the BERT model or training from scratch. Tab. 4 presents existing pre-trained language models with used corpora, size, release date, and related web pages.

We introduce some representative pre-trained language models, including encoder-only pre-trained language models like BioBERT, ClinicalBERT, SciBERT, and COVID-twitter-BERT, decoder-only pre-trained language models like MedGPT, and encoder-decoder pre-trained language models like SCIFIVE.

Table 4. Existing textual biomedical pre-trained models. The base setting is with 0.1B parameters, and the large setting is with 0.3B parameters. The date is based on the submission in arXiv or published date of the journal or conference proceeding.

Model	Corpora	Architecture	Size	Date	Link
BioBERT [156]	PubMed and PMC	BERT	base & large	2019.01	https://github.com/dmis-lab/biobert
BERT-MIMIC [269]	MIMIC III	BERT	base and large	2019.02	-
SciBERT [23]	Semantic Scholar papers	BERT	base	2019.03	https://github.com/allenai/SciBERT
BioELMo [121]	PubMed abstracts	ELMo	93.6 M	2019.04	https://github.com/Andy-jqa/bioelmo
Clinical BERT [113]	EHR (MIMIC-III)	BERT	base	2019.04	https://github.com/EmilyAlsentzer/clinicalBERT
Clinical BERT [113]	EHR (MIMIC-III)	BERT	base	2019.05	https://github.com/kexinhuang12345/clinicalBERT
BlueBERT [226]	PubMed+MIMIC-III	BERT	base & large	2019.05	https://github.com/ncbi-nlp/bluebert
G-BERT [263]	MIMIC III	BERT	-	2019.06	https://github.com/jshang123/G-Bert
BEHRT [167]	Clinical Practice Research Datalink	BERT	-	2019.07	https://github.com/deepmedicine/BEHRT
BioFLAIR [264]	PubMed abstracts	BERT	lagre	2019.08	https://github.com/zalandoresearch/flair
RadBERT [195]	RadCore radiology reports	BERT	-	2019.12	-
EhrBERT [161]	MADE corpus	BERT	base	2019.12	https://github.com/umassbento/hrbert
Clinical XLNet [114]	EHR (MIMIC-III)	XLNET	base	2019.12	https://github.com/lindvalllab/clinicalXLNet
CT-BERT [210]	Tweets about the coronavirus	BERT	large	2020.05	https://github.com/digitalepidemiologylab/covid-twitter-bert
Med-BERT [248]	Cerner Health Facts (general EHR)	BERT	-	2020.05	https://github.com/ZhiGroup/Med-BERT
ouBioBERT [304]	PubMed	BERT	base	2020.05	https://github.com/sy-wada/blue_benchmark_with_transformers
Bio-ELECTRA [222]	PubMed	ELECTRA	base	2020.05	https://github.com/SciCrunch/bio_electra
BERT-XLM	Anonymous Institution EHR system	BERT	small and base	2020.06	-
PubMedBERT [90]	PubMed	BERT	base	2020.07	https://huggingface.co/microsoft/BiomedNLP-PubMedBERT-base-uncased-abstract
MCBERT [365]	Chinese social media, wiki and EHR	BERT	base	2020.08	https://github.com/alibaba-research/ChineseBLUE
BioALBERT [215]	PubMed and PMC	ALBERT	base & large	2020.09	https://github.com/usmaann/BioALBERT
BRLTM [196]	private EHR	BERT	customized	2020.09	https://github.com/lanyexiaosa/brlrm
BioMegatron [268]	PubMed and PMC	BERT	0.3/0.8/1.2B	2020.10	https://ngc.nvidia.com/
ClinicalTransformer [348]	MIMIC III	1	base	2020.10	https://github.com/uf-hobi-informatics-lab/ClinicalTransformerNER
Bioreddit-BERT [22]	healththemed forums on Reddit	BERT	base	2020.10	https://github.com/cambridgeltl/cometa
BioRoBERTa [159]	PubMed, PMC, and MIMIC-III	RoBERTa	base & large	2020.11	https://github.com/facebookresearch/bio-lm
CODER [357]	UMLS Metathesaurus	BERT	base	2020.11	https://github.com/GanjinZero/CODER
bert-for-radiology [36]	daily clinical reports	BERT	-	2020.11	https://github.com/tAldiance/bert-for-radiology
BioMedBERT [42]	BREATHE	BERT	large	2020.12	https://github.com/BioMedBERT/biomedbert
LBERT [319]	PubMed	BERT	base	2020.12	https://github.com/warikoon/LBERT
ELECTRAME [203]	PubMed	ELECTRA	base	2021.04	https://github.com/gmpoli/electrained
SCIFIVE [232]	PubMed Abstract and PMC	T5	220/770M	2021.06	https://github.com/justinphan3110/SciFive
MedGPT [146]	King's College Hospital and MIMIC-III	GPY	customized	2021.07	https://pypi.org/project/medgpt/
Clinical-Longformer [169]	MIMIC-III	Longformer [24]	base	2022.01	https://github.com/luoyuanlab/Clinical-Longformer
Clinical-BigBird [358] [169]	MIMIC-III	BigBird	base	2022.01	https://github.com/luoyuanlab/Clinical-Longformer
BioLinkBERT [351]	PubMed with citation links	BERT	base & large	2022.03	https://github.com/michiyaunaga/LinkBERT
BioBART [355]	PubMed	BART	base & large	2022.04	https://github.com/GanjinZero/BioBART
BioGPT [187]	PubMed	GPT	GPT-2 _{medium} ²	2022.09	https://github.com/microsoft/BioGPT
PubMedGPT	PubMed	GPT	2.7B	2022.12	https://www.mosaicml.com/blog/introducing-pubmed-gpt
Flan-PaLM [270]	Instruction ³	PaLM [53]	8B, 62B and 540B	2022.12	unavailable
Med-PaLM 2 [271]	Instruction ⁴	PaLM 2 [16]	8B, 62B and 540B	2023.5	unavailable
HuatuoGPT [361]	Instruction + conversation	GPT (Bloom [258])	7B	2023.5	https://github.com/FreedomIntelligence/HuatuoGPT

¹ ClinicalTransformer [348] provides a series of biomedical models based on different architectures including BERT, RoBERTa, ALBERT, ELECTRA, DistilBERT, XLNet, Longformer, and DeBERTa.² BioGPT adopts GPT-2_{medium} as the backbone network (24 layers, 1024 hidden size and 16 attention heads), resulting 347M 355M parameters in total. Its parameter size is close to BERT-large.³ [270] adopts instruction prompt tuning on medical data. The details were not introduced.⁴ Instructions are from MedQA, MedMCQA, HealthSearchQA, LiveQA and MedicationQA.

- **BioBERT** [156] is initialized with the general BERT model and pre-trained on PubMed abstracts and PMC full-text articles.
- **ClinicalBERT** [113] is trained on clinical text from approximately 2M notes in the MIMIC-III database [126], a publicly available dataset of clinical notes.
- **SciBERT** [23] is trained on the large scale of scientific papers from a multi-domain based on the BERT. The training papers are from 1.14 M full-text papers in Semantic Scholar, in which 82% articles are from the biomedical domain.
- **COVID-twitter-BERT** [210] is a natural language model to analyze COVID-19 content on Twitter. The COVID-twitter-BERT model is trained on a corpus of 160M tweets about the coronavirus collected through the Crowdt-breaks platform during the period from January 12 to April 16, 2020.

- **MedGPT** [146] is a GPT-like language model trained by patients' medical history in the format of electronic health records (EHRs). Given the sequence of past events, MedGPT aims to predict future events like a diagnosis of a new disorder or complications of an existing disorder.
- **SCIFIVE** [232] is a domain-specific T5 model which is pre-trained on large biomedical corpora. Like T5, SCIFIVE is a typical Seq2seq paradigm to transform an input sequence into an output sequence.

3.4.2 *Discussions on Biomedical Pre-trained Language Models.* Here, we will discuss the listed models in various aspects as below:

Training corpora: EHR, literature, social media, etc., or the hybrid? Most pre-trained language models are based on scientific publications *e.g.*, PubMed, and EHR notes. Note that EHR datasets are usually relatively smaller than scientific publications datasets or Wikipedia. Hence pre-trained language models with only EHR datasets are typically trained from the initialization of well-trained BERT [13, 113], XLNET[114], etc. Furthermore, some PLMs (*e.g.*, BioRoBERTa [159]) adopt both scientific publications and EHRs. A few models such as CT-BERT and BioReddit-BERT [22, 210] adopt social media, including Twitter and Reddit.

Extra features. EHR data usually have some extra meaningful features, for example, disease codes, personal information of patients like age, gender. Such extra features can be embedded as dense vectors used in some models such as Med-BERT and BEHRT [167, 248] like word embedding, position embedding, and segment embedding that are used in the embedding layer of Transformer.

Training from scratch or continue training. The standard approach to obtain a biomedical pre-trained model is to conduct continual pre-training from a general-domain pre-trained model like BERT [66], such as the BioBERT [354]. Specifically, this approach would initialize the model with the standard BERT model, including its word vocabulary, which is pre-trained by general Wikipedia and BookCorpus. Besides, some literature demonstrated training from scratch may fully make use of in-domain data and reduce the negative effect from out-of-domain corpora, which may be beneficial for downstream tasks such as PubMedBERT [90].

Reusing existing vocabulary or building a new one. To make use of well-trained general pre-trained language models like BERT [66], one has to reuse its vocabulary [90]. However, Biomedical NLP is more challenging than general NLP because it involves jargon and abbreviations: clinical notes have different syntax and grammar than books or encyclopedias. Moreover, a totally new vocabulary necessarily leads to training from scratch due to different vocabularies that may be more computationally expensive.

Model size. Typically, big models usually have a bigger capacity that needs more data for training. However, the biomedical domain usually does have as many corpora as the general domain. Thus, biomedical pre-trained language models are relatively smaller than general pre-trained language models. Another reason is that most of them are based on BERT or BERT-like encoder-based models, while pre-trained models with decoder architecture (*e.g.*, GPT, T5) could be bigger than encoder-based pre-trained models. To the best of our knowledge, the biggest model is Biomegatron [268] with 1.2B parameters. Note that bigger models take longer for inference, which is unfriendly for those researchers without enough research computing resources.

Being publicly available. Thanks to the open-sourced tradition of computer science, most models have web pages for downloading and documents for usage. Some of them standardized their model in huggingface (<https://huggingface.co>),

Table 5. Existing biomedical vision-and-language pre-trained models. The date is based on the submission in arXiv or published data of the journal or conference proceeding.

Model	Date	Type	Image Encoder	Text Encoder	Fusion Module	Corpora	Downstream Datasets
UMRL [111]	2018.11	Dual-Encoder	DenseNet	GloVe	-	MIMIC-CXR	ICD-9-IT
ConVIRT [370]	2020.10	Dual-Encoder	ResNet	ClinicalBERT	-	MIMIC-CXR, RIH-BONE	CheXpert, COVIDx, MURA, RSNA
MullInfo [173]	2021.05	Dual-Encoder	ResNet	ClinicalBERT	-	MIMIC-CXR	Pathology9, EdemaSeverity
GLORIA [115]	2021.10	Dual-Encoder	ResNet	BioClinicalBERT	-	CheXpert	CheXpert, RSNA, SIIM
LoVT [213]	2021.12	Dual-Encoder	ResNet	ClinicalBERT	-	MIMIC-CXR	COVID-Rural, NIH-CXR, Object CXR, SIIM
BioViL [31]	2022.04	Dual-Encoder	ResNet	CXR-BERT	-	MIMIC-CXR	MS-CXR, RSNA
BFSPR [261]	2022.05	Dual-Encoder	CLIP-Image	CLIP-Text	-	MIMIC-CXR	CheXpert, MIMIC-CXR, NIH-CXR, PadChest
CheXZero [293]	2022.09	Dual-Encoder	CLIP-Image	CLIP-Text	-	MIMIC-CXR	CheXpert, PadChest
MedCLIP [318]	2022.10	Dual-Encoder	ResNet/ViT	BioClinicalBERT	-	CheXpert, MIMIC-CXR	CheXpert, COVID, MIMIC-CXR, RSNA
MGCA [310]	2022.10	Dual-Encoder	ResNet/ViT	BioClinicalBERT	-	MIMIC-CXR	CheXpert, RSNA, SIIM
Analysis [212]	2022.11	Dual-Encoder	ResNet	ClinicalBERT	-	MIMIC-CXR	COVID-Rural, NIH-CXR, Object CXR, SIIM
Analysis [168]	2020.09	Fusion-Encoder	-	-	-	MIMIC-CXR	IU-Xray, MIMIC-CXR
Analysis [312]	2021.03	Fusion-Encoder	ResNet	BERT	Dual-Stream	MIMIC-CXR, NIH14-CXR, IU-Xray	MIMIC-CXR, NIH14-CXR, IU-Xray
Med-ViLL [207]	2021.05	Fusion-Encoder	ResNet	BERT	Single-Stream	MIMIC-CXR	MIMIC-CXR, IU-Xray, VQA-RAD
Berthop [206]	2021.08	Fusion-Encoder	ResNet	BlueBERT	Single-Stream	IU-Xray	IU-Xray
LViT [171]	2022.06	Fusion-Encoder	ViT	BERT	Single-Stream	QaTa-COV19, MoNuSeg	QaTa-COV19, MoNuSeg
M3AE [51]	2022.09	Fusion-Encoder	CLIP-Image	RoBERTa	Dual-Stream	MedICaT, ROCO	VQA-RAD, SLAKE, MedVQA-2019, MELINDA, ROCO
ARL [52]	2022.09	Fusion-Encoder	CLIP-Image	RoBERTa	Dual-Stream	MedICaT, MIMIC-CXR, ROCO	VQA-RAD, SLAKE, MedVQA-2019, MELINDA, ROCO

which will largely be beneficial for its wide-spreading. However, some models are not available to the public due to privacy issues even though data might have been anonymized [157].

Biomedical pre-trained language models in other languages. Most of the biomedical pre-trained language models are in English. However, there is an increasing need for biomedical pre-trained language models in other languages. There are typically two solutions: a multilingual solution or a purely second-language solution. The former may be beneficial for low-resource languages, and the latter is usually used in some rich-resource languages like Chinese [365].

3.5 Beyond Text: Biomedical Vision-and-Language Models

Biomedical data is inherently multi-modal. It includes various types of data: text data, imaging data, tabular data, time-series data, and structured sequence data (e.g., proteins and DNA). Among them, the joint learning of text and imaging data is one of the most explored directions, and biomedical vision-and-language pre-training has emerged as an attractive direction in both artificial intelligence and clinical medicine. This owes to two facts: (i) From the technical perspective, computer vision and natural language processing have been the most popular directions in the past few years, and many models and algorithms have been proposed to process these two types of data; (ii) From the data perspective, the text and imaging data are much easier to obtain in the medical domain, and more importantly they are always pair-collected (e.g., radiology images and their corresponding diagnostic reports).

Most existing biomedical vision-and-language models are motivated by the success of the self-supervised pre-training recipe of SimCLR [48] in CV and BERT in NLP. Most recently, there have also been some studies [43, 44] applying the popular text-to-image diffusion models [243, 252, 256] to the medical domain. In this subsection, we summarize the existing biomedical vision-and-language models in 3.5.1 and describe them in detail.

3.5.1 Overview of Existing Biomedical Vision-and-Language Models. In biomedical vision-and-language pre-training, most existing studies could be categorized into two classes, i.e., dual-encoder and fusion encoder. These two types of models have different advantages and disadvantages. Dual-encoder models are able to capture the relationship between visual and linguistic elements in input by independently encoding each modality and then performing shallow iteration on the resulting vectors. This allows them to effectively learn representations that can be used for single-modal/cross-modal tasks, e.g., image classification, image captioning, and cross-modal retrieval. However, dual-encoder models are

limited in their ability to fully capture the complex interactions between visual and linguistic elements, which can limit their performance on more challenging vision-and-language tasks.

On the other hand, fusion-encoder models aim to overcome this limitation by directly incorporating visual and linguistic elements into a single encoder. This allows them to capture more complex interactions between the two modalities, which can improve their performance on tasks that require a deeper understanding of the relationship between visual and linguistic elements. They jointly process these two modalities with an early interaction to learn multi-modal representations to solve those tasks requiring multi-modal reasoning, e.g., visual question answering. However, it can be more difficult to perform single-modal tasks, as the interactions between visual and linguistic elements are not as easily separated as they are in dual-encoder models. Tab. 5 presents existing dual-encoder and fusion-encoder vision-and-language models.

In addition to dual-encoder and fusion-encoder models, there are other approaches for biomedical vision-and-language pre-training. For example, motivated by the success of diffusion models [243, 252, 256] in the general domain, several medical text-to-image diffusion models [43, 44] have been proposed in the medical domains.

3.5.2 Dual-Encoder Vision-Language Models. Dual-encoder models encode images and texts separately to learn uni-modal/cross-modal representations following a shallow interaction layer (e.g., an image-text contrastive layer). The learned models can be transferred to many single-modal/cross-modal tasks, e.g., image classification and cross-modal retrieval tasks. Next, we detail some representative dual-encoder models:

- **ConVIRT** [370] is the first study to apply contrastive learning to images and texts, inspired by its success in the vision field. For the model architecture, it adopts ResNet and BERT as the vision encoder and the language encoder, respectively. Afterward, a bidirectional contrastive loss between two modalities is used to train these two encoders. It is found that the vision encoder can be used to perform the image classification tasks, requiring much fewer annotated training data as an ImageNet-initialized counterpart to achieve comparable or better performance.
- **GLoRIA** [115] proposed to perform the representation learning of medical images from global and local perspectives. Specifically, for global contrastive learning, it is similar to that of ConVIRT. For local contrastive learning, it uses an attention mechanism to learn local representations by matching the words in radiology reports and image sub-regions.
- **MedCLIP** [318] is trained on both image-text and image-label datasets. The core idea is to pre-compute the matching scores between an image and its text or an image and its label. Subsequently, the scores are used as the target to perform the learning procedure. It is observed that much fewer data are required to learn good representations for zero-shot disease classification.
- **CheXZero** [293] is initialized with the pre-trained CLIP model and pre-trained on the medical image-text dataset. With the strong backbone model and curated designs, CheXZero can achieve comparable results in disease classification tasks in a zero-shot manner.
- **LoVT** [213] is the first dual-encoder study targeting localized medical imaging tasks. It proposed a local contrastive loss to align local representations of sentences or image regions while encouraging spatial smoothness and sensitivity. This promotes its performance on many localized downstream tasks.

3.5.3 Fusion-Encoder Vision-Language Models. Fusion-encoder models encode images and texts and then exploit a fusion module to integrate the image and text features. For the fusion module, normally, there are two types: (i) single-stream: the models use a single Transformer for early and unconstrained fusion between modalities; (ii) dual-stream: the models adopt the co-attention mechanism to interact with different modalities. For fusion-encoder models, the most common objectives are masked language modeling and image-text matching. Similarly, we detail some representative fusion-encoder studies:

- **Li et al.** [168] adopted four general-domain pre-trained vision-and-language models (i.e., LXMERT [288], VisualBERT [165], UNITER [50], and PixelBERT [116]) to learn multi-modal representations from medical images and texts. The experimental results demonstrated their effectiveness of them in disease classification tasks.
- **MedViLL** [207] adopted a single BERT-based model and designed a masking scheme to improve both vision-language understanding tasks (e.g., disease classification, cross-modal retrieval, and visual question answering) and vision-language generation tasks (e.g., radiology report generation).
- **ARL** [52] proposed to integrate medical-domain knowledge bases (e.g., UMLS) into the fusion encoder. Medical knowledge is exploited from three perspectives: (i) aligning through knowledge, (ii) reasoning using knowledge, and (iii) learning from knowledge.
- **LViT** [171] is a vision-and-language fusion-encoder model for medical image segmentation. It leverages medical text annotation to improve the quality of generated segmentation results, especially in the semi-supervised setting.

3.5.4 Other Vision-Language Models. Besides the dual-encoder and fusion-encoder models, there are also some biomedical pre-trained models involving vision and language. We mainly introduce medical text-to-image diffusion models. Diffusion models are a type of generative model inspired by non-equilibrium thermodynamics. By defining a Markov chain of diffusion steps to add random noise to data slowly, the model aims to learn to reverse the diffusion process to construct desired data samples from the noise. Recently, different text-to-image diffusion models (e.g., DALLÉ-2 [243], Stable Diffusion [252], and Imagen [256]) have been proposed and achieved excellent performance on text-based image generation. In the medical domain, RoentGen [43, 44] investigated the adaptation of Stable Diffusion to the medical domain. In specific, they exploited chest X-ray images and their corresponding reports from the MIMIC-CXR dataset to train the model. Then they explored several adaptation approaches (i.e., partially fine-tuning or fully fine-tuning) and different text encoders for adaptation (e.g., domain-agnostic and domain-specific text encoder). The experiments demonstrated the effectiveness of the model with respect to image quality and clinical accuracy.

3.6 Beyond Text: Language Models for Proteins/DNA

Various biological sequences like proteins and DNA could also be treated like linguistic tokens in natural language. Therefore, many existing works explored training language models for these biological sequences. One crucial difference between language models for biological sequences and the counterparts for natural language is tokenization (see Sec. 3.6.1), which leads to different token vocabularies. Sec. 3.6.2 will summarize the existing language models for these biological sequences.

3.6.1 Tokenization for Proteins/DNAs. Like words in the text, biological sequences such as proteins and DNA sequences could also be modeled by language models, which typically aim to predict the next token in a sequence. However,

in contrast to that words are in a relatively big vocabulary (typically 10k-100k), and the vocabularies for biological sequences are usually small.

Tokenization in Proteins. Since the structure of a protein is fully determined by its amino acid sequence [15], one can represent a protein by its amino acid sequences. Roughly 500 amino acids have been identified in nature; however, only 20 amino acids are found to make up the proteins in the human body. The vocabulary of protein sequences consists of these 20 typical amino acids.

Tokenization in DNAs. The two DNA strands are known as polynucleotides, and they are composed of simpler monomeric units (*a.k.a.* nucleotides). Each nucleotide contains one of four nitrogen-containing nucleobases (*i.e.*, cytosine [C], guanine [G], adenine [A], or thymine [T]). The two separate polynucleotides are bound together, according to deterministic base pairing rules ([A] with [T] and [C] with [G]), with hydrogen bonds. Typically, existing work [119] usually adopts a so-called ‘*K*-mer’ representation for DNA sequences¹³ for richer contextual information for DNAs. By doing so, the vocabulary size will increase to the $4^k + 5$ which is exponential to k and additionally pluses five special tokens ([CLS], [SEP], [PAD], [MASK], [UNK]).

3.6.2 Language Models for biological sequences.

Protein language models. Since the commonly-found categories of amino acids are relatively small, namely 20. Initially, some work applied character-level language models to protein to deal with limited-size amino acids. In the beginning, there were many efforts to training RNN-based language models [11, 26] for protein sequences. [102, 103] trains a deep bi-directional model ELMo for proteins¹⁴. Other than those protein sequences, protein language models usually adopt additional features for proteins, *e.g.*, global structural similarity between proteins and pairwise residue contact maps for each protein [26]. Later, [246] introduces the Tasks Assessing Protein Embeddings (TAPE), a suite of biologically relevant semi-supervised learning tasks. The authors also train language models based on LSTM, Transformer, and ResNet on the protein sequences. Bepler et al [27] also proposed a novel framework based on the LSTM model to learn protein sequence embeddings. They make their embeddings publicly available at¹⁵. [250] trains a contextual transformer-based language model¹⁶ on 250 million protein sequences. The representations learned by this LM encode multi-level information spanning from the biochemical properties of amino acids to the remote homology of proteins. Different from the above line of approaches, MSA Transformer [247] fits a model separately to each family of proteins. ProtTrans [77] trains a variety of LM models with thousands of GPUs, and also makes the trained models publicly available¹⁷. ProGen [191] is a generative LM trained on 280M protein sequences conditioned on taxonomic and keyword tags. ProteinLM [330] was recently proposed, which trained a large-scale pre-train model for evolutionary-scale protein sequences, and the trained model is available at¹⁸. More recently, DeepMind developed Alphafold2 [127] that could predict protein structures with high accuracy in the challenging 14th Critical Assessment of protein Structure Prediction (CASP14). Most interestingly, there is an embedded protein language model in Alphafold2, which makes Alphafold2 feasible to make use of unlabelled protein data. In detail, Alphafold2 adopts an auxiliary BERT-like loss to predict pre-masked residues in multiple sequence alignments (MSAs). More recently, ProteinBERT [34] was proposed

¹³‘*K*-mer’ is like a k -size convolutional window for a sequence. For example, a DNA sequence ATGGCT will be tokenized to a sequence of 3-mers {ATG TGC GGC GCT} or to a sequence of 5-mers {ATGGC TGGCT}.

¹⁴<https://github.com/Rostlab/SeqVec>

¹⁵<https://github.com/tbepler/protein-sequence-embedding-iclr2019>

¹⁶The trained model and code are available at <https://github.com/facebookresearch/esm>.

¹⁷<https://github.com/agemagician/ProtTrans>

¹⁸<https://github.com/THUDM/ProteinLM>

Table 6. The performance of different biomedical pre-trained language models on downstream tasks. For all biomedical language models, we compare the F1 score of the base model on various tasks. The BLURB Score calculates the macro average of F1 test results on all tasks.

	BERT	RoBERTa	BioBERT	SciBERT	ClinicalBERT	BlueBERT	PubMedBERT	BioM-ELECTRA	BioLinkBERT	BioGPT
NER										
BC5-chem [162]	89.99	89.43	92.85	92.51	90.80	91.19	93.33	93.1	93.75	-
BC5-disease [162]	79.92	80.65	84.70	84.70	83.04	83.69	85.62	-	86.10	-
NCBI-disease [70]	85.87	86.62	89.13	88.25	86.32	88.04	87.82	88.4	88.18	-
BC2GM [272]	81.23	80.90	83.82	83.36	81.71	81.87	84.52	-	84.90	-
JNLPBA [136]	77.51	77.86	78.55	78.51	78.07	77.71	79.10	-	79.03	-
PICO extraction										
EBM-PICO [220]	71.70	73.02	73.18	73.06	72.06	72.54	73.38	-	73.97	-
Relation extraction										
ChemProt [149]	71.54	72.98	76.14	75.00	72.04	71.46	77.24	77.6	77.57	-
DDI [106]	79.34	79.52	80.88	81.22	78.20	77.78	82.36	-	82.72	-
GAD [35]	79.61	80.63	82.36	81.34	80.48	79.15	83.96	-	84.39	-
Sentence similarity										
BIOSSES [273]	81.40	81.25	89.52	87.15	91.23	85.38	92.30	-	93.25	-
Document classification										
HoC [97]	80.12	79.66	81.54	81.16	80.74	80.48	82.32	-	84.35	85.12
Question answering										
PubMedQA [122]	49.96	52.84	60.24	51.40	49.08	48.44	55.84	-	70.20	78.2
BioASQ [217]	74.44	75.20	84.14	74.22	68.50	68.71	87.56	-	91.43	-
BLURB Score [90]	75.86	76.46	80.34	78.14	77.29	76.27	81.16	-	83.39	-

Table 7. Example for each downstream task.

Task	Input	Output	Example
Named Entity Recognition	Unannotated biomedical text	Annotated text with biomedical entities identified	E.g., identifying drug names, disease terms in text
Relation Extraction	Text with annotated entities	Text with relations between entities identified	E.g., recognizing drug-disease treatment relations
Event Extraction	Text with annotated entities and relations	Text with biomedical events identified	E.g., identifying gene-mutation-event in the literature
Text Classification	Biomedical text	Classified text into pre-defined categories	E.g., classifying medical reports based on disease types
Sentence Similarity	Pair of sentences	Similarity score between the sentences	E.g., measuring similarity between two medical findings
Question Answering	Question and context	Answer to the question based on context	E.g., answering clinical questions based on medical textbooks
Dialogue Systems	User input	System response	E.g., virtual health assistant responding to user health queries
Text Summarization	Long biomedical text	Short summary of the text	E.g., summarizing a medical research article
Natural Language Inference	Pair of sentences	Inference relation between the sentences	E.g., inferring medical conclusions from patient's symptoms

to use a novel task of Gene Ontology (GO) annotation prediction along with masked language modeling and it is also tailored to make the model highly efficient and flexible to very large sequence lengths.

DNA language models. Proteins are translated from DNA through the genetic code. There are 20 natural amino acids that are used to build the proteins that DNA encodes. Therefore, amino acids cannot be one-to-one mapped by only four nucleotides. Some work also explored the potential to build language models on DNA sequences. DNABERT [119] is a bidirectional encoder pre-trained on genomic DNA sequences with up and downstream nucleotide contexts. Yamada et al [342] pre-trains a BERT on RNA sequences and RNA-binding protein sequences. All the LMs remain largely the same as those used for human language data. Designing new architectures and pipelines tailored to protein/DNA sequences is a promising direction.

4 FINE-TUNING PLMS FOR BIOMEDICAL DOWNSTREAM TASKS

Similar to the general domain, to evaluate the effectiveness and facilitate the research development of biomedical pre-trained language models, the Biomedical Language Understanding Evaluation (BLUE) benchmark has been proposed in [226]. BLUE includes five text mining tasks in biomedical natural language processing, including sentence similarity, named entity recognition, relation extraction, text classification, and inference task. However, BLUE does not include some important biomedical application tasks such as question answering, and it mixes the applications of clinical

data and biomedical literature. To improve it, Gu et al [90] proposed a novel benchmark, the Biomedical Language Understanding & Reasoning Benchmark (BLURB). It includes named entity recognition (NER), evidence-based medical information extraction (PICO), relation extraction, sentence similarity, document classification, and question-answering tasks. Moreover, some works proposed the benchmark in other languages, such as Chinese [365].

The development of biomedical pre-trained language models has greatly boosted the performance of these downstream tasks recently. In Table 6, we show the performance when directly fine-tuning different biomedical pre-trained language models for downstream tasks. All biomedical pre-trained language models significantly outperform PLMs in the general domain including BERT and RoBERTa. Especially for sentence similarity and question-answering tasks, the biomedical PLMs such as PubMedBERT and BioLinkBERT outperform BERT and RoBERTa by more than 10% percent. PubMedBERT conducts the domain-specific pre-training from scratch and consistently outperforms other biomedical PLMs such as BioBERT, ClinicalBERT, and BlueBERT in all tasks. Most recently, BioLinkBERT [351] further utilizes the citation links of documents from PubMed abstracts in the pre-training stage, and has achieved the SOTA performance on most tasks. Specifically, for the document level task such as the question-answering task, it outperforms PubMedBERT by 15% percent in the PubMedQA dataset and 4% percent in the BioASQ dataset. In the PubMedQA dataset and another document-level task: document classification, the BioGPT [187] achieves the new SOTA, which conducts the generative pre-training on PubMed abstracts from scratch like GPT.

Besides directly fine-tuning, there is other research exploring how to better leverage and improve PLMs for various downstream tasks. In the following, we will introduce the recent progress based on PLMs on these tasks (we show the example of each downstream task in Table 7) and other critical tasks in the biomedical domain.

4.1 Information Extraction

Information extraction plays a key role in automatically extracting structured biomedical information (entities, concepts, relations, and events) from unstructured biomedical text data ranging from biomedical literature, and electronic health records (EHR) to biomedical-related social media corpus, etc (one can check a review in [316]). In the biomedical community, it generally refers to several important sub-tasks, including named entity recognition (NER), relation extraction, and event extraction.

4.1.1 Named entity recognition. NER aims to identify the common biomedical concept mentions or entity names (such as genes, drug names, adverse effects, metabolites, and diseases) of biomedical texts. Singh et al [255] proposed the first effort to investigate pre-training the bidirectional language model with the PubMed abstract dataset, and then fine-tune the model for the supervised NER task. Compared with traditional neural network based methods such as BiLSTM, it outperforms them by around 1% in the F1-score in datasets including NCBI-disease [70], BC5-disease [162], BC2GM [272], JNLPBA [136], and requires less labelled training data to achieve comparable results. Several methods have shown that further pre-training the language models on the in-domain data can consistently improve the performance. For example, Zhu et al [376]¹⁹ trained a domain-specific ELMo model in the mixture data of clinical reports and relevant Wikipedia pages, which outperforms the previous SOTA method based on BiLSTM-CRF by 3.4% in F1-score in the i2b2 2010 [300] dataset. Si et al [269] have shown that the BERT-large further pre-trained on the MIMIC-III achieves the best performance for the i2b2 2010 dataset, and improves the performance by 5% over that of the traditional neural network method based on the GloVe embedding. Sheikhshab et al [266] have shown that directly using the off-the-shelf ELMo embeddings has limited improvement on the performance, while ELMo continually pre-trained on the in-domain

¹⁹https://github.com/noc-lab/clinical_concept_extraction

data has significant improvement on the performance by 4% in the F1 score of the JNLPBA dataset. Gao et al [82]²⁰ investigated the pre-training and semi-supervised self-training of BiLSTM-CRF and BlueBERT with the in-domain corpora such as MedMentions and SemMed. They evaluated these models on BioNER with limited labeled training data, and the BlueBERT pre-trained on MedMentions has the best performance overall. Moreover, for the scenarios with very few labeled data, the semi-supervised self-training can significantly boost performance.

Some methods have explored how to utilize PLMs for BioNER with less time and computational consumption. Naseem et al [215]²¹ proposed a lightweight domain-specific language model BioALBERT trained on the biomedical domain corpora for biomedical named entity recognition, that captures inter-sentence coherence via the sentence-order-prediction (SOP) task. For eight benchmark datasets, it outperforms the BioBERT by a significant margin, such as increasing the performance of the F1-score by 7.47% in the NCBI-disease dataset and 10.63% in the BC5CDR-disease dataset. Poerner et al [233]²² proposed the time and memory saving domain-adaption method: training Word2Vec on target domain text and aligning them with the word vectors of existing PLMs, and thus propose the GreenBioBERT. On eight BioNER datasets, the GreenBioBERT covers 60% results of BioBERT but only takes 2% of its cloud compute cost. Moreover, there are methods incorporating BioNER with the relation extraction task or modeling BioNER beyond the sequence labeling problem. Khan et al [133] employed PLMs including BERT and BioBERT as the encoder for the multi-task learning of BioNER. They found that using BioBERT has moderately better performance than BERT, and it requires more training epochs for the BERT based method to achieve comparable results. Giorgi et al [86]²³ proposed the end-to-end model for jointly extracting named entities and their relations using PLMs as the encoder. However, in the i2b2 2010 [300] dataset, it has worse performance than the method proposed by Si et al [269] and BlueBERT. Sun et al [285]²⁴ proposed to model the BioNER as the machine reading comprehension (MRC) problem to incorporate the prior knowledge flexibly, and use PLMs as the text encoder. Among ClinicalBERT, BlueBERT and BioBERT, the method based on BioBERT achieves the best performance. Tong et al [295] proposed the auxiliary sentence-level prediction tasks, which can improve the F1 score by 3% in the low-resource scenario on three benchmark datasets compared with BioBERT. Banerjee et al [19]²⁵ formulated the BioNER as the knowledge-guided question-answer task (KGQA), that outperforms the SOTA by 1.78–12% on 11 biomedical NER datasets in the exact match F1 score.

Summary. In Table 8, we summarize the commonly used datasets in the BioNER task and compare performances of different methods on these datasets in Table 9. We can find that the lightweight BioALBERT [215] model pre-trained on the sentence-order-prediction (SOP) task, is the SOTA method on almost all datasets. Among various PLMs, several methods [133, 285, 295] show that using BioBERT generally shows better performance than other PLMs such as BERT, ClinicalBERT, and BlueBERT. Several methods [82, 266, 269, 376] show that pre-training various PLMs such as ELMO, BERT and BlueBERT with various in-domain data such as MIMIC-III and MedMentions, can consistently improve the performance.

4.1.2 Relation Extraction. Biomedical relation extraction (BioRE) aims to identify the relationship (semantic correlation) between biomedical entities mentioned (such as genes, proteins, and diseases) in texts and generally be considered as a

²⁰<https://code.ornl.gov/biomedner/biomedner>

²¹<https://github.com/usmaann/BioALBERT>

²²<https://github.com/npoe/covid-qa>

²³<https://github.com/bowang-lab/joint-ner-and-re>

²⁴<https://github.com/CongSun-dlut/BioBERT-MRC>

²⁵<https://github.com/kuntalkumarpal/KGQA>

Table 8. Datasets used in the BioNER task.

Dataset	Language	Entity type	Text type	Text Genre	Size
BC5-chem [162]	English	Chemical	Abstract	PubMed	1,500
BC4-chem [147]	English	Chemical	Full text	PubMed	10,000
BC5-disease [162]	English	Disease	Abstract	PubMed	1,500
NCBI-disease [70]	English	Disease	Abstract	PubMed	793
i2b2 2010 [300]	English	Disease	Report	Clinical records	871
BC2GM [272]	English	Gene/Protein	Sentence	MEDLINE	20,000
JNLPBA [136]	English	Protein,DNA,RNA,cell line	Abstract	MEDLINE	2,404
LINNAEUS [84]	English	Species	Full text	PMC	100
Species-800 [223]	English	Species	Abstract	MEDLINE	800
EBM PICO [220]	English	Participants,interventions, outcomes	Abstract	PubMed	4,993
CCKS 2017	Chinese	Body,disease,symptom,test,treatment	Report	Clinical Records	400
CCKS 2018	Chinese	Anatomy,symptom,independent,drug,operation	Report	Clinical Records	1,000
PharmaCoNER [5]	Spanish	Protein,chemical	Report	Spanish Clinical Case Corpus	1,000
CANTEMIST [204]	Spanish	Tumor morphology	Report	Spanish Clinical Case Corpus	1,301
CAS [88]	French	Terms,negation,uncertainty	Clinical cases	PubMed	100

Table 9. Performances (F1-score) of different methods on benchmark datasets.

	BC5-chem	BC4-chem	BC5-disease	i2b2 2010	BC2GM	JNLPBA	LINNAEUS	Species-800
Singh et al [255]	-	-	89.28	-	81.69	75.03	-	-
Zhu et al [376]	-	-	-	88.60	-	-	-	-
Si et al [269]	-	-	-	89.55	-	-	-	-
Sheikhshab et al [266]	-	-	-	-	89.72	70.08	-	-
Gao et al [82]	91.80	88.38	84.02	-	80.56	81.44	91.36	72.49
Naseem et al [215]	97.79	96.23	97.61	-	96.33	83.53	99.73	98.72
Poerner et al [233]	93.08	91.26	85.08	-	83.45	76.89	88.34	74.31
Khan et al [133]	90.52	-	-	-	83.01	-	-	-
Giorgi et al [86]	-	-	-	89.26	-	-	-	-
Sun et al [285]	94.11	92.70	87.56	-	85.11	78.45	-	-
Tong et al [295] ²⁶	93.98	-	-	-	84.78	-	-	-
Banerjee et al [19]	90.50	92.39	-	92.67	83.47	79.19	92.63	-

classification problem to predict the possible relation type of two identified entities in a given sentence. Recently, PLMs have been widely explored in the BioRE. Wei et al [320] conducted the first study that investigated fine-tuning BERT and combining additional BIO tag features for the clinical RE. It shows that the BERT-based model outperforms previous SOTA methods based on deep neural networks on n2c2 [105] and i2b2 [300] dataset. Similarly, Thillaisundaram et al [291] adapted the SciBERT to the BioRE via fine-tuning the representation of the classification token (CLS). However, it only compared with and outperformed a simple sampling-based baseline. To further explore the potential of utilizing full information in the last layer to improve performance, Su et al [281] proposed to utilize all outputs of the last layer when fine-tuning the BioBERT model on the BioRE task, which outperforms the BioBERT only using classification token on the DDI [106], PPI [148] and ChemProt [149] dataset. Su et al [280] proposed to employ the contrastive learning to improve fine-tuning BERT model for biomedical relation extraction, which outperforms directly fine-tuning BERT on the DDI, PPI and ChemProt dataset. Xue et al [339] proposed to fine-tune BERT for joint entity and relation extraction in Chinese medical text, which outperforms the SOTA joint model based on Bi-LSTM by 1.6%. Chen et al [49] combined BERT with the one-dimensional convolutional neural network (1D-CNN) for the medical relation extraction, which significantly outperforms the traditional 1D-CNN classifier. Lin et al [176, 177] combined the global embeddings and multi-task learning to improve BERT on the clinical temporal relation extraction. Guan et al [91] investigated several PLMs including BERT, RoBERTa, ALBERT, XLNet, BioBERT, ClinicalBERT, in predicting the relationships between clinical events and temporal expressions, and found that RoBERTa generally has the best performance. To prevent private information leakage, Sui et al [283] proposed the first privacy-preserving medical relation extraction method FedED based on BERT and federated learning, which achieved promising results on three benchmark datasets.

Table 10. Datasets used in the BioRE task.

Dataset	Entity type	Text type	Relation Size
i2b2 2010 [300]	Medical problem–treatment	Report	5,261
i2b2 2012 [287]	Event–temporal expression	Summary	8,294
TM [278]	Event–event	EHRs	355
DDI [106]	Drug–drug	Abstract	48,223
PPI [148]	Protein–protein	Abstract	5,834
ChemProt [149]	Protein–chemical	Abstract	31,784
BioC VI PM [69]	Protein–protein	Full text	1,629

Table 11. Performances (F1-score) of different methods on benchmark datasets.

	i2b2 2010	DDI	PPI	ChemProt	ib2b 2012
Wei et al [320]	76.79	-	-	-	-
Su et al [281]	-	80.7	82.5	76.8	-
Su et al [280] ²⁷	-	82.9	82.7	78.7	-
Chen et al [49] ²⁸	-	-	-	-	70.85
Guan et al [91]	-	-	-	-	70.5
Sui et al [283]	-	-	-	-	75.09

Summary. In Table 10 and Table 11, we summary the commonly used datasets and compare the performances of different methods on these datasets. In summary, fine-tuning various PLMs significantly outperforms traditional neural network based methods [49, 320], and improving the fine-tuning strategy can further improve the performance, for example Su et al [280] using the contrastive learning as the auxiliary task, achieves the best performance on DDI, PPI and ChemProt.

4.1.3 Event Extraction. Event extraction is another important task for mining structured knowledge from biomedical data, which aims to extract interactions between biological components (such as protein, gene, metabolic, drug, disease) and the consequences or effects of these interactions [14]. Similar to BioRE, it is formulated into the multi-classification problem. Many efforts have been proposed to explore the application of PLMs in biomedical event extraction recently. Trieu et al [296]²⁹ proposed the model called DeepEventMine with the BERT-based encoder, which significantly outperforms the strong baseline based on CNN. Wadden et al [305]³⁰ explored combining the BERT model and graph propagation to capture long-range cross-sentence relationships, which have been proven to improve the performance of the model-based BERT alone. Ramponi et al [245]³¹ modeled the biomedical event extraction as the sequence labeling problem, and proposed the model called BEESL with the BERT model as the encoder. It outperformed the baseline based on LSTM by 1.57% in the GENIA 2011 [137] benchmark. Wang et al [313]³² formulated the biomedical event extraction as the multi-turn question-answering problem and utilized the question-answering system based on the SciBERT. The method can form event structures from the answers to multiple questions and achieves promising results on GENIA 2011 [137] and Pathway Curation 2013 [235] dataset. In Table 12 and Table 13, we summarize commonly used datasets and compare the performance of different methods.

²⁹<https://github.com/aistairc/DeepEventMine>

³⁰<https://github.com/dwadden/dygiepp>

³¹<https://github.com/cosbi-research/beesl>

³²https://github.com/WangXII/bio_event_qa

Table 12. Datasets used in the Biomedical event extraction.

Dataset	Entities	Triggers	Relations	Events
Cancer Genetics 2013 [216]	21,683	9,790	13,613	17,248
EPI 2011 [221]	16,675	2,035	3,416	2,453
GENIA 2011 [137]	22,673	10,210	14,840	13,560
GENIA 2013 [138]	12,725	4,676	7,045	6,016
Infectious Diseases 2011 [236]	12,788	2,155	2,621	2,779
Pathway Curation 2013 [235]	15,901	6,220	10,456	8,121
Multi-level event extraction [234]	8,291	5,554	7,588	6,677

Table 13. Performances (F1-score) of different methods on benchmark datasets. Genia means the GENIA 2011 [137] dataset. PC means the Pathway Curation 2013 [235] dataset.

	Genia	PC
Trieu et al [296]	63.96	55.67
Ramponi et al [245]	60.22	-
Wang et al [313]	58.33	48.29

4.2 Text Classification

Text classification aims to classify biomedical texts into pre-defined categories, which play an important role in the statistical analysis, data management, retrieval of biomedical data et al. Fine-tuning pre-trained language models on biomedical text classification has attracted great attention recently. Gao et al [81] investigated four methods of adapting the BERT model to handle input sequences up to approximately 400 words long, for the clinical single-label and multi-label clinical document classification. However, they found that the BERT or BioBERT model generally has equal or worse performance for clinical data such as the MIMIC-III clinical notes dataset, than a simple CNN model. They suggested that it may be because BERT or BioBERT models don't capture clinical domain knowledge due to trained on the general domain or biomedical literature datasets, and can't handle too long sentences longer than 512 tokens. Mascio et al [193] made a comprehensive analysis of the performance of various word representation methods (such as Bag-of-Words, Word2Vec, GloVe, FastText, BERT, BioBERT) and classification approaches (Bi-LSTM, RNN, CNN) on the electronic health records classification. They found that the contextual embeddings from BERT and BioBERT generally outperform the traditional embeddings, and the traditional deep neural networks Bi-LSTM enriched with appropriate entity information and specific domain embeddings have better performance than BERT and BioBERT. Guo et al [92] compared the performance of three PLMs including RoBERTa-base, BERTweet, and Clinical BioBERT on 25 social media classification datasets, in which 6 datasets are biomedical related. They found that RoBERTa-base and BERTweet outperform Clinical BioBERT, in which RoBERTa-base can capture general text semantic characteristics, while BERTweet captures more domain knowledge. Gutierrez et al [95]³³ also provided an analysis of traditional deep neural networks and fine-tuning PLMs including BERT and BioBERT on the performance of multi-label document classification on the COVID-19 dataset: LitCovid. They found that BERT and BioBERT models have better performance than traditional methods such as RNN, CNN, and Bi-LSTM in the datasets, and BioBERT outperforms BERT due to domain-specific pre-training.

Summary. We summarize commonly used biomedical text classification datasets in Table 14, and show the performances of different methods in these datasets in Table 15. In summary, all these methods found that directly using fine-tuning PLMs outperforms the traditional neural network based methods. The performance of different language models depends on the target datasets, for example, BERTweet pre-trained with large scale English Tweets, significantly outperforms ClinicalBioBERT on the social media dataset PM Abuse [92]. BioBERT has promising performance on the covid-19 dataset but has worse performance than ClinicalBioBERT on the clinical data such as the MIMIC-III.

³³<https://github.com/dki-lab/covid19-classification>

Table 14. Datasets used in the biomedical text classification task.

Dataset	Label type	Label Num	Avg Label Num	Text type	Data Size
HoC [97]	Multi-label	37	-	PubMed abstracts	1,852
MeSH [297]	Multi-label	26,563	12.55	Biomedical articles	10,876,004
MIMIC-III [126]	Multi-label	6,919	11.7	Discharge summaries	49,785
LitCovid [95]	Multi-label	8	-	PubMed articles	23,038
CORD-19 Test [95]	Multi-label	8	-	PubMed articles	100
PM Abuse [9]	Multi-label	4	-	Tweets	15,100

Table 15. Performances (accuracy for PM Abuse, macro F1 score for other datasets) of different methods on classification datasets.

Dataset	HoC	MiMIC-III	PM Abuse	LitCovid	CORD-19
LSTM [4]	-	-	-	83.9	83.2
CNN [140]	-	-	-	83.3	82.7
BERT [65]	80.12	29.3	-	85.5	85.1
BioBERT [156]	81.54	32.4	-	86.3	86.2
ClinicalBioBERT [12]	-	44.4	77.4	-	-
BERTweet [218]	-	-	82.4	-	-

Table 16. Benchmark datasets in the biomedical sentence similarity task.

Dataset	Text type	Data Size
BIOSSES [273]	PubMed	100
MedSTS [314]	Clinical report	174,629
MedSTS_ann [314]	Clinical report	1,068
n2c2/OHNL [315]	Clinical report	1,642

Table 17. Performances (F1-score) of different methods on benchmark datasets.

	BIOSSES	MedSTS
Chen et al [47]	84.8	83.6
Chen et al [45]	-	83.8
Chen et al [46]	-	85.28

4.3 Sentence Similarity

The semantic similarity task is generally formulated into the regression problem to predict the similarity score of each sentence pair. Recent works have focused on fine-tuning various PLMs for this task. Chen et al [47]³⁴ proposed the first pre-trained open set sentence embeddings in the biomedical domain called BioSentVec, which is trained on over 30 million documents from both biomedical literature such as PubMed and clinical notes such as the MIMIC-III Clinical Database. Compared with existing word embeddings and sentence encoder-based methods, it yields better performance on both sentence similarity and text classification tasks, due to better capturing the sentence-level semantic information. Chen et al [45] empirically compared the performance of traditional deep learning methods such as random forest, RNN, CNN with PLMs including BERT and BioBERT, which shows that PLMs are more effective. Chen et al [46] further show the BioSentVec can improve the performance of traditional deep learning models by 2% F1 score. Yang et al [349] explored three PLMs including BERT, XLNet, and RoBERTa for the clinical semantic textual similarity task, in which the XLNet achieves the best performance among the three models. We show commonly used sentence similarity datasets and compare performances of different methods in Table 16 and Table 17. We found that Chen et al [46] using the

³⁴<https://github.com/ncbi-nlp/BioSentVec>

pre-trained sentence embedding BioSentVec with the traditional neural networks has better performance than Chen et al [45] directly fine-tuning BERT and BioBERT.

4.4 Question Answering

Biomedical question answering (BioQA) aims to extract or generate the natural language answers to the given questions, and generally be formulated into the machine reading comprehension approach focusing on predicting the text span of answers with the given questions and passages containing the answers. Recently, the fine-tuning and transfer learning of PLMs have been widely explored in the task. Yoon et al [353]³⁵ applied the BioBERT to answer biomedical questions such as factoid, list, and yes-no type questions. They show that BioBERT fine-tuned with the question-answering datasets in both the general and biomedical domains and achieved the best performance in the 7th BioASQ Challenge. Jeong et al [117]³⁶ proposed to transfer the knowledge of natural language inference (NLI) to BioQA with BioBERT, which outperforms previous methods on Yes/No, Factoid, and List type questions by 5.59%, 0.53%, and 13.58%, in the 7th BioASQ Challenge. Chakraborty et al [42]³⁷ proposed a novel language model BioMedBERT for question answering (QA) and information retrieval tasks, which is pre-trained on a large-scale biomedical literature dataset BREATHE based on BERT, and outperforms BERT in the BioQA. Kamath et al [131] compared the effectiveness of PLMs based on two different QA models including the machine-reading comprehension and open question-answering method, and show the question-answering model achieves better performance on the BioQA. Du et al [75] utilized the BERT model as the encoder and then used the scaled dot-product attention mechanism to capture the interaction between the question and passage. The proposed method outperforms the best performance for factoid questions in 2016 and 2017 BioASQ-Task B. Zhou et al [375] utilized the BioBERT and interactive transformer model for both the recognizing question entailment and question answering task, and showed significant improvements on the single task with the shared representations of both tasks. Similarly, Akdemir et al [6] also explored multi-task learning to improve the performance of BioBERT on the BioQA task with the biomedical entity recognition task, and show its improvements on the BioASQ 8B challenge. However, these models can't detect multiple spans of the passage when there are multiple answers to the question. To solve the problem, Yoon et al [352]³⁸ reformulated the BioQA task as the sequence tagging problem to detect multiple entity spans simultaneously based on the BioBERT encoder, which achieves the BioASQ 7b and 8b list-type questions.

Some works tried to incorporate domain knowledge, such as biomedical-named entities, into PLMs. He et al [101]³⁹ proposed to infuse the domain knowledge of disease into a series of PLMs including BERT, BioBERT, SciBERT, ClinicalBERT, BlueBERT, and ALBERT, to improve their performance. They found all these models can be improved by infusing the disease knowledge, and for example, the accuracy of BioBERT on the CHQ dataset can be improved by nearly 4%. Rawat et al [249]⁴⁰ incorporated the medical entity information with entity embeddings and the auxiliary task on predicting the logical form of the question to improve the accuracy and generalization of the BERT model on answering questions, which improves the BERT model by 5% F1 score on the paraphrased question answering of the emrQA dataset. Kommaraju et al [144] introduced the extra biomedical named entities prediction task to improve the BioBERT on Biomedical QA. They show the BioBERT pre-trained by the prediction task outperforms the previous best model on the 7b-Phase B of the 7th BioASQ Task challenge.

³⁵<https://github.com/dmis-lab/bioasq-biobert>

³⁶<https://github.com/dmis-lab/bioasq8b>

³⁷<https://github.com/BioMedBERT/biomedbert>

³⁸<https://github.com/dmis-lab/SeqTagQA>

³⁹<https://github.com/heyunh2015/diseaseBERT>

⁴⁰https://github.com/emrQA/bionlp_acl20

Table 18. Datasets used in the biomedical question answering task.

Dataset	Text type	Data Size
PubMedQA [122]	PubMed abstracts	1,000
BioASQ [217]	MEDLINE articles	885
MEDIQA [3]	online community	383
emrQA [224]	Clinical notes	400,000
cMedQA [367]	online community	61,343
COVID19-QA [289]	Literature review	124

Table 19. Performances of different methods. For BioASQ 6b, we compare the Mean Reciprocal Rank (MRR) score on the Factoid question. For the BioASQ, we compare the averaged MRR score on the Factoid question of all batches. For the MEDIQA and emrQA, we compare the accuracy score.

	BioASQ 6b	BioASQ 8b	MEDIQA	emrQA
Yoon et al [353]	48.41	-	-	-
Jeong et al [117]	48.05	46.65	-	-
Chakraborty et al [42]	50.50	-	-	-
Kamath et al [131]	45.70	-	-	-
Zhou et al [375]	-	-	75.8	-
Akdemir et al [6]	-	43.61	-	-
Yoon et al [352]	-	37.95	-	-
He et al [101]	-	-	79.49	-
Rawat et al [249]	-	-	-	59.00
Kommaraju et al [144]	-	43.93	-	-
Soni et al [275]	-	-	-	75.56
Soni et al [275]	-	-	-	86.97

Besides methods for biomedical literature corpora, other works have proposed question-answering models for unstructured electronic health records (EHR). Soni et al [275] investigated the performance of various PLMs including BERT, BioBERT, ClinicalBERT, and XLNet on the clinical question answering, and explored the fine-tuning methods with different datasets, including datasets in the general domain, biomedical and clinical corpora. They find that fine-tuning the open-domain dataset SQuAD consistently improves the performance across all the model variants. Mairitha et al [192] explored four different fine-tuned BERT models for personalized EHR question answering and show the extended BioBERT-QA model pre-trained on unstructured EHR data achieves the best performance. Table 18 shows commonly used datasets in BioQA, and Table 19 presents the performances of different methods on these datasets. Several methods [192, 275, 353] have shown fine-tuning PLMs with the open domain question-answering dataset, and pre-training PLMs with the in-domain datasets improves the performance of various PLMs. Generally, it assumes that fine-tuning or pre-training with more corpora is always useful, for example, Soni et al [275] fine-tuning PLMs with the general domain, biomedical and clinical corpora achieves the best performance for the clinical question answering. Moreover, incorporating the domain knowledge including disease knowledge, medical entities, and multi-task learning incorporating BioNER task, can significantly improve the performance. For example, He et al [101] improved the accuracy of BioBERT on the MEDIQA dataset by nearly 4% with the disease knowledge and achieved the best performance.

4.5 Dialogue Systems

The dialogue system aims to produce a proper response in either a selective [323, 371] or generative [184, 359, 373] way given a dialogue context for the biomedical goals of a user. The context includes historical utterances from users and systems, biomedical knowledge base, electronic health records of users, etc. The format of a response could be various,

e.g., a set of structured user goal data [321], a distribution of biomedical labels for diagnosis [180, 371] and natural language utterances [359]. For different types of contexts and responses, recent work focuses on end-to-end Dialogue System (DS) [337, 359] or parts of four typical DS modules, *i.e.*, Natural Language Understanding (NLU) [73, 267], Dialogue State Tracking (DST) [180, 321], Dialogue Policy Learning (DPL) [321, 329] and Natural Language Generation (NLG) [359]. Recently, PLMs are well-known for natural language modeling, but it is nontrivial to pre-train on task datasets that are based on a specific domain [323]. To adapt PLMs to the medical domain, the dominant solution is to pre-train a language model on a large-scale general/medical corpus and then fine-tune the model with a medical dialogue dataset. Yan et al [344]⁴¹ first explored fine-tuning PLMs including BER-WWM, BERT-MED, MT5 and GPT2 on M^2 -MedDialog dataset for understanding the intents and slots of patients, in which MT5 achieves the best performance. Zeng et al [359]⁴² pre-trained Transformer, BERT-GPT, and GPT on dialog datasets and other large-scale texts, and then fine-tune models on the Chinese MedDialog dataset for generating clinically correct and human-like medical responses. BERT-GPT has been shown to have lower perplexity compared to both Transformer and GPT, while maintaining similar diversity metrics as Transformer. Shi et al [267]⁴³ show BERT has promising performance on the medical slot-filling task, and pre-trained embedding from BERT can further improve the performance of the weak supervision method. DialoGPT [373]⁴⁴ is pre-trained based on GPT-2 [240] with a large in-domain dialogue dataset, and is able to generate more relevant, informative and coherent responses compared with the strong baseline based on the sequence to sequence model. Li et al [164]⁴⁵ proposed the dialogue-adaptive pre-training objectives (DAPO) by considering dialogue-specific features including coherence, specificity, and informativeness, which shows better performance than other language modeling objectives such as MLM and NSP.

Summary. We summarize all available biomedical dialogue datasets in Table 20. Different from using the accuracy, recall, and F1 score metrics used by previous tasks, the dialogue system task generally uses the machine translation metrics including BLEU [373], METEOR [20], and NIST [68], to measure the similarity between generated responses and the ground truth based on n-gram matching. These metrics for evaluating generated responses are limited in that they only take into account shallow lexical overlaps and do not account for paraphrasing and terminology variations. To address this, some automatic metrics based on pre-trained language models have been developed, such as BERTScore [368], which calculates the similarity between two sentences using contextual embeddings from PLMs. However, these metrics have been shown to be inadequate in evaluating the faithfulness of generated responses. While there have been efforts to develop factual consistency metrics like BARTScore [356] in the general domain, there has been less focus on developing such metrics in the biomedical domain to evaluate factual correctness. Since the aforementioned methods utilized different datasets, it is hard to compare their performances directly. In summary, they have demonstrated that creating more effective pre-training tasks, incorporating task-specific information, and pre-training with large in-domain dialogue datasets are effective strategies for improving the performance of series PLMs.

⁴¹<https://github.com/yanguojun123/Medical-Dialogue>

⁴²<https://github.com/UCSD-AI4H/Medical-Dialogue-System>

⁴³<https://github.com/xmshi-trio/MSL>

⁴⁴<https://github.com/microsoft/DialoGPT>

⁴⁵<https://github.com/lockon-n/dapo>

Table 20. Datasets used in the biomedical dialogue system tasks.

Dataset	Language	Domain	Evaluated Task	Text type	# dialogues
MZ [321]	EN	Pediatrics	DPL	Discharge summaries	710
DX [337]	CN	Pediatrics	DPL	Patient-doctor dialogues & patient reports	527
RD [172]	CN	Pediatrics	DPL	Patient-doctor dialogues & patient reports	1,490
SD [172]	CN	9 domains	DPL	Patient-doctor dialogues & patient reports	30,000
CMDD [180]	CN	Pediatrics	NLU	Patient-doctor dialogues	2,067
SAT [72]	CN	14 domains	NLU	Patient-doctor dialogues	2,950
MSL [267]	CN	Pediatrics	NLU	Patient-doctor dialogues	1,652
MIE [371]	CN	Cardiology	NLU	Patient-doctor dialogues	1,120
CovidDialog [347]	CN/EN	COVID-19	NLG	Patient-doctor dialogues	1,088/603
MedDG [184]	CN	Gastroenterology	NLG	Patient-doctor dialogues	17,000
MedDialog [359]	CN/EN	29 domains	NLG	Patient-doctor dialogues & patient reports	3,407,494/257,332
Chunyu [179]	CN	-	NLG	Patient-doctor dialogues	12,842
KaMed [160]	CN	12 domains	NLG	Patient-doctor dialogues	63,754
M^2 -MedDialog-base [344]	CN	30 domains	NLU&DPL&NLG	Patient-doctor dialogues & patient reports	1,557
M^2 -MedDialog-large [344]	CN	40 domains	NLG	Patient-doctor dialogues & patient reports	95,408

4.6 Text Summarization

Automatic text summarization aims to automatically summarize the key information of single or multiple documents with shorter and more fluent texts, which greatly decreases the time-consuming of acquiring important information. Similar to the general domain, existing methods can generally be classified into two categories: extractive summarization methods and abstractive summarization methods.

To explore the advanced PLMs in the text summarization of the biomedical domain, the domain knowledge is incorporated by existing methods via domain fine-tuning [190, 332]. For biomedical extractive summarization, Du et al [74] proposed a novel model BioBERTSum, which used the domain-aware pre-trained language model as the encoder and then fine-tuned it on the biomedical extractive summarization task. It outperforms SOTA extractive methods such as BERTSum. Xie et al [331]⁴⁶ proposed the knowledge infusion training framework to incorporate medical knowledge to improve a series of PLMs including BERT, RoBERTa, BioBERT, and PubMedBERT. The PubMedBERT-based method has the best performance and outperforms other strong baselines such as BERTSum and MatchSum. Gharebagh et al [85] utilized the domain knowledge: salient medical ontological terms to help the content selection of the SciBERT-based clinical abstractive summarization model, which improves SOTA results by around 2% in ROUGE metrics on two medical datasets MIMIC-CXR [124] and OpenI [63]. Bishop et al [28]⁴⁷ proposed unsupervised extractive summarization method for biomedical literature with T5 and BERTScore, which achieves better performance than strong supervised methods such as BERTSum. Xie et al [333]⁴⁸ incorporated the neural topic model with hierarchical transformer encoder (HTE) based on PLMs, which significantly improved the performance of RoBERTa on long biomedical document summarization.

For abstractive summarization, Wallace et al [306] utilized the Bidirectional and Auto-Regressive Transformers (BART) as the encoder for generating biomedical evidence summary of multiple clinical trials. They found that the summarizers can produce fluent and relevant synopses, but the factual accuracy can't be guarantee. Deyoug et al [67]⁴⁹ investigated the BART model for the multi-document summarization on medical studies, which can generate coherent summaries that align with the reference summaries in evidence direction approximately 50% of the time. Guo et al [93]⁵⁰ proposed a novel task of plain language summarization task on the biomedical scientific reviews, and explored

⁴⁶<https://github.com/xashely/KeBioSum>

⁴⁷<https://github.com/jbshp/gencomparesum>

⁴⁸https://github.com/xashely/GRETEL_extractive

⁴⁹<https://github.com/allenai/ms2/>

⁵⁰https://github.com/qiuweipku/Plain_language_summarization

Table 21. Datasets used in the biomedical text summarization.

Dataset	Text type	Type	Data Size
COVID-19 [311]	Biomedical literature	Single	-
MS^2 [67]	Biomedical literature	Multi	470,402
CDSR [93]	Biomedical literature	Single	7,805
RCT [306]	Clinical trials	Multi	4,528
PubMed [55]	Biomedical literature	Single	119,924
MIMIC-CXR [124]	Radiology reports	Single	124,577
OpenI [63]	Radiology reports	Single	3,599
Readability [188]	Biomedical literature	Single	28,124

pre-training BART model on general domain dataset CNN/DM and in-domain PubMed dataset. They found BART pre-trained using CNN/DM and PubMed abstracts demonstrate the strongest ROUGE scores, whereas the BART model pre-trained only using PubMed abstracts has the lowest level of readability. Luo et al [188]⁵¹ proposed the new task of readability controllable summarization for biomedical documents, and explored the language model Longformer-Encoder-Decoder (LED) with the advanced controllable techniques including prompts and multi-head. They demonstrate that the method can generate fluent summaries, but it lacks the capability to effectively control for readability. Hu et al [112] incorporated the additional knowledge with graph encoder and contrastive learning, to enhance the performance of the BioBERT. The proposed method achieves state-of-the-art results in radiology report summarization. For the information acquisition of COVID-19 related scientific literature, Kieuvongngam et al [135] proposed the BERT and GPT-2 based model for both extractive and abstractive summarization of COVID-19 research literature. There are also works to build the multi-document summarization system for the information retrieval of COVID-19 research literature with the Siamese-BERT [78], BioBERT, and XLNet [279].

Similar to the dialogue system task, the commonly used automatic metrics in the text summarization task including ROUGE [178], and BERTScore [368], usually evaluate the relevance and similarity between the generated summaries and the gold summaries. Moreover, the factuality metrics have attracted much attention recently to evaluate the factual correctness of generated summaries [189, 332, 335]. Deyoung et al [67] introduce the ΔEI metric to determine the degree of the factual accuracy of generated summaries in relation to the input medical studies. Zhang et al [372] introduced the ChexBERT F1 score to evaluate the factual correctness between generated summaries and gold summaries of radiology reports. In Table 21, we summarize datasets used in the biomedical text summarization task, and report performances of different methods on these datasets with the aforementioned evaluation metrics in Table 22. We can find that the method incorporating the domain knowledge [331] has better performance than directly fine-tuning PLMs [74], and the method [333] for long biomedical text summarization achieves the best performance on two biomedical literature datasets PubMed and CORD-19.

4.7 Natural Language Inference

Natural language inference (NLI, also known as text entailment) is a basic task for the natural language understanding of biomedical texts. It aims to infer the relation such as entailment, neutral and contradiction, between two sentences, named as the premise and hypothesis, which can further benefit biomedical downstream tasks such as commonsense comprehension, question answering and evidence inference.

⁵¹<http://www.nactem.ac.uk/readability/>

Table 22. Performances (ROUGE-L score, counting the longest common subsequence (LCS) between the generated summary and the reference summary.) of different methods.

	PubMed	CORD-19	MS [^] 2	RCT	MIMIC-CXR
Du et al [72]	29.58	-	-	-	-
Xie et al [331]	33.28	29.10	-	-	-
Bishop et al [28]	35.65	33.35	-	-	-
Xie et al [333]	38.61	40.01	-	-	-
Wallace et al [306]	-	-	-	0.265	-
Deyoug et al [67]	-	-	20.80	0.1760	-
Hu et al [112]	-	-	-	-	46.65

Table 23. Datasets used in the biomedical natural language inference.

Dataset	Text type	Relation Type	Data Size
MedNLI [251]	Clinical notes	Entailment, contradiction, or neutral	14,049
RQE [1]	Consumer health questions	Entailment, contradiction	9,120
CMFAQ [377]	Consumer health questions	Entailment, contradiction	53,822

To facilitate the development of methods for text inference and entailment in the medical domain, participants in the MEDIQA 2019 shared task [3] investigated the SciBERT, BioBERT, and ClinicalBERT in the medical NLI task. Among these participants, Wu et al [328]⁵² achieves the best performance with 98% accuracy in the REQ dataset [1], which ensembled results of different base models and incorporated the syntax information. Sharma et al [265]⁵³ incorporated the embedding of knowledge graph (UMLS) into the BioELMo to improve its performance, which shows an improvement of 0.8% regarding the accuracy to the base BioELMo model. Yadav et al [341]⁵⁴ a novel framework Sem-KGN for the medical textual entailment task, which infused the medical entity information from the medical knowledge bases into the BERT model. They show the medical knowledge information improves the SOTA language model ClinicalBERT by 8.27% on the REQ dataset. He et al [101]⁵⁵ proposed to infuse the domain knowledge of disease into a series of PLMs including BERT, BioBERT, SciBERT, ClinicalBERT, BlueBERT, and ALBERT, which improves performances of these models in all cases. Zhu et al [377] utilized the neural architecture search (NAS) to automatically find a better transformer structure for language models, which improves the performance of the Chinese BERT-wwm-ext model [61] on two Chinese NLI datasets. We summarize all available datasets in Table 23, and compare performances of different methods in Table 24. We can find that Wu et al [328] using the ensemble method significantly outperforms other methods in RQE. Among various PLMs including BioELMo, BERT, BioBERT, SciBERT, ClinicalBERT, BlueBERT, and ALBERT, ALBERT achieved the best performance on the MedNLI dataset.

4.8 Proteins/DNAs Prediction

In this section, we only list some applications that have been well-investigated or have potential, although there are much bigger spaces in biomedical domains to make use of PLMs.

⁵²https://github.com/ZhaofengWu/MEDIQA_WTMED

⁵³<https://github.com/soumyaah/KGMedNLI>

⁵⁴<https://github.com/VishalPallagani/Medical-Knowledge-enriched-Textual-Entailment>

⁵⁵<https://github.com/heyunh2015/diseaseBERT>

Table 24. Performances (accuracy for MedNLI and REQ, F1 score for CMFAQ) of different methods.

	MedNLI	RQE	CMFAQ
Wu et al [328]	-	98.00	-
Sharma et al [265]	79.04	-	-
Yadav et al [341]	-	56.17	-
He et al [101]	79.49	-	-
Zhu et al [377]	-	-	88.9

4.8.1 Protein structure predictions. Proteins are essential to life, and knowing their structure can facilitate our understanding of their function. However, the structure of only a small fraction of proteins is known [127]. Predicting the 3D structure of a protein is based solely on its amino acid sequence, a.k.a, ‘protein folding problem’ [15]. To evaluate protein structure predictions, CASP (Critical Assessment of Structure Prediction) uses proteins with recently solved structures that have not been deposited in the PDB or publicly disclosed; it therefore, is a blind test for the participants, which is the gold-standard assessment for protein structure predictions [151, 209]. In CASP14, AlphaFold 2 [127], a model designed by DeepMind achieves much better performance than other participating methods (e.g. template-based methods). The authors claim that AlphaFold 2 could provide precise estimates and could be confidently used for protein structure predictions with high reliability. However, predictions of existing methods, including the AlphaFold 2 are more family-specific than protein-specific, and rely on the evolutionary information captured in multiple sequence alignments (MSAs). To solve these issues, citeweissenow2021protein proposed to use the attention head from the pre-trained protein language model ProtT5 without MSAs. Recently, Sturmfels et al [277] presented a new biologically-informed pre-training task: predicting protein profiles derived from multiple sequence alignments, which can improve the downstream protein structure prediction task.

4.8.2 DNA related applications. There are few works in DNA pre-training, among which DNABERT [119] is the representative one. DNABERT not only achieved SOTA performance on promoter prediction, splice sites and transcription factor binding sites, but also identify functional genetic variants. Hong et al [109] proposed to pre-train DNA vectors to encode enhancers and promoters, and then Incorporated the attention mechanism to predict long-range enhancer–promoter interactions (EPs). Yamada et al [343] proposed a novel method based on the BERT to predict the interactions between RNA sequences and RNA-binding proteins (RBPs), in which BERT is pre-trained on the human reference genome. Mock et al [205] presented the BERTax based on BERT, for the taxonomic classification of DNA sequences.

4.9 Competitions and Venues

To facilitate the technological developments in biomedical text mining, many shared tasks and competitions have been organized since several years ago, focusing on various important tasks in the biomedical domain.

- **BioNLP workshop.** The BioNLP workshop⁵⁶ has been organized for 20 years and continually promoted the development of the biomedical domain, in which the community proposed a series of shared tasks and benchmark datasets. In BioNLP 2019, the BioNLP Open Shared Tasks (BioNLP-OST) 2019 [123] and the MEDIQA 2019 Shared Task [3] were organized. The BioNLP-OST 2019 proposed six tasks, including the information extraction on the bacterial biotopes and phenotypes, event extraction of genetic and molecular mechanisms, pharmacological

⁵⁶<https://aclweb.org/aclwiki/SIGBIOMED>

substances, compounds and proteins named entity recognition, integrated structure, semantics and coreference task, concept extraction for drug repurposing, and the information retrieval task for neuroscience. The MEDIQA 2019 aims to explore the method development on the natural language inference (NLI), recognizing question entailment (RQE), and question answering (QA) in the medical domain. In bioNLP 2021, the MEDIQA 2021 [2] shared tasks have three tasks related to the summarization of medical documents, including the question summarization task, the multi-answer summarization task, and the radiology report summarization task.

- **BioNLP-OST.** The BioNLP Open Shared Tasks (BioNLP-OST)⁵⁷ has been proposed since 2009 and was motivated to facilitate the development and sharing of methods on various tasks of biomedical text mining. It is organized every two years and organized at different conferences such as BioNLP and EMNLP. The latest BioNLP-OST 2019 is organized at the BioNLP 2019 as introduced aforementioned.
- **BioASQ.** The BioASQ⁵⁸ organizes workshops and challenges on biomedical semantic indexing and question answering. It has been held annually since 2013. In BioASQ 2019, the large-scale biomedical semantic indexing task, the biomedical information retrieval and question-answering task, and corresponding benchmark datasets are proposed.
- **BioCreAtIvE.** The Critical Assessment of Information Extraction systems in Biology (BioCreAtIvE)⁵⁹ organized challenge evaluations for the text mining and information extraction method on the biological domain since 2004. The latest BioCreative VII Challenge proposed five tracks, of which two tracks are related to COVID-19, including text mining and multi-label topic classification.
- **TREC.** The Text REtrieval Conference (TREC⁶⁰) organizes workshops for supporting the development of information retrieval methods based on large test collections. It was started in 1992 and held annually. It has biomedical tracks focusing on clinical decision support, precision medicine, and clinical trials et al.
- **eHealth-KD.** The eHealth-KD⁶¹ organizes challenges on the structure knowledge extraction of eHealth documents in the Spanish Language. The eHealth-KD Challenge 2019 proposed the key phrases identification and classification task, and the semantic relations detection task.
- **#SMM4H.** The Social Media Mining for Health Applications (#SMM4H)⁶² held workshops and shared tasks related to natural language processing challenges in social media data for health research since 2015 annually. The shared tasks in the #SMM4H '21 involve the information processing methods on Twitter related to COVID-19, self-report of breast cancer, adverse effect mentions, medication regimen, and adverse pregnancy outcomes.

Moreover, there are some challenges proposed recently, such as the COVID-19 Open Research Dataset Challenge (CORD-19)⁶³ in response to the COVID-19 pandemic, EHR DREAM Challenge⁶⁴ proposed in October 2019 and focusing on using electronic health record data to predict patient mortality, and ICLR 2021 workshop⁶⁵ devoting to propose machine learning methods for preventing and combating pandemics. Furthermore, since the continual development of pre-trained language models from 2018, in recently organized challenges, most participants proposed pre-trained language model-based methods for different tasks.

⁵⁷<https://2019.bionlp-ost.org/home>

⁵⁸<http://bioasq.org>

⁵⁹<https://biocreative.bioinformatics.udel.edu>

⁶⁰<https://trec.nist.gov>

⁶¹<https://knowledge-learning.github.io/ehealthkd-2019/>

⁶²<https://healthlanguageprocessing.org/smm4h-2021/>

⁶³<https://www.kaggle.com/allen-institute-for-ai/CORD-19-research-challenge>

⁶⁴<https://www.synapse.org/#!/Synapse:syn18405991/wiki/589657>

⁶⁵<https://mlpcp21.github.io/pages/challenge>

5 DISCUSSION

5.1 Limitations and Concerns

In this subsection, we will mainly discuss the limitations of biomedical PLMs and raise some concerns about them.

Misinformation. The training corpora consist of EHR, and social media may include wrong information. Thus, pre-trained language models pre-trained on them may convey some misinformation [335]. Furthermore, the biomedical domain itself may have misclassified disease definitions during its development process. Misinformation has become much more serious in the biomedical domain than in the general domain since this may lead to fatal biomedical decision-making consequences. However, researchers must be aware of the complexity of routinely collected electronic health records, including ways to manage variable completeness. We believe that the predictions from pre-trained language models should be artificially calibrated by biomedical experts before it is used by end-users like patient or the public.

Interpretation issues. Along with the power of neural networks, there is a growing concern about the interpretability of deep neural networks (DNNs). While in the biomedical domain, the consequence of bad decisions/predictions may be deadly; thus, a well-interpreted model is more crucial. The interpretation in the biomedical domain may come from two aspects: (1) biomedical models should be easily understood, and the predictions could be simulated from the raw input, (2) a (textual) reason should be provided for each prediction. The basic example of the former (a.k.a, transparency [181]) is decision trees that could clearly illustrate the decision path. However, such a transparency goal is hardly achieved in modern natural language processing, especially with pre-trained language models. More efforts could be made for the latter; one has to find some textual explanation for each prediction/decision, based on what doctors and patients could make their own decisions.

Identifying causalities from correlations. Similar to interpretability, causality may provide the underlying explanation of the model decisions. Causality is crucial in many tasks of biomedical knowledge, e.g., diagnosis, pathology, or systems biology. Causal associations between biological entities, events, and processes are central to most claims of interest; see an early review from [143]. With automatic causality recognition, it could suggest possible causal connections that may be beneficial for biomedical decisions, which hence greatly reduces the human workload [199].

Trade-off between coverage or quality? There are no large-scaled and high-quality training corpora in the biomedical domain. This means one has to sacrifice its coverage to obtain a high-quality vertical application, or train a general model with large-scaled yet low-quality corpora. Pre-trained language models typically consist of many transformer layers that have many parameters, which usually require a massive amount of plain text. This may lead to a general model with great coverage, but a smaller proportion of high-quality expert knowledge.

Heterogeneous training data. For biomedical understandings, there is heterogeneous information, including tables, figures, graphs (fMRI), etc. For example, tables and numbers are crucial in scientific literature. But most PLMs are unable to interpret tables and numbers well. To deeply capture the information in these heterogeneous data, both in-depth data preprocessing and model adaption may be needed. Especially, multi-domain pre-trained language models in biomedical should be paid much more attention.

Ethics and bias. With the rapid development of AI systems and applications in industrial products, it should be aware that they should not introduce any bias for special groups or populations [194], and some of the efforts were taken in the

NLP field [29, 83, 286, 374]. This becomes more crucial in these sensitive environments in the biomedical domain that involves life-changing decisions, like surgery [254]. It should ensure that the decisions cannot reflect discriminatory or biased behavior toward specific groups or certain populations. A few works have quantified the ethics and bias issues in the domain of pre-trained language models. [362] quantifies biases in clinical contextual word embeddings. The reason behind this may arise due to the training itself is biased with respect to various attributes like gender, language, race, age, ethnicity, and marital status. For example, in the MIMIC-III dataset [126], one can find: 1) gender bias: males have more heart disease than females, and 2) ethnicity bias: black patients have fewer clinical studies than other groups [129]. Considering the complexity of directly reducing biases in training corpora, existing works explore identifying bias by adversarial training [362], or data augmentation [202].

Privacy. Although most corpora used in biomedical pre-training like scientific publications and social medical are open-access. Some EHRs are private since some organizations do not want to expose their data. For example, clinical records may contain patient visits and medical history; these sensitive information may bring some physical and mental harm to patients if exposed [214]. Note that de-identification of these sensitive information in EHR records (like MIMIC III) is not always safe; recent works showed that there is data leakage from pre-trained models in the general domain, *i.e.* recovering Personal Health Information (PHI) from pre-trained models trained from is possible [157]. Therefore, we warn the public release of pre-trained models, if PHI is risky to be exposed. Recently, Nakamura et al [214] proposed a framework called ‘KART’ to assess the sensitive information from pre-trained biomedical language models using various attacks. Also, the federated learning [166, 345] framework may help when different organizations and end-users could collaboratively learn a shared prediction model while keeping all the training data on a private side.

5.2 Future trends

We further suggest some future trends in this subsection.

Standardized benchmark. In general NLP fields, evaluation criteria and standard benchmarks are a driving force for the NLP community. For example, BERT [66] was widely accepted in benchmarks [242, 307] makes it spread to various tasks in NLP. On the other hand, lacking an effective evaluation criterion is one of the bottlenecks of text generation [41]. In the biomedical domain, various pre-trained models and their fine-tuning applications have been proposed (as introduced in Sec. 3 and Sec. 4). However, they are generally not well-compared. Although a few efforts have been made to standardize benchmarks for biomedical pre-trained models, which include but are not limited to [90, 364]. This becomes much more difficult in the cross-discipline domain like the biomedical domain since papers are usually from different communities like informatics, medicine, and computer science. An open standardized and well-categorized benchmark (like in [159, 163]) should be proposed to make use of the advantages of each work and collaboratively push the development of biomedical NLP. This survey is the first step to introducing the biomedical pre-trained language models and their applications in downstream tasks. More efforts are expected to be made to design fine-grained taxonomy and define each SOTA approach in various applications, based on what incremental work could be better evaluated.

Open culture. In general NLP fields, a lot of effort is made to make better-available resources, including open-source resources (released training data and models), and fairly implemented approaches. In addition, open culture makes researchers could easily contribute to the community. For example, the NLP community has been largely developed thanks to the model collections [79, 325]. In addition, most accepted papers in top conferences tend to release codes,

models, and data. Biomedical NLP fields also benefit a lot from such open culture and standard systematic evaluations. For instance, pre-trained models in Huggingface ⁶⁶ largely fascinated their applications in the biomedical domain.

Efficiency on pre-trained language models. Compared to previous SOTA methods training from scratch based on neural networks such as LSTM or CNN, before Transformer, pre-trained language models are much bigger in terms of model scale and much slower due to the increasing number of parameters. This is more expensive for deployment that requires more computing resources. One may have to refer to [290] for efficient transformers. For example, current work explores quantization [18, 369], weights pruning [110], and knowledge distillation [120, 257] for BERT. Therefore, in the biomedical domain, pre-training language models with lower computation complexity are a direction needed to pay more attention.

Generation-based PLMs are under-investigated. Most works focused on encoder-based models, and a few works involve decoder or encoder-decoder architectures that enable generations. This may be due to the fact that classification tasks may be widely used in downstream biomedical tasks. Very recently, [146] proposes GPT models using temporal electronic health records and [232] trained a T5-based biomedical pre-trained model. We believe that generation-based PLMs (e.g. GPT, T5, and BART) have great potential in the biomedical domain, but it is currently under-investigated. Very recently, we have witnessed some work that uses large generation-based PLMs in the biomedical domain, see especially BIOGPT [187], PubMedGPT ⁶⁷, and Flan-PaLM [270].

Few-shot learning. [229] evaluates the few-shot ability of LMs when held-out examples are unavailable for choosing hyperparameters or prompts and finds that LMs do not perform well compared to random selection and under-perform selection based on held-out examples. In other words, previous methods overestimate the few-shot capability of LMs based on more realistic settings. This might be even worse for biomedical LMs.

In non-English or low-resource language. Most works in biomedical pre-trained language models are with English corpora, and a few about Chinese [365], German [36], Japanese [132, 304], Spanish [7, 8, 186, 204], Korean [141], Russian [299], Italian [40], Arabic [17, 33], French [60], Portuguese [259, 260] etc. For non-English biomedical tasks, there are two mainstream solutions: a single non-English language paradigm and a multi-linguistic paradigm. The former uses a single language, while the latter uses multiple languages. The multi-linguistic paradigm could be more beneficial for low-resource since biomedical knowledge itself is language-independent, and information in a second language could be complementary.

Multi-modal pre-training. Multi-modal pre-training [239, 244] has attracted much attention in image classification and generation tasks, because it only needs cheap but large-scale publicly available online resources. This shows great potential in machine learning since less human annotation is needed. It is expected that various modalities could provide complementary information. For example, making use of biomedical codes, medical images, waveforms, and genomics in pre-training models would be beneficial but challenging due to its multi-modal nature.

Injecting biomedical knowledge in pre-trained language model. Before the pre-training age, some works [231] have explored injecting medical knowledge into embeddings that provide potentially better ML features. Recently, existing work claims that pre-trained language models could be a soft knowledge base that captures knowledge. Despite this, [59, 338, 363] also tried to inject knowledge into pre-trained language models explicitly. In the biomedical domain,

⁶⁶<https://huggingface.co/>

⁶⁷<https://crfm.stanford.edu/2022/12/15/pubmedgpt.html>

which is knowledge-intensive; knowledge-injected models could have great potential in the future. For example, [198] integrates domain knowledge (*i.e.*, Unified Medical Language System (UMLS) Metathesaurus) in pre-training via a knowledge augmentation strategy.

Interpretability in biomedical PLMs. Neural networks were criticized for having limited interpretability. Pre-trained language models are typically huge neural network models, which is more challenging in terms of interpretability. One may expect to understand the working mechanism related to the medical characteristics in pre-trained language models. For example, probing pre-trained language models have been widely used to understand pre-trained language models, see [145, 175, 303, 327]. For biomedical pre-trained language models, [10] aims to evaluate pre-trained language models about the disease knowledge. [302] exhaustively analyzing attention in protein Transformer models, providing many interesting findings to understand the working mechanisms better. [121] conducts some probing experiments to determine what additional information is carried intrinsically by BioELMo and BioBERT. Another direction of interpretability in the biomedical field is to mine the causality (rather than correlation) due to its crucial relevance in establishing clinical interventions and public health policies. Correlation merely indicates a statistical relationship between two variables, which is valuable in generating hypotheses, but provides limited insights into the underlying mechanisms. Conversely, causality moves beyond associative relationships to delineate direct cause-effect relationships. This deeper understanding is pivotal in biomedical research, as it provides the foundation for intervention studies and enables the development of effective treatments. Identifying a causal relationship, for instance, between a specific genetic mutation and a disease, allows for targeted therapies and personalized medicine. Thus, while correlation provides a starting point for scientific exploration, it is the discernment of causality that truly advances biomedical knowledge and contributes to the development of life-saving interventions.

Dialogue-based medical consultation. Transitional medical consultation is to obtain medical suggestions and treatment from clinicians. Recently, AI communities have tried to solve medical consultations through online ways using artificial intelligence tools, especially for pre-consultation and psychological treatment. Meanwhile, online medical consultation is another natural playground for current state-of-the-art AI algorithms under the ‘AI for Science’ Trend. Some existing work formulate medical consultation as a question-answering task in the sense that it could leverage many existing question-answering pairs. However, medical consultation is complicated in the sense that static and single-turn question-answering pairs could not solve individually-dependent consultation; especially, medical consultations are more likely to be dependent on individual backgrounds, like historical diseases and treatment, genes, and dietary habits. We believe dialogue-based consultation systems could better fit medical scenarios than single-question-answering systems. Existing medical dialogue systems have shown some potential but also perform much worse than the expectation. Very recently, motivated by the great success of Open AI ChatGPT which uses giant language models to meet human consultation needs, we believe using giant medical language models could largely improve the quality of medical consultation. More optimistically, we believe this might, at least to some extent, revolutionize the current medical industry, see [153, 219] as some preliminary work.

Scale law in medical PLMs. Not only in dialogue systems, large-scale PLMs are as popular as it in the general domain. The reasons are twofold. First, the adaption of SOTA PLMs to the medical domain takes time, and it is usually more than half a year late after a general PLM is released. Secondly, non-generative language models are insensitive to huger scales, and their performance becomes saturated when they are beyond 24 layers (the scale of BERT-large). Meanwhile, most works use non-generative language models (*e.g.*, BERT, RoBERTa, and Electra) in the biomedical domain while

very few generative language models are used. With huger language models, we might see some emergent abilities in medical applications. Fortunately, we have witnessed a preliminary sign that we started to investigate large language models in medical/clinical tasks [174, 187, 270], especially BIOGPT [187], PubMedGPT⁶⁸, and Flan-PaLM [270].

Data collection and sharing protocol. The need for data in biomedicine is tremendous since data is the fuel for learning. The reasons that hinder medical data collection and sharing are manifold. Firstly, it has a legal risk regarding privacy issues, especially because this also involves cross-border or cross-organization data transfer. Secondly, an individual hospital might adopt different standards in terminology, this issue becomes more severe in developing countries than in developed countries. The merge between two data sources will be difficult due to the inconsistency of terminology. Therefore, it requires a well-defined protocol to deal with this, including solving terminology inconsistency and data privacy. From an NLP perspective, we need to normalize word terminology and data desensitization. For other perspectives, this needs some high-level data-sharing protocol, e.g., federated learning [346].

Dealing with long sequences. The computation of self-attentions in Transformers is quadratic to the length of sequences. This means the longer sequences would necessarily make transformer-based PLMs much more time-consuming. Sequences in biomedicine are usually long; it varies from DNA/protein sequences to texts. First, DNA/protein sequences are long especially for big protein sequence which has lengths that are longer than 4096, i.e., the typical maximum sequence length in language models. Biomedical texts, including EHRs, biomedical encyclopedias, and biomedical literature, are usually longer than the general domain (e.g., the maximum sentence length used in GLUE is usually 128); for instance, there is usually text redundancy in clinical notes. Therefore, we need to design more efficient and effective models tailored to long sequences, see some existing recent works [169, 170].

6 CONCLUSION

This paper systematically summarizes recent advances of pre-trained language models in the biomedical domain, including background, why and how pre-trained language models are used in the biomedical domain, existing biomedical pre-trained language models, data sources in the biomedical domain, application of pre-trained language models in various biomedical downstream tasks. Furthermore, we also discuss some limitations and future trends. Finally, we expect that the pre-trained language model in the general NLP domain could also help the specific biomedical domain.

ACKNOWLEDGMENT

This work is supported by Chinese Key-Area Research and Development Program of Guangdong Province (2020B0101350001), the Shenzhen Science and Technology Program (JCYJ20220818103001002), the Guangdong Provincial Key Laboratory of Big Data Computing, The Chinese University of Hong Kong, Shenzhen, Shenzhen Key Research Project (C10120230151) and Shenzhen Doctoral Startup Funding (RCBS20221008093330065).

REFERENCES

- [1] A. B. Abacha and D. Demner-Fushman. Recognizing question entailment for medical question answering. In *AMIA Annual Symposium Proceedings*, volume 2016, page 310. American Medical Informatics Association, 2016.
- [2] A. B. Abacha, Y. M’rabet, Y. Zhang, C. Shivade, C. Langlotz, and D. Demner-Fushman. Overview of the mediqua 2021 shared task on summarization in the medical domain. In *Workshop on Biomedical Language Processing*, pages 74–85, 2021.
- [3] A. B. Abacha, C. Shivade, and D. Demner-Fushman. Overview of the mediqua 2019 shared task on textual inference, question entailment and question answering. In *BioNLP Workshop and Shared Task*, pages 370–379, 2019.

⁶⁸<https://crfm.stanford.edu/2022/12/15/pubmedgpt.html>

- [4] A. Adhikari, A. Ram, R. Tang, and J. Lin. Rethinking complex neural network architectures for document classification. In *NAACL*, pages 4046–4051, 2019.
- [5] A. G. Agirre, M. Marimon, A. Intxaurreondo, O. Rabal, M. Villegas, and M. Krallinger. Pharmaconer: pharmacological substances, compounds and proteins named entity recognition track. In *Workshop on BioNLP Open Shared Tasks*, pages 1–10, 2019.
- [6] A. Akdemir and T. Shibuya. Transfer learning for biomedical question answering. In *CLEF (Working Notes)*, 2020.
- [7] L. Akhtyamova. Named entity recognition in spanish biomedical literature: short review and bert model. In *FRUCT*, pages 1–7. IEEE, 2020.
- [8] L. Akhtyamova, P. Martínez, K. Verspoor, and J. Cardiff. Testing contextualized word embeddings to improve ner in spanish clinical case narratives. *IEEE Access*, 8:164717–164726, 2020.
- [9] M. A. Al-Garadi, Y.-C. Yang, H. Cai, Y. Ruan, K. O’Connor, G.-H. Graciela, J. Perrone, and A. Sarker. Text classification models for the automatic detection of nonmedical prescription medication use from social media. *BMC medical informatics and decision making*, 21(1):1–13, 2021.
- [10] I. Alghanmi, L. Espinosa-Anke, and S. Schockaert. Probing pre-trained language models for disease knowledge. *arXiv preprint arXiv:2106.07285*, 2021.
- [11] E. C. Alley, G. Khimulya, S. Biswas, M. AlQuraishi, and G. M. Church. Unified rational protein engineering with sequence-only deep representation learning. *bioRxiv*, page 589333, 2019.
- [12] E. Alsentzer, J. Murphy, W. Boag, W.-H. Weng, D. Jindi, T. Naumann, and M. McDermott. Publicly available clinical bert embeddings. In *Proceedings of the 2nd Clinical Natural Language Processing Workshop*, pages 72–78, 2019.
- [13] E. Alsentzer, J. R. Murphy, W. Boag, W. Weng, D. Jin, T. Naumann, and M. B. A. McDermott. Publicly available clinical Bert embeddings. *CoRR*, abs/1904.03323, 2019.
- [14] S. Ananiadou, S. Pyysalo, J. Tsujii, and D. B. Kell. Event extraction for systems biology by text mining the literature. *Trends in biotechnology*, 28(7):381–390, 2010.
- [15] C. B. Anfinsen. Principles that govern the folding of protein chains. *Science*, 181(4096):223–230, 1973.
- [16] R. Anil, A. M. Dai, O. Firat, M. Johnson, D. Lepikhin, A. Passos, S. Shakeri, E. Taropa, P. Bailey, Z. Chen, et al. Palm 2 technical report. *arXiv preprint arXiv:2305.10403*, 2023.
- [17] W. Antoun, F. Baly, and H. Hajj. Arabert: transformer-based model for arabic language understanding. *arXiv preprint arXiv:2003.00104*, 2020.
- [18] H. Bai, W. Zhang, L. Hou, L. Shang, J. Jin, X. Jiang, Q. Liu, M. Lyu, and I. King. Binarybert: pushing the limit of bert quantization, 2020.
- [19] P. Banerjee, K. K. Pal, M. Devarakonda, and C. Baral. Biomedical named entity recognition via knowledge guidance and question answering. *ACM Transactions on Computing for Healthcare*, 2(4):1–24, 2021.
- [20] S. Banerjee and A. Lavie. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pages 65–72, 2005.
- [21] H. Bao, L. Dong, S. Piao, and F. Wei. Beit: Bert pre-training of image transformers. In *International Conference on Learning Representations*, 2021.
- [22] M. Basaldella, F. Liu, E. Shareghi, and N. Collier. Cometa: a corpus for medical entity linking in the social media. *arXiv preprint arXiv:2010.03295*, 2020.
- [23] I. Beltagy, K. Lo, and A. Cohan. Scibert: A pretrained language model for scientific text. In K. Inui, J. Jiang, V. Ng, and X. Wan, editors, *EMNLP-IJCNLP*, pages 3613–3618. Association for Computational Linguistics, 2019.
- [24] I. Beltagy, M. E. Peters, and A. Cohan. Longformer: the long-document transformer. *arXiv preprint arXiv:2004.05150*, 2020.
- [25] Y. Bengio, R. Ducharme, P. Vincent, and C. Jauvin. A neural probabilistic language model. *Journal of machine learning research*, 3(Feb):1137–1155, 2003.
- [26] T. Bepler and B. Berger. Learning protein sequence embeddings using information from structure. In *ICLR*, 2018.
- [27] T. Bepler and B. Berger. Learning protein sequence embeddings using information from structure. *arXiv preprint arXiv:1902.08661*, 2019.
- [28] J. Bishop, Q. Xie, and S. Ananiadou. Gencomparesum: a hybrid unsupervised summarization method using salience. In *Proceedings of the 21st Workshop on Biomedical Language Processing*, pages 220–240, 2022.
- [29] S. L. Blodgett, S. Barocas, H. Daumé III, and H. Wallach. Language (technology) is power: a critical survey of “bias” in nlp. *arXiv preprint arXiv:2005.14050*, 2020.
- [30] O. Bodenreider. The unified medical language system (umls): integrating biomedical terminology. *Nucleic acids research*, 32(suppl_1):D267–D270, 2004.
- [31] B. Boecking, N. Usuyama, S. Bannur, D. C. Castro, A. Schwaighofer, S. Hyland, M. Wetscherek, T. Naumann, A. Nori, J. Alvarez-Valle, et al. Making the most of text semantics to improve biomedical vision–language processing. *arXiv preprint arXiv:2204.09817*, 2022.
- [32] R. Bommasani and et al. On the opportunities and risks of foundation models, 2021.
- [33] N. Boudjellal, H. Zhang, A. Khan, A. Ahmad, R. Naseem, J. Shang, and L. Dai. Abioner: a bert-based model for arabic biomedical named-entity recognition. *Complexity*, 2021, 2021.
- [34] N. Brandes, D. Ofer, Y. Peleg, N. Rappoport, and M. Linial. Proteinbert: A universal deep-learning model of protein sequence and function. *Bioinformatics*, 38(8):2102–2110, 2022.
- [35] Á. Bravo, J. Piñero, N. Queralt-Rosinach, M. Rautschka, and L. I. Furlong. Extraction of relations between genes and diseases from text and large-scale data analysis: implications for translational research. *BMC bioinformatics*, 16(1):1–17, 2015.
- [36] K. K. Bressem, L. C. Adams, R. A. Gaudin, D. Tröltzsch, B. Hamm, M. R. Makowski, C.-Y. Schüle, J. L. Vahldiek, and S. M. Niehues. Highly accurate classification of chest radiographic reports using a deep learning natural language model pre-trained on 3.8 million text reports. *Bioinformatics*,

- 36(21):5255–5261, 2020.
- [37] E. Brochu, V. M. Cora, and N. De Freitas. A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv preprint arXiv:1012.2599*, 2010.
 - [38] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020.
 - [39] B. Cao, H. Lin, X. Han, L. Sun, L. Yan, M. Liao, T. Xue, and J. Xu. Knowledgeable or educated guess? revisiting language models as knowledge bases. In *ACL*, pages 1860–1874, 2021.
 - [40] R. Catelli, F. Gargiulo, V. Casola, G. De Pietro, H. Fujita, and M. Esposito. Crosslingual named entity recognition for clinical de-identification applied to a covid-19 italian data set. *Applied Soft Computing*, 97:106779, 2020.
 - [41] A. Celikyilmaz, E. Clark, and J. Gao. Evaluation of text generation: a survey. *arXiv preprint arXiv:2006.14799*, 2020.
 - [42] S. Chakraborty, E. Bisong, S. Bhatt, T. Wagner, R. Elliott, and F. Mosconi. Biomedbert: a pre-trained biomedical language model for qa and ir. In *ICCL*, pages 669–679, 2020.
 - [43] P. Chambon, C. Bluethgen, J.-B. Delbrouck, R. Van der Sluijs, M. Polacin, J. M. Z. Chaves, T. M. Abraham, S. Purohit, C. P. Langlotz, and A. Chaudhari. Roentgen: Vision-language foundation model for chest x-ray generation. *arXiv preprint arXiv:2211.12737*, 2022.
 - [44] P. Chambon, C. Bluethgen, C. P. Langlotz, and A. Chaudhari. Adapting pretrained vision-language foundational models to medical imaging domains. *arXiv preprint arXiv:2210.04133*, 2022.
 - [45] Q. Chen, J. Du, S. Kim, W. J. Wilbur, and Z. Lu. Evaluation of five sentence similarity models on electronic medical records. In *ACM-BCB*, pages 533–533, 2019.
 - [46] Q. Chen, J. Du, S. Kim, W. J. Wilbur, and Z. Lu. Deep learning with sentence embeddings pre-trained on biomedical corpora improves the performance of finding similar sentences in electronic medical records. *BMC medical informatics and decision making*, 20:1–10, 2020.
 - [47] Q. Chen, Y. Peng, and Z. Lu. Biosentvec: creating sentence embeddings for biomedical texts. In *ICHI*, pages 1–5. IEEE, 2019.
 - [48] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020.
 - [49] T. Chen, M. Wu, and H. Li. A general approach for improving deep learning-based medical relation extraction using a pre-trained model and fine-tuning. *Database*, 2019, 2019.
 - [50] Y.-C. Chen, L. Li, L. Yu, A. El Kholy, F. Ahmed, Z. Gan, Y. Cheng, and J. Liu. Uniter: Universal image-text representation learning. In *European conference on computer vision*, pages 104–120. Springer, 2020.
 - [51] Z. Chen, Y. Du, J. Hu, Y. Liu, G. Li, X. Wan, and T.-H. Chang. Multi-modal masked autoencoders for medical vision-and-language pre-training. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 679–689. Springer, 2022.
 - [52] Z. Chen, G. Li, and X. Wan. Align, reason and learn: Enhancing medical vision-and-language pre-training with knowledge. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 5152–5161, 2022.
 - [53] A. Chowdhery, S. Narang, J. Devlin, M. Bosma, G. Mishra, A. Roberts, P. Barham, H. W. Chung, C. Sutton, S. Gehrmann, et al. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*, 2022.
 - [54] K. Clark, M.-T. Luong, Q. V. Le, and C. D. Manning. Electra: pre-training text encoders as discriminators rather than generators, 2020.
 - [55] A. Cohan, F. Dernoncourt, D. S. Kim, T. Bui, S. Kim, W. Chang, and N. Goharian. A discourse-aware attention model for abstractive summarization of long documents. In *NAACL-HLT*, pages 615–621, 2018.
 - [56] A. M. Cohen and W. R. Hersh. A survey of current work in biomedical text mining. *Briefings in bioinformatics*, 6(1):57–71, 2005.
 - [57] J. Cohen. Bioinformatics—an introduction for computer scientists. *ACM Computing Surveys (CSUR)*, 36(2):122–158, June 2004.
 - [58] K. B. Cohen and D. Demner-Fushman. *Biomedical natural language processing*, volume 11. John Benjamins Publishing Company, 2014.
 - [59] P. Colon-Hernandez, C. Havasi, J. Alonso, M. Huggins, and C. Breazeal. Combining pre-trained language models and structured knowledge. *arXiv preprint arXiv:2101.12294*, 2021.
 - [60] J. Copara, J. Knafo, N. Naderi, C. Moro, P. Ruch, and D. Teodoro. Contextualized french language models for biomedical named entity recognition. In *6e conférence conjointe Journées d’Études sur la Parole (JEP, 31e édition), Traitement Automatique des Langues Naturelles (TALN, 27e édition), Rencontre des Étudiants Chercheurs en Informatique pour le Traitement Automatique des Langues (RÉCITAL, 22e édition)*, pages 36–48. ATALA, 2020.
 - [61] Y. Cui, W. Che, T. Liu, B. Qin, Z. Yang, S. Wang, and G. Hu. Pre-training with whole word masking for chinese bert. *arXiv preprint arXiv:1906.08101*, 2019.
 - [62] D. Demner-Fushman, M. D. Kohli, M. B. Rosenman, S. E. Shooshan, L. Rodriguez, S. Antani, G. R. Thoma, and C. J. McDonald. Preparing a collection of radiology examinations for distribution and retrieval. *Journal of the American Medical Informatics Association*, 23(2):304–310, 2016.
 - [63] D. Demner-Fushman, M. D. Kohli, M. B. Rosenman, S. E. Shooshan, L. Rodriguez, S. Antani, G. R. Thoma, and C. J. McDonald. Preparing a collection of radiology examinations for distribution and retrieval. *Journal of the American Medical Informatics Association*, 23(2):304–310, 2016.
 - [64] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: a large-scale hierarchical image database. In *CVPR*, pages 248–255. Ieee, 2009.
 - [65] J. Devlin, M. Chang, K. Lee, and K. Toutanova. Bert: pre-training of deep bidirectional transformers for language understanding. In J. Burstein, C. Doran, and T. Solorio, editors, *NAACL-HLT*, pages 4171–4186. Association for Computational Linguistics, 2019.
 - [66] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. Bert: pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.

- [67] J. DeYoung, I. Beltagy, M. van Zuylen, B. Kuehl, and L. Wang. Ms²: Multi-document summarization of medical studies. In *EMNLP*, pages 7494–7513, 2021.
- [68] G. Doddington. Automatic evaluation of machine translation quality using n-gram co-occurrence statistics. In *Proceedings of the second international conference on Human Language Technology Research*, pages 138–145, 2002.
- [69] R. I. Dogan, A. Chatr-aryamontri, S. Kim, C.-H. Wei, Y. Peng, D. C. Comeau, and Z. Lu. Biocreative vi precision medicine track: creating a training corpus for mining protein-protein interactions affected by mutations. In *BioNLP*, pages 171–175, 2017.
- [70] R. I. Doğan, R. Leaman, and Z. Lu. Ncbi disease corpus: a resource for disease name recognition and concept normalization. *Journal of biomedical informatics*, 47:1–10, 2014.
- [71] Y. Dong, J.-B. Cordonnier, and A. Loukas. Attention is not all you need: pure attention loses rank doubly exponentially with depth. *arXiv preprint arXiv:2103.03404*, 2021.
- [72] N. Du, K. Chen, A. Kannan, L. Tran, Y. Chen, and I. Shafran. Extracting symptoms and their status from clinical conversations. In *ACL*, pages 915–925. Association for Computational Linguistics, 2019.
- [73] N. Du, M. Wang, L. Tran, G. Li, and I. Shafran. Learning to infer entities, properties and their relations from clinical conversations. In *EMNLP-IJCNLP*, pages 4978–4989. Association for Computational Linguistics, 2020.
- [74] Y. Du, Q. Li, L. Wang, and Y. He. Biomedical-domain pre-trained language model for extractive summarization. *Knowledge-Based Systems*, 199:105964, 2020.
- [75] Y. Du, B. Pei, X. Zhao, and J. Ji. Deep scaled dot-product attention based domain adaptation model for biomedical question answering. *Methods*, 173:69–74, 2020.
- [76] M. Eichelberg, T. Aden, J. Riesmeier, A. Dogac, and G. B. Laleci. A survey and analysis of electronic healthcare record standards. *ACM Computing Surveys (CSUR)*, 37(4):277–315, Dec. 2005.
- [77] A. Elnaggar, M. Heinzinger, C. Dallago, G. Rihawi, Y. Wang, L. Jones, T. Gibbs, T. Feher, C. Angerer, M. Steinegger, et al. Prottrans: towards cracking the language of life’s code through self-supervised deep learning and high performance computing. *arXiv preprint arXiv:2007.06225*, 2020.
- [78] A. Esteva, A. Kale, R. Paulus, K. Hashimoto, W. Yin, D. Radev, and R. Socher. Co-search: covid-19 information retrieval with semantic search, question answering, and abstractive summarization. *arXiv preprint arXiv:2006.09595*, 2020.
- [79] Y. Fan, L. Pang, J. Hou, J. Guo, Y. Lan, and X. Cheng. Matchzoo: a toolkit for deep text matching. *arXiv preprint arXiv:1707.07270*, 2017.
- [80] Z. Feng, D. Guo, D. Tang, N. Duan, X. Feng, M. Gong, L. Shou, B. Qin, T. Liu, D. Jiang, et al. Codebert: a pre-trained model for programming and natural languages. *arXiv preprint arXiv:2002.08155*, 2020.
- [81] S. Gao, M. Alawad, M. T. Young, J. Gounley, N. Schaefferkoetter, H.-J. Yoon, X.-C. Wu, E. B. Durbin, J. Doherty, A. Stroup, et al. Limitations of transformers on clinical text classification. *IEEE journal of biomedical and health informatics*, 2021.
- [82] S. Gao, O. Kotevska, A. Sorokine, and J. B. Christian. A pre-training and self-training approach for biomedical named entity recognition. *PloS one*, 16(2):e0246310, 2021.
- [83] I. Garrido-Muñoz, A. Montejo-Ráez, F. Martínez-Santiago, and L. A. Ureña-López. A survey on bias in deep nlp. *Applied Sciences*, 11(7):3184, 2021.
- [84] M. Gerner, G. Nenadic, and C. M. Bergman. Linnaeus: a species name identification system for biomedical literature. *BMC bioinformatics*, 11(1):1–17, 2010.
- [85] S. S. Gharebagh, N. Goharian, and R. Filice. Attend to medical ontologies: content selection for clinical abstractive summarization. In *ACL*, pages 1899–1905, 2020.
- [86] J. Giorgi, X. Wang, N. Sahar, W. Y. Shin, G. D. Bader, and B. Wang. End-to-end named entity recognition and relation extraction using pre-trained language models. *arXiv preprint arXiv:1912.13415*, 2019.
- [87] G. Gonzalez-Hernandez, A. Sarker, K. O’Connor, and G. Savova. Capturing the patient’s perspective: a review of advances in natural language processing of health-related text. *Yearbook of medical informatics*, 26(1):214, 2017.
- [88] N. Grabar, V. Claveau, and C. Dalloux. Cas: french corpus with clinical cases. In *Workshop on Health Text Mining and Information Analysis*, pages 122–128, 2018.
- [89] J.-B. Grill, F. Strub, F. Althché, C. Tallec, P. Richemond, E. Buchatskaya, C. Doersch, B. Avila Pires, Z. Guo, M. Gheshlaghi Azar, et al. Bootstrap your own latent—a new approach to self-supervised learning. *Advances in neural information processing systems*, 33:21271–21284, 2020.
- [90] Y. Gu, R. Tinn, H. Cheng, M. Lucas, N. Usuyama, X. Liu, T. Naumann, J. Gao, and H. Poon. Domain-specific language model pretraining for biomedical natural language processing. *Special Issue on Computational Methods for Biomedical Natural Language Processing, ACM Transactions on Computing for Healthcare (HEALTH)*, 3(1):1–23, 2021.
- [91] H. Guan, J. Li, H. Xu, and M. Devarakonda. Robustly pre-trained neural model for direct temporal relation extraction. *arXiv preprint arXiv:2004.06216*, 2020.
- [92] Y. Guo, X. Dong, M. A. Al-Garadi, A. Sarker, C. Paris, and D. M. Aliod. Benchmarking of transformer-based pre-trained models on social media text classification datasets. In *Workshop of the Australasian Language Technology Association*, pages 86–91, 2020.
- [93] Y. Guo, W. Qiu, Y. Wang, and T. Cohen. Automated lay language summarization of biomedical scientific reviews. *arXiv preprint arXiv:2012.12573*, 2020.
- [94] S. Gururangan, A. Marasovic, S. Swayamdipta, K. Lo, I. Beltagy, D. Downey, and N. A. Smith. Don’t stop pretraining: adapt language models to domains and tasks. In D. Jurafsky, J. Chai, N. Schluter, and J. R. Tetreault, editors, *ACL*, pages 8342–8360. Association for Computational Linguistics, 2020.

- [95] B. J. Gutierrez, J. Zeng, D. Zhang, P. Zhang, and Y. Su. Document classification for covid-19 literature. In *EMNLP: Findings*, pages 3715–3722, 2020.
- [96] X. Han, Z. Zhang, N. Ding, Y. Gu, X. Liu, Y. Huo, J. Qiu, L. Zhang, W. Han, M. Huang, Q. Jin, Y. Lan, Y. Liu, Z. Liu, Z. Lu, X. Qiu, R. Song, J. Tang, J.-R. Wen, J. Yuan, W. X. Zhao, and J. Zhu. Pre-trained models: past, present and future, 2021.
- [97] D. Hanahan and R. A. Weinberg. The hallmarks of cancer. *cell*, 100(1):57–70, 2000.
- [98] K. He, X. Chen, S. Xie, Y. Li, P. Dollár, and R. Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009, 2022.
- [99] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738, 2020.
- [100] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [101] Y. He, Z. Zhu, Y. Zhang, Q. Chen, and J. Caverlee. Infusing disease knowledge into bert for health question answering, medical inference and disease name recognition. In *EMNLP*, pages 4604–4614, 2020.
- [102] M. Heinzinger, A. Elnaggar, Y. Wang, C. Dallago, D. Nechaev, F. Matthes, and B. Rost. Modeling aspects of the language of life through transfer-learning protein sequences. *BMC bioinformatics*, 20(1):1–17, 2019.
- [103] M. Heinzinger, A. Elnaggar, Y. Wang, C. Dallago, D. Nechaev, F. Matthes, and B. Rost. Modeling the language of life–deep learning protein sequences. *bioRxiv*, page 614313, 2019.
- [104] D. Hendrycks and K. Gimpel. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*, 2016.
- [105] S. Henry, K. Buchan, M. Filannino, A. Stubbs, and O. Uzuner. 2018 n2c2 shared task on adverse drug events and medication extraction in electronic health records. *Journal of the American Medical Informatics Association*, 27(1):3–12, 2020.
- [106] M. Herrero-Zazo, I. Segura-Bedmar, P. Martinez, and T. Declerck. The ddi corpus: an annotated corpus with pharmacological substances and drug–drug interactions. *Journal of biomedical informatics*, 46(5):914–920, 2013.
- [107] E. Herrett, A. M. Gallagher, K. Bhaskaran, H. Forbes, R. Mathur, T. Van Staa, and L. Smeeth. Data resource profile: clinical practice research datalink (cprd). *International journal of epidemiology*, 44(3):827–836, 2015.
- [108] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [109] Z. Hong, X. Zeng, L. Wei, and X. Liu. Identifying enhancer–promoter interactions with neural network based on pre-trained dna vectors and attention mechanism. *Bioinformatics*, 36(4):1037–1043, 2020.
- [110] L. Hou, Z. Huang, L. Shang, X. Jiang, X. Chen, and Q. Liu. Dynabert: dynamic bert with adaptive width and depth, 2020.
- [111] T.-M. H. Hsu, W.-H. Weng, W. Boag, M. McDermott, and P. Szolovits. Unsupervised multimodal representation learning across medical images and reports. *arXiv preprint arXiv:1811.08615*, 2018.
- [112] J. Hu, Z. Li, Z. Chen, Z. Li, X. Wan, and T.-H. Chang. Graph enhanced contrastive learning for radiology findings summarization. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4677–4688, 2022.
- [113] K. Huang, J. Altsaas, and R. Ranganath. Clinicalbert: modeling clinical notes and predicting hospital readmission. *CoRR*, abs/1904.05342, 2019.
- [114] K. Huang, A. Singh, S. Chen, E. Moseley, C. ying Deng, N. George, and C. Lindvall. Clinical xlnet: modeling sequential clinical notes and predicting prolonged mechanical ventilation. *arXiv:1912.11975*, 2019.
- [115] S.-C. Huang, L. Shen, M. P. Lungren, and S. Yeung. Gloria: A multimodal global-local representation learning framework for label-efficient medical image recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3942–3951, 2021.
- [116] Z. Huang, Z. Zeng, B. Liu, D. Fu, and J. Fu. Pixel-bert: Aligning image pixels with text by deep multi-modal transformers. *arXiv preprint arXiv:2004.00849*, 2020.
- [117] M. Jeong, M. Sung, G. Kim, D. Kim, W. Yoon, J. Yoo, and J. Kang. Transferability of natural language inference to biomedical question answering. *arXiv preprint arXiv:2007.00217*, 2020.
- [118] K. Jha and A. Zhang. Continual knowledge infusion into pre-trained biomedical language models. *Bioinformatics*, 38(2):494–502, 2022.
- [119] Y. Ji, Z. Zhou, H. Liu, and R. V. Davuluri. Dnabert: pre-trained bidirectional encoder representations from transformers model for dna-language in genome. *bioRxiv*, 2020.
- [120] X. Jiao, Y. Yin, L. Shang, X. Jiang, X. Chen, L. Li, F. Wang, and Q. Liu. Tinybert: distilling bert for natural language understanding, 2020.
- [121] Q. Jin, B. Dhingra, W. Cohen, and X. Lu. Probing biomedical embeddings from language models. In *Workshop on Evaluating Vector Space Representations for NLP*, pages 82–89, 2019.
- [122] Q. Jin, B. Dhingra, Z. Liu, W. Cohen, and X. Lu. Pubmedqa: a dataset for biomedical research question answering. In *EMNLP-IJCNLP*, pages 2567–2577, 2019.
- [123] K. Jin-Dong, N. Claire, B. Robert, and D. Louise. Proceedings of the 5th workshop on bionlp open shared tasks. In *BioNLP Open Shared Tasks*, 2019.
- [124] A. E. Johnson, T. J. Pollard, S. J. Berkowitz, N. R. Greenbaum, M. P. Lungren, C.-y. Deng, R. G. Mark, and S. Horng. Mimic-cxr, a de-identified publicly available database of chest radiographs with free-text reports. *Scientific data*, 6(1):317, 2019.
- [125] A. E. Johnson, T. J. Pollard, N. R. Greenbaum, M. P. Lungren, C.-y. Deng, Y. Peng, Z. Lu, R. G. Mark, S. J. Berkowitz, and S. Horng. Mimic-cxr-jpg, a large publicly available database of labeled chest radiographs. *arXiv preprint arXiv:1901.07042*, 2019.
- [126] A. E. Johnson, T. J. Pollard, L. Shen, H. L. Li-Wei, M. Feng, M. Ghassemi, B. Moody, P. Szolovits, L. A. Celi, and R. G. Mark. Mimic-iii, a freely accessible critical care database. *Scientific data*, 3(1):1–9, 2016.

- [127] J. Jumper, R. Evans, A. Pritzel, T. Green, M. Figurnov, O. Ronneberger, K. Tunyasuvunakool, R. Bates, A. Židek, A. Potapenko, et al. Highly accurate protein structure prediction with alphafold. *Nature*, page 1, 2021.
- [128] D. Jurafsky. *Speech & language processing*. Pearson Education India, 2000.
- [129] K. S. Kalyan, A. Rajasekharan, and S. Sangeetha. Ammu—a survey of transformer-based biomedical pretrained language models. *arXiv preprint arXiv:2105.00827*, 2021.
- [130] K. S. Kalyan and S. Sangeetha. Secnlp: a survey of embeddings in clinical natural language processing. *Journal of biomedical informatics*, 101:103323, 2020.
- [131] S. Kamath, B. Grau, and Y. Ma. How to pre-train your model? comparison of different pre-training models for biomedical question answering. In *ECML-PKDD*, pages 646–660. Springer, 2019.
- [132] Y. Kawazoe, D. Shibata, E. Shinohara, E. Aramaki, and K. Ohe. A clinical specific bert developed with huge size of japanese clinical narrative. *medRxiv*, 2020.
- [133] M. R. Khan, M. Ziyadi, and M. AbdelHady. Mt-bioner: multi-task learning for biomedical named entity recognition using deep bidirectional transformers. *arXiv preprint arXiv:2001.08904*, 2020.
- [134] F. K. Khattak, S. Jebblee, C. Pou-Prom, M. Abdalla, C. Meaney, and F. Rudzicz. A survey of word embeddings for clinical text. *Journal of Biomedical Informatics*, X, 4:100057, 2019.
- [135] V. Kieuvoongngam, B. Tan, and Y. Niu. Automatic text summarization of covid-19 medical research articles using bert and gpt-2. *arXiv preprint arXiv:2006.01997*, 2020.
- [136] J.-D. Kim, T. Ohta, Y. Tsuruoka, Y. Tateisi, and N. Collier. Introduction to the bio-entity recognition task at jnlpba. In *NLPBA/BioNLP*, pages 70–75. Citeseer, 2004.
- [137] J.-D. Kim, Y. Wang, T. Takagi, and A. Yonezawa. Overview of genia event task in bionlp shared task 2011. In *BioNLP shared task 2011 workshop*, pages 7–15, 2011.
- [138] J.-D. Kim, Y. Wang, and Y. Yasunori. The genia event extraction shared task, 2013 edition-overview. In *Proceedings of the BioNLP Shared Task 2013 Workshop*, pages 8–15, 2013.
- [139] Y. Kim. Convolutional neural networks for sentence classification. In *EMNLP*, pages 1746–1751, 2014.
- [140] Y. Kim. Convolutional neural networks for sentence classification. *arXiv preprint arXiv:1408.5882*, 2014.
- [141] Y.-M. Kim and T.-H. Lee. Korean clinical entity recognition from diagnosis text using bert. *BMC Medical Informatics and Decision Making*, 20(7):1–9, 2020.
- [142] D. T. Kingsbury. Computational biology. *ACM Computing Surveys (CSUR)*, 28(1):101–103, Mar. 1996.
- [143] S. Kleinberg and G. Hripcsak. A review of causal inference for biomedical informatics. *Journal of biomedical informatics*, 44(6):1102–1112, 2011.
- [144] V. Kommaraju, K. Gunasekaran, K. Li, T. Bansal, A. McCallum, I. Williams, and A.-M. Istrate. Unsupervised pre-training for biomedical question answering. *arXiv preprint arXiv:2009.12952*, 2020.
- [145] F. Koto, J. H. Lau, and T. Baldwin. Discourse probing of pretrained language models. *arXiv preprint arXiv:2104.05882*, 2021.
- [146] Z. Kraljevic, A. Shek, D. Bean, R. Bendayan, J. Teo, and R. Dobson. Medgpt: medical concept prediction from clinical narratives. *arXiv preprint arXiv:2107.03134*, 2021.
- [147] M. Krallinger, F. Leitner, O. Rabal, M. Vazquez, J. Oyarzabal, and A. Valencia. Chemdner: the drugs and chemical names extraction challenge. *Journal of cheminformatics*, 7(1):1–11, 2015.
- [148] M. Krallinger, F. Leitner, C. Rodriguez-Penagos, and A. Valencia. Overview of the protein-protein interaction annotation extraction task of biocreative ii. *Genome biology*, 9(2):1–19, 2008.
- [149] M. Krallinger, O. Rabal, S. A. Akhondi, M. P. Pérez, J. Santamaría, G. P. Rodríguez, et al. Overview of the biocreative vi chemical-protein interaction track. In *BioCreative challenge evaluation workshop*, volume 1, pages 141–146, 2017.
- [150] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90, 2017.
- [151] A. Kryshchuk, T. Schwede, M. Topf, K. Fidelis, and J. Moulton. Critical assessment of methods of protein structure prediction (casp)—round xiii. *Proteins: Structure, Function, and Bioinformatics*, 87(12):1011–1020, 2019.
- [152] T. Kudo and J. Richardson. Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 66–71, 2018.
- [153] T. H. Kung, M. Cheatham, A. Medinilla, ChatGPT, C. Sillos, L. De Leon, C. Elepano, M. Madriaga, R. Aggabao, G. Diaz-Candido, et al. Performance of chatgpt on usmle: Potential for ai-assisted medical education using large language models. *medRxiv*, pages 2022–12, 2022.
- [154] Z. Lan, M. Chen, S. Goodman, K. Gimpel, P. Sharma, and R. Soricut. Albert: a lite bert for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942*, 2019.
- [155] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [156] J. Lee, W. Yoon, S. Kim, D. Kim, S. Kim, C. H. So, and J. Kang. Biobert: a pre-trained biomedical language representation model for biomedical text mining. *Bioinform.*, 36(4):1234–1240, 2020.
- [157] E. Lehman, S. Jain, K. Pichotta, Y. Goldberg, and B. C. Wallace. Does bert pretrained on clinical notes reveal sensitive data? *arXiv preprint arXiv:2104.07762*, 2021.

- [158] M. Lewis, Y. Liu, N. Goyal, M. Ghazvininejad, A. Mohamed, O. Levy, V. Stoyanov, and L. Zettlemoyer. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*, 2019.
- [159] P. Lewis, M. Ott, J. Du, and V. Stoyanov. Pretrained language models for biomedical and clinical tasks: understanding and extending the state-of-the-art. In *Clinical Natural Language Processing Workshop*, pages 146–157, 2020.
- [160] D. Li, Z. Ren, P. Ren, Z. Chen, M. Fan, J. Ma, and M. de Rijke. Semi-supervised variational reasoning for medical dialogue generation. *SIGIR*, pages 544–554, 2021.
- [161] F. Li, Y. Jin, W. Liu, B. P. S. Rawat, P. Cai, and H. Yu. Fine-tuning bidirectional encoder representations from transformers (bert)-based models on large-scale electronic health record notes: an empirical study. *JMIR medical informatics*, 7(3):e14830, 2019.
- [162] J. Li, Y. Sun, R. J. Johnson, D. Sciaky, C.-H. Wei, R. Leaman, A. P. Davis, C. J. Mattingly, T. C. Wiegiers, and Z. Lu. Biocreative v cdr task corpus: a resource for chemical disease relation extraction. *Database*, 2016, 2016.
- [163] J. Li, X. Wang, X. Wu, Z. Zhang, X. Xu, J. Fu, P. Tiwari, X. Wan, and B. Wang. Huatuo-26m, a large-scale chinese medical qa dataset. *arXiv preprint arXiv:2305.01526*, 2023.
- [164] J. Li, Z. Zhang, H. Zhao, X. Zhou, and X. Zhou. Task-specific objectives of pre-trained language models for dialogue adaptation. *arXiv preprint arXiv:2009.04984*, 2020.
- [165] L. H. Li, M. Yatskar, D. Yin, C.-J. Hsieh, and K.-W. Chang. Visualbert: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557*, 2019.
- [166] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith. Federated learning: challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3):50–60, 2020.
- [167] Y. Li, S. Rao, J. R. A. Solares, A. Hassaine, R. Ramakrishnan, D. Canoy, Y. Zhu, K. Rahimi, and G. Salimi-Khorshidi. Behrt: transformer for electronic health records. *Scientific reports*, 10(1):1–12, 2020.
- [168] Y. Li, H. Wang, and Y. Luo. A comparison of pre-trained vision-and-language models for multimodal representation learning across medical images and reports. In *2020 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pages 1999–2004. IEEE, 2020.
- [169] Y. Li, R. M. Wehbe, F. S. Ahmad, H. Wang, and Y. Luo. Clinical-longformer and clinical-bigbird: Transformers for long clinical sequences. *arXiv preprint arXiv:2201.11838*, 2022.
- [170] Y. Li, R. M. Wehbe, F. S. Ahmad, H. Wang, and Y. Luo. A comparative study of pretrained language models for long clinical text. *Journal of the American Medical Informatics Association*, 30(2):340–347, 2023.
- [171] Z. Li, Y. Li, Q. Li, Y. Zhang, P. Wang, D. Guo, L. Lu, D. Jin, and Q. Hong. Lvit: language meets vision transformer in medical image segmentation. *arXiv preprint arXiv:2206.14718*, 2022.
- [172] K. Liao, Q. Liu, Z. Wei, B. Peng, Q. Chen, W. Sun, and X. Huang. Task-oriented dialogue system for automatic disease diagnosis via hierarchical reinforcement learning. *arXiv preprint arXiv:2004.14254*, 2020.
- [173] R. Liao, D. Moyer, M. Cha, K. Quigley, S. Berkowitz, S. Horng, P. Golland, and W. M. Wells. Multimodal representation learning via maximization of local mutual information. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 273–283. Springer, 2021.
- [174] V. Liévin, C. E. Hother, and O. Winther. Can large language models reason about medical questions? *arXiv preprint arXiv:2207.08143*, 2022.
- [175] B. Y. Lin, S. Lee, R. Khanna, and X. Ren. Birds have four legs?! numersense: probing numerical commonsense knowledge of pre-trained language models. *arXiv preprint arXiv:2005.00683*, 2020.
- [176] C. Lin, T. Miller, D. Dligach, S. Bethard, and G. Savova. A bert-based universal model for both within-and cross-sentence clinical temporal relation extraction. In *Clinical Natural Language Processing Workshop*, pages 65–71, 2019.
- [177] C. Lin, T. Miller, D. Dligach, F. Sadeque, S. Bethard, and G. Savova. A bert-based one-pass multi-task model for clinical temporal relation extraction. In *SIGBioMed Workshop on Biomedical Language Processing*, pages 70–75, 2020.
- [178] C.-Y. Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81, 2004.
- [179] S. Lin, P. Zhou, X. Liang, J. Tang, R. Zhao, Z. Chen, and L. Lin. Graph-evolving meta-learning for low-resource medical dialogue generation. In *AAAI*, pages 13362–13370. AAAI Press, 2021.
- [180] X. Lin, X. He, Q. Chen, H. Tou, Z. Wei, and T. Chen. Enhancing dialogue symptom diagnosis with global attention and symptom graph. In *EMNLP-IJCNLP*, pages 5033–5042, 2019.
- [181] Z. C. Lipton. The myths of model interpretability: in machine learning, the concept of interpretability is both important and slippery. *Queue*, 16(3):31–57, 2018.
- [182] P. Liu, W. Yuan, J. Fu, Z. Jiang, H. Hayashi, and G. Neubig. Pre-train, prompt, and predict: a systematic survey of prompting methods in natural language processing. *arXiv preprint arXiv:2107.13586*, 2021.
- [183] Q. Liu, M. J. Kusner, and P. Blunsom. A survey on contextual embeddings, 2020.
- [184] W. Liu, J. Tang, J. Qin, L. Xu, Z. Li, and X. Liang. Meddg: a large-scale medical consultation dataset for building medical dialogue system. *arXiv preprint arXiv:2010.07497*, 2020.
- [185] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov. Roberta: a robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.
- [186] P. López-Úbeda, M. C. Díaz-Galiano, L. A. Ureña-López, and M. T. Martín-Valdivia. Pre-trained language models to extract information from radiological reports. *CLEF eHealth*, 2021.

- [187] R. Luo, L. Sun, Y. Xia, T. Qin, S. Zhang, H. Poon, and T.-Y. Liu. Biogpt: generative pre-trained transformer for biomedical text generation and mining. *Briefings in Bioinformatics*, 23(6), 2022.
- [188] Z. Luo, Q. Xie, and S. Ananiadou. Readability controllable biomedical document summarization. *arXiv preprint arXiv:2210.04705*, 2022.
- [189] Z. Luo, Q. Xie, and S. Ananiadou. Chatgpt as a factual inconsistency evaluator for abstractive text summarization. *arXiv preprint arXiv:2303.15621*, 2023.
- [190] Z. Luo, Q. Xie, and S. Ananiadou. Citationsum: Citation-aware graph contrastive learning for scientific paper summarization. *arXiv preprint arXiv:2301.11223*, 2023.
- [191] A. Madani, B. McCann, N. Naik, N. S. Keskar, N. Anand, R. R. Eguchi, P.-S. Huang, and R. Socher. Progen: language modeling for protein generation. *arXiv preprint arXiv:2004.03497*, 2020.
- [192] T. Mairitha, N. Mairitha, and S. Inoue. Improving fine-tuned question answering models for electronic health records. In *UBICOMP*, pages 688–691, 2020.
- [193] A. Mascio, Z. Kraljevic, D. Bean, R. Dobson, R. Stewart, R. Bendayan, and A. Roberts. Comparative analysis of text classification approaches in electronic health records. In *Proceedings of the 19th SIGBioMed Workshop on Biomedical Language Processing*, pages 86–94, 2020.
- [194] N. Mehrabi, F. Morstatter, N. Saxena, K. Lerman, and A. Galstyan. A survey on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)*, 54(6):1–35, 2021.
- [195] X. Meng, C. H. Ganoe, R. T. Sieberg, Y. Y. Cheung, and S. Hassanpour. Self-supervised contextual language representation of radiology reports to improve the identification of communication urgency, 2019.
- [196] Y. Meng, W. F. Speier, M. K. Ong, and C. Arnold. Bidirectional representation learning from transformers using multimodal electronic health record data to predict depression. *IEEE Journal of Biomedical and Health Informatics*, pages 1–1, 2021.
- [197] A. Merchant, E. Rahimtoroghi, E. Pavlick, and I. Tenney. What happens to bert embeddings during fine-tuning? *arXiv preprint arXiv:2004.14448*, 2020.
- [198] G. Michalopoulos, Y. Wang, H. Kaka, H. Chen, and A. Wong. Umlsbert: clinical domain knowledge augmentation of contextual embeddings using the unified medical language system metathesaurus. *arXiv preprint arXiv:2010.10391*, 2020.
- [199] C. Mihăilă, T. Ohta, S. Pyysalo, and S. Ananiadou. Biocause: annotating and analysing causality in the biomedical domain. *BMC bioinformatics*, 14(1):1–18, 2013.
- [200] T. Mikolov, K. Chen, G. Corrado, and J. Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013.
- [201] T. Mikolov, I. Sutskever, K. Chen, G. Corrado, and J. Dean. Distributed representations of words and phrases and their compositionality. *NeuIPS*, 26:3111–3119, 2013.
- [202] J. R. Minot, N. Cheney, M. Maier, D. C. Elbers, C. M. Danforth, and P. S. Dodds. Interpretable bias mitigation for textual data: reducing gender bias in patient notes while maintaining classification performance. *arXiv preprint arXiv:2103.05841*, 2021.
- [203] G. Miolo, G. Mantoan, and C. Orsenigo. Electramed: a new pre-trained language representation model for biomedical nlp, 2021.
- [204] A. Miranda-Escalada, E. Farré, and M. Krallinger. Named entity recognition, concept normalization and clinical coding: overview of the cantemist track for cancer text mining in spanish, corpus, guidelines, methods and results. In *IberLEF*, 2020.
- [205] F. Mock, F. Kretschmer, A. Krieser, S. Böcker, and M. Marz. Bertax: taxonomic classification of dna sequences with deep neural networks. *BioRxiv*, 2021.
- [206] M. Monajatipoor, M. Rouhsedaghat, L. H. Li, C.-C. Jay Kuo, A. Chien, and K.-W. Chang. Berthop: An effective vision-and-language model for chest x-ray disease diagnosis. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 725–734. Springer, 2022.
- [207] J. H. Moon, H. Lee, W. Shin, Y.-H. Kim, and E. Choi. Multi-modal understanding and generation for medical images and text via vision-language pre-training. *IEEE Journal of Biomedical and Health Informatics*, 2022.
- [208] M. Mosbach, M. Andriushchenko, and D. Klakow. On the stability of fine-tuning bert: misconceptions, explanations, and strong baselines. *arXiv preprint arXiv:2006.04884*, 2020.
- [209] J. Moul, J. T. Pedersen, R. Judson, and K. Fidelis. A large-scale experiment to assess protein structure prediction methods, 1995.
- [210] M. Müller, M. Salathé, and P. E. Kummervold. Covid-twitter-bert: A natural language processing model to analyse Covid-19 content on twitter. *CoRR*, abs/2005.07503, 2020.
- [211] M. M. Müller and M. Salathé. Crowdbreaks: tracking health trends using public social media data and crowdsourcing. *Frontiers in public health*, 7:81, 2019.
- [212] P. Müller, G. Kaissis, and D. Rueckert. The role of local alignment and uniformity in image-text contrastive learning on medical images. *arXiv preprint arXiv:2211.07254*, 2022.
- [213] P. Müller, G. Kaissis, C. Zou, and D. Rueckert. Joint learning of localized representations from medical images and reports. In *European Conference on Computer Vision*, pages 685–701. Springer, 2022.
- [214] Y. Nakamura, S. Hanaoka, Y. Nomura, N. Hayashi, O. Abe, S. Yada, S. Wakamiya, and E. Aramaki. Kart: privacy leakage framework of language models pre-trained with clinical records. *arXiv preprint arXiv:2101.00036*, 2020.
- [215] U. Naseem, M. Khushi, V. Reddy, S. Rajendran, I. Razzak, and J. Kim. Bioalbert: a simple and effective pre-trained language model for biomedical named entity recognition. *arXiv preprint arXiv:2009.09223*, 2020.
- [216] C. Nédellec, R. Bossy, J.-D. Kim, J.-J. Kim, T. Ohta, S. Pyysalo, and P. Zweigenbaum. Overview of bionlp shared task 2013. In *Proceedings of the BioNLP shared task 2013 workshop*, pages 1–7, 2013.

- [217] A. Nentidis, K. Bougiatiotis, A. Krithara, and G. Paliouras. Results of the seventh edition of the bioasq challenge. In *ECML PKDD*, pages 553–568. Springer, 2019.
- [218] D. Q. Nguyen, T. Vu, and A. T. Nguyen. Bertweet: A pre-trained language model for english tweets. In *EMNLP*, pages 9–14, 2020.
- [219] O. Nov, N. Singh, and D. Mann. Putting chatgpt’s medical advice to the (turing) test, 2023.
- [220] B. Nye, J. J. Li, R. Patel, Y. Yang, I. J. Marshall, A. Nenkova, and B. C. Wallace. A corpus with multi-level annotations of patients, interventions and outcomes to support language processing for medical literature. In *ACL*, volume 2018, page 197. NIH Public Access, 2018.
- [221] T. Ohta, S. Pyysalo, and J. Tsujii. Overview of the epigenetics and post-translational modifications (epi) task of bionlp shared task 2011. In *BioNLP Shared Task 2011 Workshop*, pages 16–25, 2011.
- [222] I. B. Ozyurt. On the effectiveness of small, discriminatively pre-trained language representation models for biomedical text mining. In *Workshop on Scholarly Document Processing*, pages 104–112, 2020.
- [223] E. Pafilis, S. P. Frankild, L. Fanini, S. Faulwetter, C. Pavloudi, A. Vasileiadou, C. Arvanitidis, and L. J. Jensen. The species and organisms resources for fast and accurate identification of taxonomic names in text. *PLoS one*, 8(6):e65390, 2013.
- [224] A. Pampari, P. Raghavan, J. J. Liang, and J. Peng. Emrqa: a large corpus for question answering on electronic medical records. In *EMNLP*, 2018.
- [225] O. Pelka, S. Koitka, J. Rückert, F. Nensa, and C. M. Friedrich. Radiology objects in context (roco): a multimodal image dataset. In *Intravascular Imaging and Computer Assisted Stenting and Large-Scale Annotation of Biomedical Data and Expert Label Synthesis*, pages 180–189. Springer, 2018.
- [226] Y. Peng, S. Yan, and Z. Lu. Transfer learning in biomedical natural language processing: an evaluation of bert and elmo on ten benchmarking datasets. In *BioNLP Workshop and Shared Task*, pages 58–65, 2019.
- [227] J. Pennington, R. Socher, and C. D. Manning. Glove: global vectors for word representation. In *EMNLP*, pages 1532–1543, 2014.
- [228] B. Percha. Modern clinical text mining: A guide and review. *Annual Review of Biomedical Data Science*, 4:165–187, 2021.
- [229] E. Perez, D. Kiela, and K. Cho. True few-shot learning with language models. *arXiv preprint arXiv:2105.11447*, 2021.
- [230] M. E. Peters, M. Neumann, M. Iyyer, M. Gardner, C. Clark, K. Lee, and L. Zettlemoyer. Deep contextualized word representations. In *NAACL*, pages 2227–2237, 2018.
- [231] F. Petroni, T. Rocktäschel, S. Riedel, P. Lewis, A. Bakhtin, Y. Wu, and A. Miller. Language models as knowledge bases? In *EMNLP-IJCNLP*, pages 2463–2473, 2019.
- [232] L. N. Phan, J. T. Anibal, H. Tran, S. Chanana, E. Bahadroglu, A. Peltekian, and G. Altan-Bonnet. Scifive: a text-to-text transformer model for biomedical literature, 2021.
- [233] N. Poerner, U. Waltinger, and H. Schütze. Inexpensive domain adaptation of pretrained language models: Case studies on biomedical ner and covid-19 qa. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1482–1490, 2020.
- [234] S. Pyysalo, T. Ohta, M. Miwa, H.-C. Cho, J. Tsujii, and S. Ananiadou. Event extraction across multiple levels of biological organization. *Bioinformatics*, 28(18):i575–i581, 2012.
- [235] S. Pyysalo, T. Ohta, R. Rak, A. Rowley, H.-W. Chun, S.-J. Jung, S.-P. Choi, J. Tsujii, and S. Ananiadou. Overview of the cancer genetics and pathway curation tasks of bionlp shared task 2013. *BMC bioinformatics*, 16(10):1–19, 2015.
- [236] S. Pyysalo, T. Ohta, R. Rak, D. Sullivan, C. Mao, C. Wang, B. Sobral, J. Tsujii, and S. Ananiadou. Overview of the infectious diseases (id) task of bionlp shared task 2011. In *BioNLP Shared Task 2011 Workshop*, pages 26–35, 2011.
- [237] X. Qiu, T. Sun, Y. Xu, Y. Shao, N. Dai, and X. Huang. Pre-trained models for natural language processing: a survey. *Science China Technological Sciences*, pages 1–26, 2020.
- [238] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021.
- [239] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, et al. Learning transferable visual models from natural language supervision. *arXiv preprint arXiv:2103.00020*, 2021.
- [240] A. Radford, K. Narasimhan, T. Salimans, and I. Sutskever. Improving language understanding by generative pre-training. *OpenAI Technical Report*, 2018.
- [241] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer, 2020.
- [242] P. Rajpurkar, J. Zhang, K. Lopyrev, and P. Liang. Squad: 100,000+ questions for machine comprehension of text. *arXiv preprint arXiv:1606.05250*, 2016.
- [243] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, and M. Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022.
- [244] A. Ramesh, M. Pavlov, G. Goh, S. Gray, C. Voss, A. Radford, M. Chen, and I. Sutskever. Zero-shot text-to-image generation, 2021.
- [245] A. Ramponi, R. van der Goot, R. Lombardo, and B. Plank. Biomedical event extraction as sequence labeling. In *EMNLP*, pages 5357–5367, 2020.
- [246] R. Rao, N. Bhattacharya, N. Thomas, Y. Duan, X. Chen, J. Canny, P. Abbeel, and Y. S. Song. Evaluating protein transfer learning with tape. *NeuIPS*, 32:9689, 2019.
- [247] R. Rao, J. Liu, R. Verkuil, J. Meier, J. F. Canny, P. Abbeel, T. Sercu, and A. Rives. Msa transformer. *bioRxiv*, 2021.
- [248] L. Rasmy, Y. Xiang, Z. Xie, C. Tao, and D. Zhi. Med-bert: pre-trained contextualized embeddings on large-scale structured electronic health records for disease prediction. *CoRR*, abs/2005.12833, 2020.

- [249] B. P. S. Rawat, W.-H. Weng, S. Y. Min, P. Raghavan, and P. Szolovits. Entity-enriched neural models for clinical question answering. In *SIGBioMed Workshop on Biomedical Language Processing*, pages 112–122, 2020.
- [250] A. Rives, J. Meier, T. Sercu, S. Goyal, Z. Lin, J. Liu, D. Guo, M. Ott, C. L. Zitnick, J. Ma, et al. Biological structure and function emerge from scaling unsupervised learning to 250 million protein sequences. *Proceedings of the National Academy of Sciences*, 118(15), 2021.
- [251] A. Romanov and C. Shivade. Lessons from natural language inference in the clinical domain. In *EMNLP*, pages 1586–1596, 2018.
- [252] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022.
- [253] S. Rongali, A. Jagannatha, B. P. S. Rawat, and H. Yu. Continual domain-tuning for pretrained language models, 2021.
- [254] F. Rudzicz and R. Saqr. Ethics of artificial intelligence in surgery. *arXiv preprint arXiv:2007.14302*, 2020.
- [255] D. S. Sachan, P. Xie, M. Sachan, and E. P. Xing. Effective use of bidirectional language modeling for transfer learning in biomedical named entity recognition. In *Machine Learning for Healthcare Conference*, pages 383–402. PMLR, 2018.
- [256] C. Saharia, W. Chan, S. Saxena, L. Li, J. Whang, E. Denton, S. K. S. Ghasemipour, B. K. Ayan, S. S. Mahdavi, R. G. Lopes, et al. Photorealistic text-to-image diffusion models with deep language understanding. *arXiv preprint arXiv:2205.11487*, 2022.
- [257] V. Sanh, L. Debut, J. Chaumond, and T. Wolf. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter, 2020.
- [258] T. L. Scao, A. Fan, C. Akiki, E. Pavlick, S. Ilić, D. Hesslow, R. Castagné, A. S. Luccioni, F. Yvon, M. Gallé, et al. Bloom: A 176b-parameter open-access multilingual language model. *arXiv preprint arXiv:2211.05100*, 2022.
- [259] E. T. R. Schneider, J. V. A. de Souza, Y. B. Gumiel, C. Moro, and E. C. Paraiso. A gpt-2 language model for biomedical texts in portuguese. In *CBMS*, pages 474–479. IEEE, 2021.
- [260] E. T. R. Schneider, J. V. A. de Souza, J. Knafou, L. E. S. e. Oliveira, J. Copara, Y. B. Gumiel, L. F. A. d. Oliveira, E. C. Paraiso, D. Teodoro, and C. M. C. M. Barra. BioBERTpt - a Portuguese neural language model for clinical named entity recognition. In *Proceedings of the 3rd Clinical Natural Language Processing Workshop*, pages 65–72, Online, Nov. 2020. Association for Computational Linguistics.
- [261] C. Seibold, S. Reiß, M. S. Sarfraz, R. Stiefelhagen, and J. Kleesiek. Breaking with fixed set pathology recognition through report-guided contrastive training. *arXiv preprint arXiv:2205.07139*, 2022.
- [262] R. Sennrich, B. Haddow, and A. Birch. Neural machine translation of rare words with subword units. In *ACL*, pages 1715–1725. Association for Computational Linguistics (ACL), 2016.
- [263] J. Shang, T. Ma, C. Xiao, and J. Sun. Pre-training of graph augmented transformers for medication recommendation. *arXiv preprint arXiv:1906.00346*, 2019.
- [264] S. Sharma and R. D. J. au2. Bioflair: pretrained pooled contextualized embeddings for biomedical sequence labeling tasks, 2019.
- [265] S. Sharma, B. Santra, A. Jana, S. Tokala, N. Ganguly, and P. Goyal. Incorporating domain knowledge into medical NLI using knowledge graphs. In *EMNLP-IJCNLP*, pages 6092–6097, Nov. 2019.
- [266] G. Sheikhshab, I. Birol, and A. Sarkar. In-domain context-aware token embeddings improve biomedical named entity recognition. In *Workshop on Health Text Mining and Information Analysis*, pages 160–164, 2018.
- [267] X. Shi, H. Hu, W. Che, Z. Sun, T. Liu, and J. Huang. Understanding medical conversations with scattered keyword attention and weak supervision from responses. In *AAAI*, volume 34, pages 8838–8845, 2020.
- [268] H.-C. Shin, Y. Zhang, E. Bakhturina, R. Puri, M. Patwary, M. Shoenybi, and R. Mani. Biomegatron: larger biomedical domain language model. In *EMNLP*, pages 4700–4706, 2020.
- [269] Y. Si, J. Wang, H. Xu, and K. Roberts. Enhancing clinical concept extraction with contextual embeddings. *Journal of the American Medical Informatics Association*, 26(11):1297–1304, 2019.
- [270] K. Singhal, S. Azizi, T. Tu, S. S. Mahdavi, J. Wei, H. W. Chung, N. Scales, A. Tanwani, H. Cole-Lewis, S. Pfohl, et al. Large language models encode clinical knowledge. *arXiv preprint arXiv:2212.13138*, 2022.
- [271] K. Singhal, T. Tu, J. Gottweis, R. Sayres, E. Wulczyn, L. Hou, K. Clark, S. Pfohl, H. Cole-Lewis, D. Neal, et al. Towards expert-level medical question answering with large language models. *arXiv preprint arXiv:2305.09617*, 2023.
- [272] L. Smith, L. K. Tanabe, R. J. nee Ando, C.-J. Kuo, I.-F. Chung, C.-N. Hsu, Y.-S. Lin, R. Klinger, C. M. Friedrich, K. Ganchev, et al. Overview of biocreative ii gene mention recognition. *Genome biology*, 9(2):1–19, 2008.
- [273] G. Soğancıoğlu, H. Öztürk, and A. Özgür. Biosses: a semantic sentence similarity estimation system for the biomedical domain. *Bioinformatics*, 33(14):i49–i58, 2017.
- [274] J. R. A. Solares, F. E. D. Raimondi, Y. Zhu, F. Rahimian, D. Canoy, J. Tran, A. C. P. Gomes, A. H. Payberah, M. Zottoli, M. Nazarzadeh, et al. Deep learning for electronic health records: a comparative review of multiple deep neural architectures. *Journal of biomedical informatics*, 101:103337, 2020.
- [275] S. Soni and K. Roberts. Evaluation of dataset selection for pre-training and fine-tuning transformer language models for clinical question answering. In *LREC*, pages 5532–5538, 2020.
- [276] P. Spyns. Natural language processing in medicine: an overview. *Methods of information in medicine*, 35(4-5):285–301, 1996.
- [277] P. Sturmfels, J. Vig, A. Madani, and N. F. Rajani. Profile prediction: An alignment-based pre-training task for protein sequence models. *arXiv preprint arXiv:2012.00195*, 2020.
- [278] W. F. Styler, S. Bethard, S. Finan, M. Palmer, S. Pradhan, P. C. De Groen, B. Erickson, T. Miller, C. Lin, G. Savova, et al. Temporal annotation in the clinical domain. *Transactions of the association for computational linguistics*, 2:143–154, 2014.

- [279] D. Su, Y. Xu, T. Yu, F. B. Siddique, E. Barezi, and P. Fung. Caire-covid: A question answering and query-focused multi-document summarization system for covid-19 scholarly information management. In *Proceedings of the 1st Workshop on NLP for COVID-19 (Part 2) at EMNLP 2020*, 2020.
- [280] P. Su, Y. Peng, and K. Vijay-Shanker. Improving bert model using contrastive learning for biomedical relation extraction. In *Proceedings of the 20th Workshop on Biomedical Language Processing*, pages 1–10, 2021.
- [281] P. Su and K. Vijay-Shanker. Investigation of bert model on biomedical relation extraction based on revised fine-tuning mechanism. In *BIBM*, pages 2522–2529. IEEE, 2020.
- [282] S. Subramanian, L. L. Wang, B. Bogin, S. Mehta, M. van Zuylen, S. Parasa, S. Singh, M. Gardner, and H. Hajishirzi. Medicat: A dataset of medical images, captions, and textual references. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2112–2120, 2020.
- [283] D. Sui, Y. Chen, J. Zhao, Y. Jia, Y. Xie, and W. Sun. Feded: federated learning via ensemble distillation for medical relation extraction. In *EMNLP*, pages 2118–2128, 2020.
- [284] C. Sun, X. Qiu, Y. Xu, and X. Huang. How to fine-tune bert for text classification? In *China National Conference on Chinese Computational Linguistics*, pages 194–206. Springer, 2019.
- [285] C. Sun, Z. Yang, L. Wang, Y. Zhang, H. Lin, and J. Wang. Biomedical named entity recognition using bert in the machine reading comprehension framework. *arXiv preprint arXiv:2009.01560*, 2020.
- [286] T. Sun, A. Gaut, S. Tang, Y. Huang, M. ElSherief, J. Zhao, D. Mirza, E. Belding, K.-W. Chang, and W. Y. Wang. Mitigating gender bias in natural language processing: literature review. *arXiv preprint arXiv:1906.08976*, 2019.
- [287] W. Sun, A. Rumshisky, and O. Uzuner. Evaluating temporal relations in clinical text: 2012 i2b2 challenge. *Journal of the American Medical Informatics Association*, 20(5):806–813, 2013.
- [288] H. Tan and M. Bansal. Lxmert: Learning cross-modality encoder representations from transformers. In *EMNLP-IJCNLP*, pages 5100–5111, 2019.
- [289] R. Tang, R. Nogueira, E. Zhang, N. Gupta, P. Cam, K. Cho, and J. Lin. Rapidly bootstrapping a question answering dataset for covid-19. *arXiv preprint arXiv:2004.11339*, 2020.
- [290] Y. Tay, M. Dehghani, D. Bahri, and D. Metzler. Efficient transformers: a survey. *arXiv preprint arXiv:2009.06732*, 2020.
- [291] A. Thillaisundaram and T. Togia. Biomedical relation extraction with pre-trained language representations and minimal task-specific architecture. In *Workshop on BioNLP Open Shared Tasks*, pages 84–89, 2019.
- [292] R. Tinn, H. Cheng, Y. Gu, N. Usuyama, X. Liu, T. Naumann, J. Gao, and H. Poon. Fine-tuning large neural language models for biomedical natural language processing. *arXiv preprint arXiv:2112.07869*, 2021.
- [293] E. Tiu, E. Taluis, P. Patel, C. P. Langlotz, A. Y. Ng, and P. Rajpurkar. Expert-level detection of pathologies from unannotated chest x-ray images via self-supervised learning. *Nature Biomedical Engineering*, pages 1–8, 2022.
- [294] E. F. Tjong Kim Sang and F. De Meulder. Introduction to the CoNLL-2003 shared task: language-independent named entity recognition. In *NAACL-HLT*, pages 142–147, 2003.
- [295] Y. Tong, Y. Chen, and X. Shi. A multi-task approach for improving biomedical named entity recognition by incorporating multi-granularity information. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 4804–4813, 2021.
- [296] H.-L. Trieu, T. T. Tran, K. N. Duong, A. Nguyen, M. Miwa, and S. Ananiadou. Deepeventmine: end-to-end neural nested event extraction from biomedical texts. *Bioinformatics*, 36(19):4910–4917, 2020.
- [297] G. Tsatsaronis, G. Balikas, P. Malakasiotis, I. Partalas, M. Zschunke, M. R. Alvers, D. Weissenborn, A. Krithara, S. Petridis, D. Polychronopoulos, et al. An overview of the bioasq large-scale biomedical semantic indexing and question answering competition. *BMC bioinformatics*, 16(1):1–28, 2015.
- [298] R. Turner, D. Eriksson, M. McCourt, J. Kiili, E. Laaksonen, Z. Xu, and I. Guyon. Bayesian optimization is superior to random search for machine learning hyperparameter tuning: analysis of the black-box optimization challenge 2020. *arXiv preprint arXiv:2104.10201*, 2021.
- [299] E. Tutubalina, I. Alimova, Z. Miftahutdinov, A. Sakhovskiy, V. Malykh, and S. Nikolenko. The russian drug reaction corpus and neural models for drug reactions and effectiveness detection in user reviews. *Bioinformatics*, 37(2):243–249, 2021.
- [300] Ö. Uzuner, B. R. South, S. Shen, and S. L. DuVall. 2010 i2b2/va challenge on concepts, assertions, and relations in clinical text. *Journal of the American Medical Informatics Association*, 18(5):552–556, 2011.
- [301] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. In *NeurIPS*, pages 5998–6008, 2017.
- [302] J. Vig, A. Madani, L. R. Varshney, C. Xiong, R. Socher, and N. F. Rajani. Bertology meets biology: interpreting attention in protein language models, 2021.
- [303] I. Vulić, E. M. Ponti, R. Litschko, G. Glavaš, and A. Korhonen. Probing pretrained language models for lexical semantics. *arXiv preprint arXiv:2010.05731*, 2020.
- [304] S. Wada, T. Takeda, S. Manabe, S. Konishi, J. Kamohara, and Y. Matsumura. Pre-training technique to localize medical bert and enhance biomedical bert. *arXiv preprint arXiv:2005.07202*, 2020.
- [305] D. Wadden, U. Wennberg, Y. Luan, and H. Hajishirzi. Entity, relation, and event extraction with contextualized span representations. In *EMNLP-IJCNLP*, pages 5788–5793, 2019.
- [306] B. C. Wallace, S. Saha, F. Soboczenski, and I. J. Marshall. Generating (factual?) narrative summaries of rcts: experiments with neural multi-document summarization. *arXiv preprint arXiv:2008.11293*, 2020.
- [307] A. Wang, A. Singh, J. Michael, F. Hill, O. Levy, and S. R. Bowman. Glue: a multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*, 2018.

- [308] B. Wang, L. Shang, C. Lioma, X. Jiang, H. Yang, Q. Liu, and J. G. Simonsen. On position embeddings in bert. In *ICLR*, volume 2, pages 12–13, 2021.
- [309] B. Wang, D. Zhao, C. Lioma, Q. Li, P. Zhang, and J. G. Simonsen. Encoding word order in complex embeddings. *ICLR 2020 spotlight*, 2019.
- [310] F. Wang, Y. Zhou, S. Wang, V. Vardhanabhuti, and L. Yu. Multi-granularity cross-modal alignment for generalized medical visual representation learning. In *Advances in Neural Information Processing Systems*, 2022.
- [311] L. L. Wang, K. Lo, Y. Chandrasekhar, R. Reas, J. Yang, D. Burdick, D. Eide, K. Funk, Y. Katsis, R. M. Kinney, et al. Cord-19: the covid-19 open research dataset. In *ACL Workshop on NLP for COVID-19*, 2020.
- [312] X. Wang, Z. Xu, L. Tam, D. Yang, and D. Xu. Self-supervised image-text pre-training with mixed data in chest x-rays. *arXiv preprint arXiv:2103.16022*, 2021.
- [313] X. D. Wang, L. Weber, and U. Leser. Biomedical event extraction as multi-turn question answering. In *ACL Workshop on Health Text Mining and Information Analysis*, pages 88–96, 2020.
- [314] Y. Wang, N. Afzal, S. Fu, L. Wang, F. Shen, M. Rastegar-Mojarad, and H. Liu. Medsts: a resource for clinical semantic textual similarity. *Language Resources and Evaluation*, 54(1):57–72, 2020.
- [315] Y. Wang, S. Fu, F. Shen, S. Henry, O. Uzuner, and H. Liu. The 2019 n2c2/ohnlp track on clinical semantic textual similarity: overview. *JMIR Medical Informatics*, 8(11):e23375, 2020.
- [316] Y. Wang, L. Wang, M. Rastegar-Mojarad, S. Moon, F. Shen, N. Afzal, S. Liu, Y. Zeng, S. Mehrabi, S. Sohn, et al. Clinical information extraction applications: a literature review. *Journal of biomedical informatics*, 77:34–49, 2018.
- [317] Y.-A. Wang and Y.-N. Chen. What do position embeddings learn? an empirical study of pre-trained language model positional encoding. In *EMNLP*, pages 6840–6849, 2020.
- [318] Z. Wang, Z. Wu, D. Agarwal, and J. Sun. Medclip: Contrastive learning from unpaired medical images and text. *arXiv preprint arXiv:2210.10163*, 2022.
- [319] N. Warikoo, Y.-C. Chang, and W.-L. Hsu. Lbert: lexically aware transformer-based bidirectional encoder representation model for learning universal bio-entity relations. *Bioinformatics*, 37(3):404–412, 2021.
- [320] Q. Wei, Z. Ji, Y. Si, J. Du, J. Wang, F. Tiriyaki, S. Wu, C. Tao, K. Roberts, and H. Xu. Relation extraction from clinical narratives using pre-trained language models. In *AMIA Annual Symposium Proceedings*, volume 2019, page 1236. American Medical Informatics Association, 2019.
- [321] Z. Wei, Q. Liu, B. Peng, H. Tou, T. Chen, X.-J. Huang, K.-F. Wong, and X. Dai. Task-oriented dialogue system for automatic diagnosis. In *ACL*, pages 201–207, 2018.
- [322] W.-H. Weng and P. Szolovits. Representation learning for electronic health records. *arXiv preprint arXiv:1909.09248*, 2019.
- [323] T. Whang, D. Lee, C. Lee, K. Yang, D. Oh, and H. Lim. An effective domain adaptive post-training method for bert in response selection. In *INTERSPEECH*, pages 1585–1589, 2020.
- [324] A. Williams, N. Nangia, and S. Bowman. A broad-coverage challenge corpus for sentence understanding through inference. In *NAACL-HLT*, pages 1112–1122. Association for Computational Linguistics, 2018.
- [325] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, J. Davison, S. Shleifer, P. von Platen, C. Ma, Y. Jernite, J. Plu, C. Xu, T. L. Scao, S. Gugger, M. Drame, Q. Lhoest, and A. M. Rush. Huggingface’s transformers: state-of-the-art natural language processing, 2020.
- [326] S. Wu, K. Roberts, S. Datta, J. Du, Z. Ji, Y. Si, S. Soni, Q. Wang, Q. Wei, Y. Xiang, et al. Deep learning in clinical natural language processing: a methodical review. *Journal of the American Medical Informatics Association*, 27(3):457–470, 2020.
- [327] Z. Wu, Y. Chen, B. Kao, and Q. Liu. Perturbed masking: parameter-free probing for analyzing and interpreting bert. In *ACL*, pages 4166–4176, 2020.
- [328] Z. Wu, Y. Song, S. Huang, Y. Tian, and F. Xia. Wtmed at medqa 2019: A hybrid approach to biomedical natural language inference. In *Proceedings of the 18th BioNLP Workshop and Shared Task*, pages 415–426, 2019.
- [329] Y. Xia, J. Zhou, Z. Shi, C. Lu, and H. Huang. Generative adversarial regularized mutual information policy gradient framework for automatic diagnosis. In *AAAI*, volume 34, pages 1062–1069, 2020.
- [330] Y. Xiao, J. Qiu, Z. Li, C.-Y. Hsieh, and J. Tang. Modeling protein using large-scale pretrain language model. *arXiv preprint arXiv:2108.07435*, 2021.
- [331] Q. Xie, J. A. Bishop, P. Tiwari, and S. Ananiadou. Pre-trained language models with domain knowledge for biomedical extractive summarization. *Knowledge-Based Systems*, 252:109460, 2022.
- [332] Q. Xie, J. Hu, J. Zhou, Y. Peng, and F. Wang. Factreranker: Fact-guided reranker for faithful radiology report summarization. *arXiv preprint arXiv:2303.08335*, 2023.
- [333] Q. Xie, J. Huang, T. Saha, and S. Ananiadou. Gretel: Graph contrastive topic enhanced language model for long document extractive summarization. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 6259–6269, 2022.
- [334] Q. Xie, Z. Luo, B. Wang, and S. Ananiadou. A survey on biomedical text summarization with pre-trained language model. *arXiv preprint arXiv:2304.08763*, 2023.
- [335] Q. Xie and F. Wang. Faithful ai in healthcare and medicine. *medRxiv*, pages 2023–04, 2023.
- [336] Z. Xie, Z. Zhang, Y. Cao, Y. Lin, J. Bao, Z. Yao, Q. Dai, and H. Hu. Simmim: A simple framework for masked image modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9653–9663, 2022.
- [337] L. Xu, Q. Zhou, K. Gong, X. Liang, J. Tang, and L. Lin. End-to-end knowledge-routed relational dialogue system for automatic diagnosis. In *AAAI*, volume 33, pages 7346–7353, 2019.

- [338] S. Xu, H. Li, P. Yuan, Y. Wang, Y. Wu, X. He, Y. Liu, and B. Zhou. K-plug: knowledge-injected pre-trained language model for natural language understanding and generation in e-commerce. *arXiv preprint arXiv:2104.06960*, 2021.
- [339] K. Xue, Y. Zhou, Z. Ma, T. Ruan, H. Zhang, and P. He. Fine-tuning bert for joint entity and relation extraction in chinese medical text. In *BIBM*, pages 892–897. IEEE, 2019.
- [340] P. Yadav, M. Steinbach, V. Kumar, and G. Simon. Mining electronic health records (ehrs) a survey. *ACM Computing Surveys (CSUR)*, 50(6):1–40, 2018.
- [341] S. Yadav, V. Pallagani, and A. Sheth. Medical knowledge-enriched textual entailment framework. In *ICCL*, pages 1795–1801, 2020.
- [342] K. Yamada and M. Hamada. Prediction of rna-protein interactions using a nucleotide language model. *bioRxiv*, 2021.
- [343] K. Yamada and M. Hamada. Prediction of rna-protein interactions using a nucleotide language model. *bioRxiv*, 2021.
- [344] G. Yan, J. Pei, P. Ren, Z. Ren, X. Xin, H. Liang, M. de Rijke, and Z. Chen. Remedi: Resources for multi-domain, multi-service, medical dialogues. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 3013–3024, 2022.
- [345] Q. Yang, Y. Liu, T. Chen, and Y. Tong. Federated machine learning: concept and applications. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 10(2):1–19, 2019.
- [346] Q. Yang, Y. Liu, T. Chen, and Y. Tong. Federated machine learning: Concept and applications. *ACM Trans. Intell. Syst. Technol.*, 10(2), jan 2019.
- [347] W. Yang, G. Zeng, B. Tan, Z. Ju, S. Chakravorty, X. He, S. Chen, X. Yang, Q. Wu, Z. Yu, et al. On the generation of medical dialogues for covid-19. *arXiv preprint arXiv:2005.05442*, 2020.
- [348] X. Yang, J. Bian, W. R. Hogan, and Y. Wu. Clinical concept extraction using transformers. *Journal of the American Medical Informatics Association*, 10 2020. ocaa189.
- [349] X. Yang, X. He, H. Zhang, Y. Ma, J. Bian, and Y. Wu. Measurement of semantic textual similarity in clinical texts: comparison of transformer-based models. *JMIR Medical Informatics*, 8(11):e19735, 2020.
- [350] Z. Yang, Z. Dai, Y. Yang, J. Carbonell, R. Salakhutdinov, and Q. V. Le. Xlnet: generalized autoregressive pretraining for language understanding, 2020.
- [351] M. Yasunaga, J. Leskovec, and P. Liang. 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, 2022.
- [352] W. Yoon, R. Jackson, J. Kang, and A. Lagerberg. Sequence tagging for biomedical extractive question answering. *arXiv preprint arXiv:2104.07535*, 2021.
- [353] W. Yoon, J. Lee, D. Kim, M. Jeong, and J. Kang. Pre-trained language model for biomedical question answering. In *ECML PKDD*, pages 727–740. Springer, 2019.
- [354] X. Yu, W. Hu, S. Lu, X. Sun, and Z. Yuan. Biobert based named entity recognition in electronic medical record. In *ITME*, pages 49–52. IEEE, 2019.
- [355] H. Yuan, Z. Yuan, R. Gan, J. Zhang, Y. Xie, and S. Yu. Biobart: Pretraining and evaluation of a biomedical generative language model. *BioNLP 2022@ACL 2022*, page 97, 2022.
- [356] W. Yuan, G. Neubig, and P. Liu. Bartscore: Evaluating generated text as text generation. *Advances in Neural Information Processing Systems*, 34:27263–27277, 2021.
- [357] Z. Yuan, Z. Zhao, and S. Yu. Coder: knowledge infused cross-lingual medical term embedding for term normalization. *arXiv preprint arXiv:2011.02947*, 2020.
- [358] M. Zaheer, G. Guruganesh, A. Dubey, J. Ainslie, C. Alberti, S. Ontanon, P. Pham, A. Ravula, Q. Wang, L. Yang, et al. Big bird: transformers for longer sequences. *arXiv preprint arXiv:2007.14062*, 2020.
- [359] G. Zeng, W. Yang, Z. Ju, Y. Yang, S. Wang, R. Zhang, M. Zhou, J. Zeng, X. Dong, R. Zhang, et al. Meddialog: a large-scale medical dialogue dataset. In *EMNLP*, pages 9241–9250, 2020.
- [360] Z. Zeng, H. Shi, Y. Wu, and Z. Hong. Survey of natural language processing techniques in bioinformatics. *Computational and mathematical methods in medicine*, 2015, 2015.
- [361] H. Zhang, J. Chen, F. Jiang, F. Yu, Z. Chen, J. Li, G. Chen, X. Wu, Z. Zhang, Q. Xiao, X. Wan, B. Wang, and H. Li. Huatuogpt, towards taming language models to be a doctor. *arXiv preprint arXiv:2305.15075*, 2023.
- [362] H. Zhang, A. X. Lu, M. Abdalla, M. McDermott, and M. Ghassemi. Hurtful words: quantifying biases in clinical contextual word embeddings. In *CHIL*, pages 110–120, 2020.
- [363] H. Zhang, X. Wan, and B. Wang. Injecting knowledge into biomedical pre-trained models via polymorphism and synonymous substitution. *arXiv preprint arXiv:2305.15010*, 2023.
- [364] N. Zhang, Z. Bi, X. Liang, L. Li, X. Chen, S. Deng, L. Li, X. Xie, H. Ye, X. Shang, K. Yin, C. Tan, J. Xu, M. Chen, F. Huang, L. Si, Y. Ni, G. Xie, Z. Sui, B. Chang, H. Zong, Z. Yuan, L. Li, J. Yan, H. Zan, K. Zhang, H. Chen, B. Tang, and Q. Chen. Cblue: a chinese biomedical language understanding evaluation benchmark, 2021.
- [365] N. Zhang, Q. Jia, K. Yin, L. Dong, F. Gao, and N. Hua. Conceptualized representation learning for chinese biomedical text mining. *arXiv preprint arXiv:2008.10813*, 2020.
- [366] R. Zhang, R. G. Reddy, M. A. Sultan, V. Castelli, A. Ferritto, R. Florian, E. S. Kayi, S. Roukos, A. Sil, and T. Ward. Multi-stage pre-training for low-resource domain adaptation, 2020.
- [367] S. Zhang, X. Zhang, H. Wang, L. Guo, and S. Liu. Multi-scale attentive interaction networks for chinese medical question answer selection. *IEEE Access*, 6:74061–74071, 2018.

- [368] T. Zhang, V. Kishore, F. Wu, K. Q. Weinberger, and Y. Artzi. Bertscore: Evaluating text generation with bert. In *International Conference on Learning Representations*.
- [369] W. Zhang, L. Hou, Y. Yin, L. Shang, X. Chen, X. Jiang, and Q. Liu. Ternarybert: distillation-aware ultra-low bit bert, 2020.
- [370] Y. Zhang, H. Jiang, Y. Miura, C. D. Manning, and C. P. Langlotz. Contrastive learning of medical visual representations from paired images and text. *arXiv preprint arXiv:2010.00747*, 2020.
- [371] Y. Zhang, Z. Jiang, T. Zhang, S. Liu, J. Cao, K. Liu, S. Liu, and J. Zhao. Mie: a medical information extractor towards medical dialogues. In *ACL*, pages 6460–6469, 2020.
- [372] Y. Zhang, D. Merck, E. Tsai, C. D. Manning, and C. Langlotz. Optimizing the factual correctness of a summary: A study of summarizing radiology reports. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5108–5120, 2020.
- [373] Y. Zhang, S. Sun, M. Galley, Y.-C. Chen, C. Brockett, X. Gao, J. Gao, J. Liu, and W. B. Dolan. Dialogpt: Large-scale generative pre-training for conversational response generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 270–278, 2020.
- [374] J. Zhao, T. Wang, M. Yatskar, V. Ordonez, and K.-W. Chang. Men also like shopping: reducing gender bias amplification using corpus-level constraints. *arXiv preprint arXiv:1707.09457*, 2017.
- [375] H. Zhou, X. Li, W. Yao, C. Lang, and S. Ning. Dut-nlp at mediqa 2019: an adversarial multi-task network to jointly model recognizing question entailment and question answering. In *BioNLP Workshop and Shared Task*, pages 437–445, 2019.
- [376] H. Zhu, I. C. Paschalidis, and A. M. Tahmasebi. Clinical concept extraction with contextual word embedding. In *NeuIPS Workshop on Machine Learning for Health*, 2018.
- [377] W. Zhu, Y. Ni, X. Wang, and G. Xie. Discovering better model architectures for medical query understanding. In *NAACL-HLT*, pages 230–237, 2021.
- [378] Y. Zhu, R. Kiros, R. Zemel, R. Salakhutdinov, R. Urtasun, A. Torralba, and S. Fidler. Aligning books and movies: towards story-like visual explanations by watching movies and reading books. In *ICCV*, pages 19–27, 2015.